A Sum of Squares Theorem for Visibility Complexes and Applications

Pierre Angelier and Michel Pocchiola

G - LIENS - 01 - 1

Rapport de Recherche du Liens
Laboratoire d'Informatique de l'Ecole Normale Supérieure, UMR 8548

LIENS, Département d'Informatique,
École Normale Supérieure, 45 rue d'Ulm, F-75230 Paris Cedex 05, France.

Téléphone : (33) (1) 44 32 20 34 Mél: direction@di.ens.fr Télécopie: (33) (1) 44 32 20 75

A sum of squares theorem for visibility complexes and applications

Pierre Angelier
Michel Pocchiola

Abstract

We present a new method to implement in constant amortized time the flip operation of the so-called Greedy Flip Algorithm, an optimal algorithm to compute the visibility graph or the visibility complex of a collection of pairwise disjoint bounded convex sets of constant complexity. The method uses simple data structures and only the left-turn or counterclockwise predicate; it relies, among other things, on a sum of squares like theorem for visibility complexes stated and proved in this paper. (The sum of squares theorem for a simple arrangement of lines states that the average value of the square of the number of vertices of a face of the arrangement is a $O(1)$.)

1 Introduction

....the examples discussed so far suggested that oriented matroids may help in [research in computational geometry]. They provide a combinatorial model and good ways to organize the combinatorial data, and perhaps they offer a better understanding of what algorithms do, why they work, and how to improve them.
—A. Björner, M. Las Vergnas, B. Sturmfels, N. White and G. M. Ziegler, Oriented Matroids, 1993.

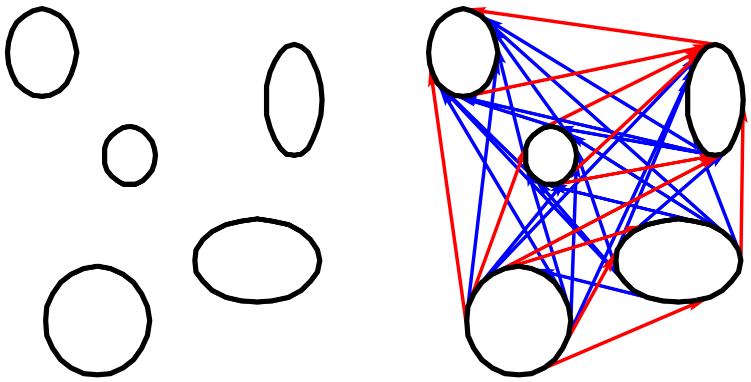

Figure 1: A collection of five convex objects and its set of free bitangents.

This paper is concerned with visibility among convex objects in the plane [32, 5]. A better understanding of this is intimately involved in the practical concerns of computer graphics and motion planning in robotics. An abstract object related to this problem is the visibility complex, a 2D regular cell complex whose underlying topological space is the quotient of the space of rays by the visibility relation : two rays are said visible if they are supported by the same line and if the line segment joining the origins of the two rays lies in free space, the complement of the objects [40]; the one-skeleton of the visibility complex is combinatorially equivalent to the tangent visibility graph; in particular its set of vertices is in one-to-one correspondence with the set of free bitangents. Applications of visibility graphs and visibility complexes are found in path planning [31, 27, 13], network optimization [29, 30], global visibility [38, 44, 11, 14], and

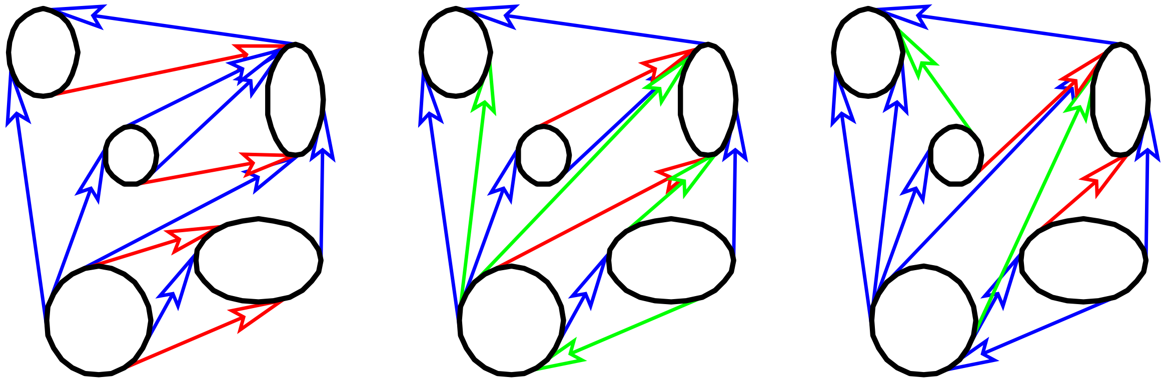

Figure 2: Six steps of the GFA.

pattern recognition [20]; visibility complexes of simple polygons have been elegantly characterized in the framework of oriented matroid theory [1, 7] as a class of bipartite graphs [34, 35, 46]. The related notions of pseudo-triangles and pseudo-triangulations, which have played a key role in the design of efficient visibility graphs and visibility complexes algorithms [40, 39], have applications to ray-shooting [9, 19, 38], visibility graph or visibility complex counting [37, 38, 36], covering and separation [41], collision detection [2, 25, 24, 23], stretchability and realizability of pseudoline arrangements [36, 38, 18], and, last but not least, planning expansive motions of polygonal chains [33, 12, 47]. A key feature of pseudo-triangulations that is largely exploited in the applications mentioned above is the existence of a flip operation e.g., in Figure 2 the rightmost pseudo-triangulation is obtained from the leftmost pseudo-triangulation after a sequence of six flip operations (see also left part of Figure 3); properties of the flip graph of pseudo-triangulations of a given point set are reported in [42, 8, 45]. A challenging question is to extend these notions and/or applications to 3D environments.

Source	Obstacles	Data structures	Predicate
Edelsbrunner-Guibas [16]	points	Simple	left-turn
Pocchiola-Vegter [39]	disks	Splittable queues	χ_2, χ_3
Rivière [43]	line segments	Simple	left-turn
This paper	disks	Simple	left-turn

Table 1: Optimal time and linear space visibility graph or visibility complex algorithms for obstacles with constant complexity.

In this paper we focus on the algorithmic of visibility graphs and visibility complexes : we describe an optimal time and linear space algorithm that uses only simple data structures and, once the extremal points of the disks are sorted with respect to some canonical direction, only the counterclockwise or left-turn predicate χ_1 for disks defined by

$$\chi_1(b_1,b_2) \iff \begin{cases} b_1 \text{ and } b_2 \text{ enter or leave the} \\ \text{same disk on the same side} \\ \text{and } \det(\vec{b}_1,\vec{b}_2) > 0; \end{cases} \tag{1}$$

here $\vec{b} \in \mathbb{S}$ denotes the direction of the bitangent b; note that in the case where the disks are reduced to points the χ_1 predicate coincides with the counterclockwise or left-turn predicate pqr on triplets of points which asserts that the points p, q, r are encountered in cyclic order $p, q, r, p, q \ldots$ when traversing counterclockwise the boundary of the circle through p, q and r. Our technique, specialized to point obstacles or circles with an infinitesimal radius [22], subsumes the topological sweep technique of Edelsbrunner and Guibas [16] and its extensions developed in [6, 43]. This last point will be explained in a different paper. For a bibliography on the algorithmic of visibility graphs we refer to [39, 30, 5].

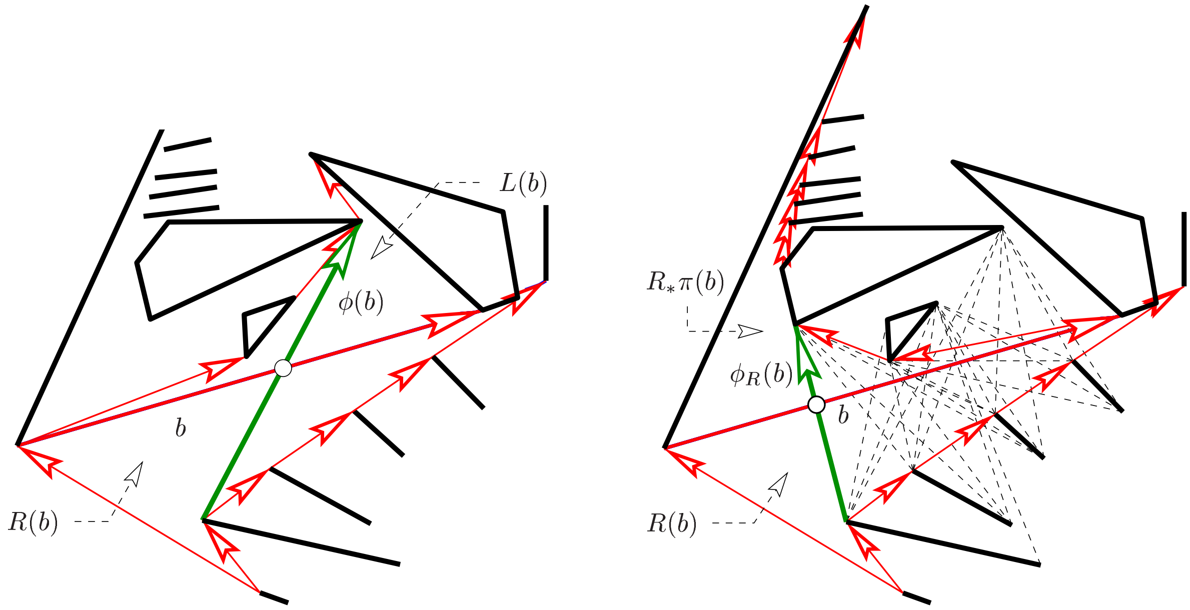

Figure 3: Flipping the bitangent b to obtain $\phi(b)$.

Our starting point is the so-called Greedy Flip Algorithm (GFA for short), an optimal time and linear working space output sensitive algorithm to compute the visibility complex of obstacles of constant complexity like line segments, circles, ellipses, etc [39]. The GFA realizes a topological sweep of the visibility complex guided by some partial order $\prec$ on the set of directed bitangents properly defined, once a canonical direction is chosen, as the transitive closure of the χ_1 relations i.e., $b \prec b'$ if there exists a sequence

$$b_1(=b), b_2, \ldots, b_r(=b') \tag{2}$$

of bitangents such that (1-) $\chi_1(b_i, b_{i+1})$ and (2-) the (counterclockwise) interval of directions $[\vec{b}_i, \vec{b}_{i+1}]$ avoids the canonical direction for all $i = 1, \ldots, r-1$.

The sweep structure S, some sub-complex of the visibility complex, is associated with a pseudo-triangulation $G(S)$ via the bijective map that associates with a 2-cell of the visibility complex the bitangent of its boundary which has the greater angle with respect to the canonical direction. The sweep of a vertex of the visibility complex boils down to some local change in the sweep structure S plus the flip of the corresponding bitangent b in the pseudo-triangulation $G(S)$ i.e., the computation of the bitangent that leaves the pseudo-triangle R of $G(S)$ incident upon b on its right and enters the pseudo-triangle L of $G(S)$ incident upon b on its left. It turns out that the pseudo-triangle R (resp. L) is *greedy* in the sense that considered as the $\prec$-increasing sequence of bitangents that appear in its boundary it is the lexicographic smallest pseudo-triangle incident upon b on its right (resp. left). Therefore R and L depend only on b and so does the bitangent denoted $\phi(b)$ that leaves R and enters L. Alternatively $\phi(b)$ can be defined as the bitangent crossing b from right-to-left whose angle with b is minimal; it is known that the operator ϕ is one-to-one. (See left part of Figure 3.)

Besides linked structures to represent the incidence relations of the atoms of the sweep structure S and its associated pseudo-triangulation $G(S)$ the method described in [39] uses auxiliary data structures to implement efficiently the flip operation: two parts of the boundary of each pseudo-triangle of $G(S)$, called Awake and Asleep, are represented by splittable queues: a data structure for ordered lists that supports the enqueue and dequeue operations either at the head or at the tail of the list, and the split

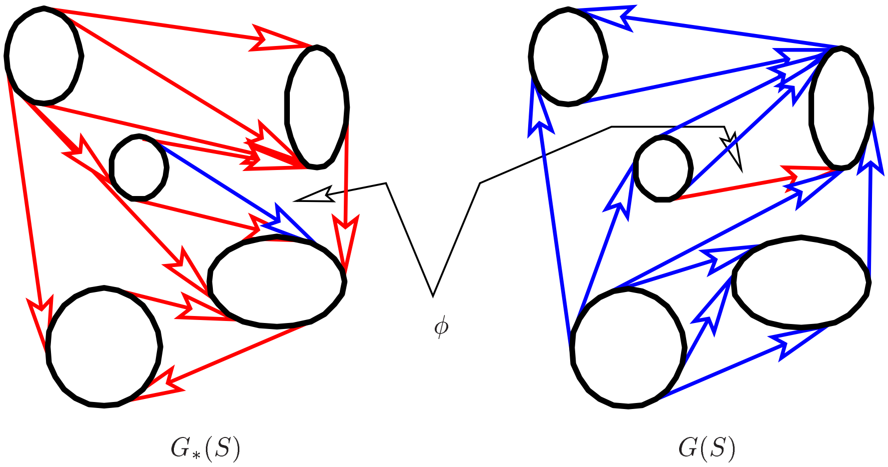

Figure 4: The greedy pseudotriangulation $G(I)$ and its dual $G_*(I)$.

operation at an atom preceded by a search of the atom; the implementation of the splittable queues should guarantee a $O(1)$ amortized time complexity per operation on a sequence of operations (this can be done for example with red-black trees with parent pointers or with randomized search trees); in [39] it is shown that the flip operation requires a constant number of enqueue and split operations and thus is constant amortized; the dequeue and enqueue operations are guided by the evaluation of the geometric predicate $\chi_2(b_1, b_2)$ defined on pairs of directed bitangents not necessarily tangent to the same disk by $\det(\vec{b}_1, \vec{b}_2) > 0$, while the split operations are guided by the evaluation of the predicate $\chi_3(b_1, b_2)$ which states that the tail of b_1, the head of b_1 and the head of b_2 are in counterclockwise order.

Since the visibility complex depends only on the χ_1 relations and not on the χ_2 nor χ_3 relations it is natural to ask whether or not it can be computed using only the left-turn predicate χ_1. A related question—maybe more application oriented—is whether or not the flip operation can be implemented without having recourse to the splittable queue data structure. We answer affirmatively both questions. This result simplifies greatly the implementation of the GFA and should improve its practical performance. The key idea leading to these results is to introduce in the picture the dual partial order $\prec_*$ of $\prec$—defined on the same set of bitangents by $b \prec_* b'$ if $b' \prec b$—, and in particular

1. the dual pseudo-triangulation $G_*(S)$ defined via the bijective map that associates with a 2-cell of the visibility complex the bitangent of its boundary which has the *minimal* angle with respect to the canonical direction; it is known that $G_*(S) = \phi_*(G(S))$ where $\phi_*(b) = \phi^{-1}(b)$ is the bitangent crossing b from left-to-right and whose angle with b is minimal; see Figure 4.

2. the operator ϕ_R defined on the set of bitangents by $\phi_R(b)$ is the bitangent crossing b from right-to-left and whose crossing point with b is the closest to the tail of b; we will show that $\phi_R(b)$ is also the bitangent that leaves $R(b)$ and enters the dual greedy pseudo-triangle $R_*\pi(b)$, —here π is the operator that changes the direction of a bitangent; see the right part of Figure 3.

Thus the main result of the paper can be state as follows.

Theorem 1. *The flip operation of the greedy flip algorithm can be implemented in constant amortized time using only the left-turn predicate for disks. Furthermore the only necessary data structures are (1-) linked structures to represent the atoms and their incidences of the sweep structure S and its associated pseudo-triangulations $G(S)$ and $G_*(S)$, and (2-) a list to store the minimal bitangents of $G(S)$ and for each minimal bitangent t a pointer to the arc of $G(S)$ (resp. $G_*(S)$) that $\phi_R(t)$ leaves (resp. enters) plus similar informations for the left version ϕ_L of ϕ_R formally defined by $\phi_L(t) = \pi\phi_R\pi(t)$ where π is the operator that reverses the direction of a bitangent.* □

Once the disks are sorted with respect to the canonical direction the sweep structure S and its associated pseudo-triangulations $G(S)$ and $G_*(S)$ can be computed in linear time using only the left-turn predicate (see [3, 40]). Thus the GFA can be implemented using only the left-turn predicate modulo some initial sorting.

Some comments on the bit-complexity of the predicates χ_1, χ_2 and χ_3 are in order at this point. If the bit-complexity of the predicates χ_2 and χ_3 for polygonal obstacles are identical (they boil down to the counterclockwise predicate on vertices and difference of vertices of the input polygons) the situation is different in general. To fix the ideas the evaluation of the afore mentioned predicates for circles given by center and radius requires the evaluation of the sign of expressions of the following type

$$a_1 + a_2\sqrt{b_2} + a_3\sqrt{b_3} + a_4\sqrt{b_2 b_3} \tag{3}$$

where the a_i, b_i are polynomials in the inputs whose degrees depend on the predicate (see [3] for the explicit formulas). By a repeated application of the principle that if $a > 0$ and $b < 0$ then the expressions $a + b\sqrt{c}$ and $a^2 - b^2 c$ have the same sign one can show that the evaluation of χ_3 requires in the evaluation of the sign of an irreducible polynomial of degree at most 16 while χ_2 (or χ_1) requires the evaluation of an irreducible polynomial of degree at most 6. Therefore the bit-complexities of χ_2 and χ_3 are $6b + O(1)$ and $16b + O(1)$ where b is the bit-complexity of the inputs (radius and coordinates of the center of the circle) : this is a substantial difference that should reinforce the practical significance of our results.

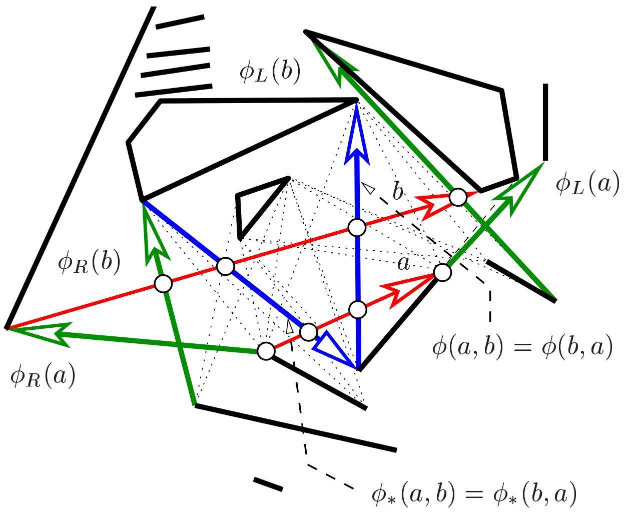

Figure 5: A pair (a, b) such that $\phi(a, b) = \phi(b, a)$ and $\phi_*(a, b) = \phi_*(b, a)$.

The complexity analysis of our flip method is based on a theorem similar to the sum of squares theorem for a simple arrangement of lines in the plane; this last theorem states that the average value of the square of the number of vertices of a face of the arrangement is a $O(1)$: this is a well-known consequence of the linear bound on the complexity of the so-called zone of a line in an arrangement of lines; see [10, 17, 13, 15] and [21, 4] for an higher dimensional analogue. To state our sum of squares theorem we need to introduce new operators. First we extend the definition of the ϕ operator to the case where the set of obstacles is augmented with a set of bitangent obstacles : for H a set of pairwise disjoint bitangents and for b a bitangent which crosses no bitangent in H we denote by $\phi(b; H)$ the bitangent
(1-) crossing no bitangent in H,
(2-) crossing b from right-to-left, and
(3-) whose angle with b is minimal.
If no bitangent crossing b crosses a bitangent in H then $\phi(b; H) = \phi(b)$ and in particular $\phi(b; \emptyset) = \phi(b)$. Finally we introduce the operator

$$\phi(a, b) = \phi(a; \{\phi_R(b), \phi_L(b)\}), \tag{4}$$

which is defined on pairs of bitangents (a, b) such that a crosses nor $\phi_R(b)$ nor $\phi_L(b)$. See Figure 5 for an illustration. We are now ready to state our sum of squares theorem for visibility complexes.

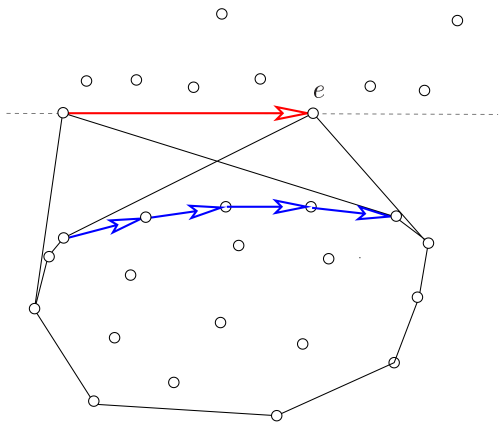

Figure 6: Here $Z(e)$ is composed of 4 edges.

Theorem 2 (Sum of Squares Theorem). *The number of pairs of free bitangents (a, b) such that $\phi(a, b) = \phi(b, a)$ and $\phi_*(a, b) = \phi_*(b, a)$ is a big-O of the number of free bitangents. (Here $\phi_* = \phi^{-1}$.)* $\square$

In the case of point obstacles our theorem can be restate in terms that make no reference to the ϕ and ϕ_* operators and thus are easier to grasp. Let P be a set of n points and let E be the set of $k \leq n(n-1)$ directed edges with endpoints in P that contain no point in P. For e in E we denote by $Z(e)$ the set of edges e' of the convex hull of the points, lying in the right half-plane of the supporting line of e, whose endpoints are visible from each of the endpoints of e. See Figure 6 for an illustration. Then our theorem asserts that

$$\sum_{e \in E} |Z(e)| = O(k). \tag{5}$$

This bound implies a weak form of the sum of squares theorem for simple arrangement of lines : indeed its is no hard to relate $Z(e)$ to the face below the vertex e^* of the arrangement of the dual lines p^* of the points p in P, and thus deduce from the upper bound (5) that the average value of the product of the numbers of vertices of the right and left boundaries of a face of the dual arrangement of the points in P is a $O(1)$; here left and right boundaries are defined with respect to some canonical direction. Finally we mention that the bound (5) is essentially the bound used by Edelsbrunner and Guibas to analyze the complexity of their topological sweep method [16].

The paper is organized as follows. In Section 2 we recall the definition of the visibility complex $\mathbb{V}_H$ of a set of convex obstacles augmented with a set H of bitangent obstacles [40, 38], we introduce notations for the incidence relations of its cells, and we give a characterization of the chains of bitangents and arcs that are boundary of a 2-cell of the visibility complex (Theorem 3); then we relate the ϕ operator to the visibility complex and we we show that this operator satisfies the so-called Minimum Element, Right-to-Left and Flip Properties (Theorems 4 and 5) which were stated in [39] in the case where H is the emptyset. Finally we describe an optimal flip method that uses simple data structures and only the χ_2 predicate. We do not claim originality for the method but for its analysis. We have implemented this method for polygons and circles using the C++ library **CGAL** (binaries are available at the URL http://www.di.ens.fr/˜gecoal). Experimental results are discussed and reported in [3]. Section 3 is devoted to the study of the dual order $\prec_*$, the ϕ_R operator and to the proof of the sum of square theorem. We successively

(1-) identify key properties of the left-turn predicate for disks (Theorem 8);

(2-) describe the greedy pseudo-triangles in terms of the incidence relations of the visibility complex (Theorems 9 and 10) and establish a duality relation between the greedy pseudo-triangles defined with respect to the partial order $\prec$ and the ones defined with respect to its dual partial order $\prec_*$ (Theorem 11);

(3-) give several characterizations of the operator ϕ_R (Theorems 12 and 13), study the equation $\phi(b) = \phi_R(b)$ (Theorem 14), and explain how $\phi_R(b)$ can be efficiently computed in the dynamic setting of the

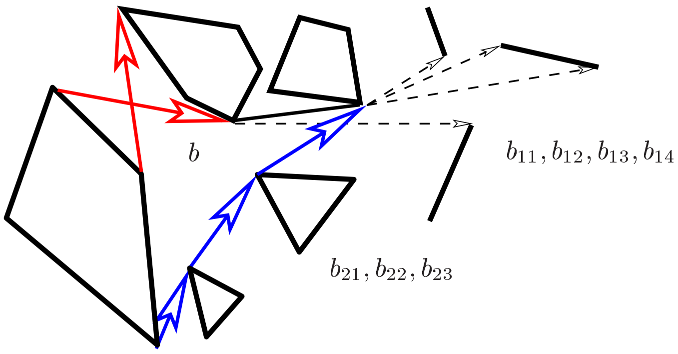

Figure 7: Bitangents, arcs, chains and a face of the visibility complex.

GFA using only the left-turn predicate (Theorems 15 and 16);
(4-) prove the Sum of Squares Theorem and state a sufficient condition, used in our flip method, for a pair (a, b) of bitangents to satisfy the conditions of the Sum of Squares Theorem (Theorem 19).
In the fourth section we describe the optimal flip method based on the left-turn predicate χ_1 and analyze its complexity, thus proving Theorem 1. In the fifth and last section we mention further research suggested by our results.

2 The Greedy Flip Algorithm

2.1 Terminology, pseudo-triangle and pseudo-triangulation

We consider a finite family $O_1, O_2, \ldots, O_n$ of $n \geq 2$ pairwise disjoint bounded closed convex subsets of the Euclidean plane, disks or obstacles for short; for the ease of the exposition we assume that the disks are (1-) smooth (i.e., there is a well-defined tangent line through each boundary point)—in particular the disks are two-dimensional—, and (2-) in generic position (i.e., there is no line tangent to three disks). The general case can be treated using standard symbolic perturbation techniques; for example, to fulfill the smoothness assumption, the obstacles are replaced by their Minkowski sums with an infinitesimally small circle (see [3] for more details). Free space $\mathbb{F}$ is by definition the complement of the interiors of the disks. A *signed* disk is a disk with a sign positive or negative; a directed line is said tangent to a positive (resp. negative) disk if the disk lies on its left (resp. right).

A *bitangent* is a closed directed line segment whose supporting directed line is tangent to two disks at its endpoints and whose interior is disjoint from the two disks (this last condition is added to treat properly the case where the boundaries of the disks contain line segments). Observe that there is exactly one bitangent that leaves a first given signed disk and enters a second given signed disk; a *right-xx* (resp. *xx-right*) bitangent is a bitangent that leaves (resp. enters) a disk on its right i.e., a negative disk; similarly we introduce left-xx and xx-left bitangents. We stress that, thanks to our smoothness and generic position assumptions, two bitangents are disjoint or intersect transversally (i.e., not at their endpoints).

An *arc* is a signed open connected piece of the boundary of a disk whose endpoints are endpoints of free bitangents. Walking in counterclockwise order around the boundary of a disk we traverse an arc from its tail to its head. The source (resp. sink) of an arc is the bitangent whose supporting line is tangent to the supporting signed disk of the arc at its tail (resp. head). The operator that reverses the direction of a bitangent or changes the sign of an arc is denoted π. We denote by B the set of free bitangents. A chain in B is a sequence $s_1 s_2 \cdots s_m$ of alternate bitangents and arcs such that the bitangent that precedes (resp. follows) an arc is the source (resp. sink) of this arc, and two consecutive arcs lie on different disks.

A pseudo-triangulation is a maximal subset of pairwise disjoint free bitangents; it is known that a pseudo-triangulation contains $3n-3$ bitangents, that subdivide free space into simply connected regions, called pseudo-triangles; the boundary of a pseudo-triangle consists of a simple closed curve with three cusp points whose convex hull contains the pseudo-triangle. Following [40] we observe that two tangent lines to a pseudo-triangle crosses inside the pseudo-triangle from which we deduce, using the intermediate value theorem, that given two interior disjoint pseudo-triangles T and T' there exists a unique bitangent that leaves tangentially the boundary of T and enters tangentially the boundary of T'. We will use the following variation of this observation whose proof is left to the reader.

Lemma 1. *Let T and T' be two pseudo-triangles adjacent along a common bitangent b and non overlapping locally around b. Then there exists a unique bitangent that leaves T and enters T'; furthermore this bitangent is free and pierces b. This bitangent is denoted $\phi(T, T')$.* □

2.2 The visibility complex

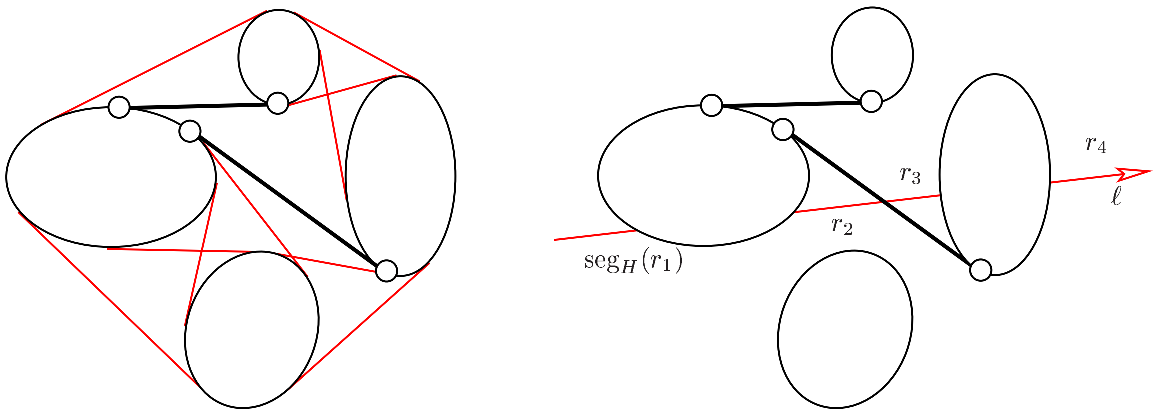

Figure 8: Here H is reduced to two bitangents and B_H has 24 elements; there are 5 canonical chains one per cusp point and one for the lowest disk. The pre-image of the line l under the mapping ω consists of 4 rays in $\mathbb{V}_H$ denoted r_1, r_2, r_3, r_4 represented by the line segments $\mathrm{seg}_H(r_i)$.

We recall now concisely the definition of the visibility complex (we refer to [40, 39] for a smoother introduction) and we introduce notations for the incidence relations of its cells. Let H be a subset of pairwise disjoint or non-crossing free bitangents and let H' be the set H augmented with the hull-bitangents; B_H denotes the set of free bitangents that cross no bitangent in H augmented with the bitangents in $H' \cup \pi(H')$, and A_H denotes the set of arcs induced by the endpoints of the bitangents in B_H. Let $\mathbb{F}_H$ be the topological space obtained by cutting $\mathbb{F}$ along the bitangents in H (a point lying on a bitangent in H appears twice in $\mathbb{F}_H$) and let J_H be the set of cusp points of the boundary of $\mathbb{F}_H$. The cusp points in J_H subdivide the boundary of $\mathbb{F}_H$ into convex chains of arcs and bitangents in H. These chains are called the *canonical chains* associated with H. The visibility complex (associated with H) is a two dimensional cell complex whose underlying topological space $\mathbb{V}_H$ is the quotient space of the space of rays $r = (o(r), \vec{r}) \in \mathbb{F}_H \times \mathbb{S}^1$ under the transitive and the topological closures of the equivalence relation $\sim_H$ defined on Interior$(\mathbb{F}_H) \times \mathbb{S}^1$ as follows: $r \sim_H r'$ iff. (1-) r and r' are supported by the same directed line $\omega(r)$, and (2-) the line segment in the plane joining the origins of r and r' is free and crosses transversally no bitangent in H. We denote by $\mathrm{seg}_H(r)$ the (directed) line segment in $\mathbb{F}$ joining the origins $o(r')$ of the rays r' in the equivalence class of r; the *backward (resp. forward) view* of the ray r is the canonical chain that contains the source (resp. target) point of $\mathrm{seg}_H(r)$; we denote by γ_{+i} (resp. γ_{-i}) the curve of rays in $\mathbb{V}_H$ that leave tangentially the disk O_i on their left (resp. right) and by γ_+ (resp. γ_-) the curve of rays that leaves tangentially the convex hull of the disks on their left (resp. right); we denote by γ_p the curve of rays in $\mathbb{V}_H$ that emanate from the cusp point p; these curves are oriented counterclockwise; The map that associates with $r \in \mathbb{V}_H$ the common supporting line $\omega(r)$ of its rays is a kind of branched covering space: the pre-image $\omega^{-1}(l)$ of the line l consists of a finite set of rays $r_1, \ldots, r_s \in \mathbb{V}_H$ supported by l such that the $\mathrm{seg}_H(r_i)$ are the projections in $\mathbb{F}$ of the connected components on $\mathbb{F}_{H(l)}$ of the line l — here $H(l)$ is the set of line segments of H crossed transversally by

l. Therefore $\mathbb{V}_H$ is locally a two-dimensional set except on the curves γ_i and γ_p; these curves induce a natural structure of regular two-dimensional cell complex on $\mathbb{V}_H$:

1. its 0-cells (or vertices) are the intersection points of the curves γ_i. The map that associates with the vertex v the bitangent $\varepsilon(v) \in B_H$ included in the line segment $\mathrm{seg}_H(v)$ is one-to-one. In the sequel we will speak freely of the tail, head, etc of v for the tail, head, etc of $\varepsilon(v)$.

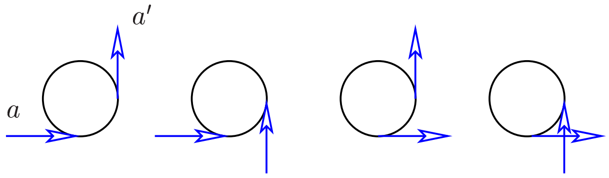

Figure 9: The four types of edges.

2. its 1-cells (or edges) are the connected components of the curves γ_i, γ_p minus the vertices. The map that associates with the edge $e \subset \gamma_i$ the arc $\varepsilon(e) \in A_H$ defined as the envelope of the supporting directed lines of the rays in e is one-to-one. The map that associates with the edge $e \subset \gamma_p$ the cusp point p is two-to-one. An edge $e \subset \gamma_i$ with source a and sink a' is said to be a 0-,1-, or 2-cusp edge depending on the number 0,1 or 2 of cusp points of the curve $\varepsilon(a)\varepsilon(e)\varepsilon(a')$; a source (resp. sink) cusp edge is a 1-cusp edge whose cusp point is realized by the source (resp.sink) of the edge. An edge is said *positive (resp. negative)* if it lies on some γ_{+i} (resp. γ_{-i}). A *chain* in $\mathbb{V}_H$ is a sequence $\sigma_1 \cdots \sigma_m$ of alternate vertices and edges such that the vertex that precedes (resp. follows) an edge is the sink (resp. source) of this edge. The *envelope* $\epsilon(c)$ of the chain $c = \sigma_1\sigma_2 \dots \sigma_k$ is the curve in the plane obtained by concatenation of the curves τ_i where τ_i is defined to be the arc $\epsilon(\sigma_i)$ if σ_i is an edge, the bitangent $\epsilon(\sigma_i)$ if σ_i is a vertex whose adjacent edges σ_{i-1} and σ_{i+1} are not supported by the same disk, or the empty curve otherwise; in this later case the arcs $\epsilon(\sigma_{i-1})$ and $\epsilon(\sigma_{i+1})$ share an endpoint which is either the source point of $\epsilon(\sigma_i)$ — and in that case σ_i is said to leave the chain c — or its sink point — and in that case σ_i is said to enter the chain c. Conversely a curve γ in the plane composed of bitangents in B_H and arcs is the envelope of a unique chain $\epsilon^{-1}(\gamma)$ in $\mathbb{V}(H)$ called its lift in $\mathbb{V}(H)$. A chain is said *convex* if its envelope is a convex curve.

3. its 2-cells (or faces) are the connected components of $\mathbb{V}_H$ minus the vertices and edges; we denote by $\sigma_{+\mathrm{hull}}$ (resp. $s\sigma_{-\mathrm{hull}}$) the 2-cells of rays whose supporting lines let the disks on their left (resp. right). Note that the faces are exactly the interiors of the connected sets of rays with the same backward and forward views. The boundary of a 2-cell $\sigma \neq \sigma_{\pm\mathrm{hull}}$ is composed of two chains with the same source and the same sink, called the source and the sink of the 2-cell. These chains are denoted $rc(\sigma)$ and $lc(\sigma)$ with the convention that in the vicinity of $\mathrm{sour}(\sigma)$ the supporting line of a ray $r \in rc(\sigma)$ lies on the right side of the supporting line of the corresponding parallel ray $r' \in lc(\sigma)$.

An edge $e \subset \gamma_i$ is incident to three 2-cells denoted σ_{righ} or σ_{left} depending on whether i is positive or negative, σ_{forw}, and σ_{back} as indicated in Figure 10. An edge $e \subset \gamma_p$ is incident to one 2-cell denoted σ_{righ} or σ_{left}. A vertex v is incident to six or eight 2-cells depending on whether the bitangent $b := \varepsilon(v)$ does not belong to H or belongs to H. These 2-cells are defined as follows (see Figure 10 for an illustration) :

1. σ_{righ} (resp. σ_{left}) is the 2-cell defined by a ray with direction the direction of v and origin a translated of the middle point of b by a sufficiently small vector orthogonal to b pointing on the right (resp. left) side of the supporting line of v; similar constructions done with the segments $\mathrm{seg}_{\mathrm{back}}, \mathrm{seg}_{\mathrm{forw}}$ left after removing from $\mathrm{seg}_H(v)$ the bitangent b, where $\mathrm{seg}_{\mathrm{back}}$ comes before $\mathrm{seg}_{\mathrm{forw}}$ along $\mathrm{seg}_H(r)$ give rise to two new cells denoted σ_{back} and σ_{forw}.

2. For $b \notin H$ the ray with origin the middle point of b and with direction a direction slightly larger (resp. smaller) then the direction of v defines a 2-cell, denoted σ_{sour} (resp. σ_{sink}) whose source (resp. sink) is v.

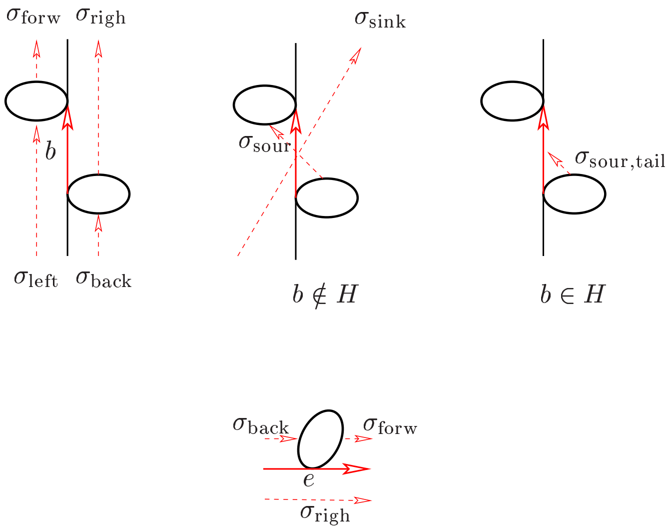

Figure 10: Topology around a vertex and an edge of the visibility complex.

3. For $b \in H$ similar constructions give rise to two 2-cells whose source (resp. sink) is v. For example if b leaves a disk on its right (resp. left) at point x — the tail of b — then the set of rays r with *positive* direction with respect to v (i.e., $\det(\vec{r}, \vec{v}) > 0\}$) and with origin $o(r) \neq x$ in a sufficiently small neighborhood of x on the right (resp. left) of the supporting line of v and on the right of the supporting line of $(x, \vec{r})$ define a 2-face whose source is v. A similar construction done with the head of b instead of its tail gives rise to a second face with source v. Finally similar constructions done with the negative directions with respect to v instead of the positive directions give rise to two faces with sink v. These four faces are denoted $\sigma_{x,y}$ with $x \in \{\text{sour}, \text{sink}\}$ and $y \in \{\text{tail}, \text{head}\}$ with the obvious conventions.

For v a vertex or an edge and xxx $\in \{\text{back}, \text{forw}, \text{righ}, \text{left}\}$ the sink of the cell $\sigma_{\text{xxx}}(c)$ (if defined) is denoted $\phi_{\text{xxx}}(c; H)$. The sink of the cell(s) which source b is denoted $\phi_{\text{sink}}(b; H)$. (Note that if $b \in H$ $\phi_{\text{sink}}(b; H) = \pi(b)$.) We close this subsection with a characterization of chains which are boundary of a 2-face of the visibility complex. The proof is left to the reader.

Theorem 3 (Face Description). *Let σ be a 2-face of $\mathbb{V}(H)$ with source not in H'. Then its right boundary $rc(\sigma)$ is decomposed into three non empty convex subchains*

$$rc(\sigma) = rc_1(\sigma) rc_2(\sigma) rc_3(\sigma) \tag{6}$$

whose elements $a \neq sour(\sigma), sink(\sigma)$ are characterized by $\phi_{back}(a; H)$, $\phi_{left}(a; H)$ and $\phi_{forw}(a; H) = sink(\sigma)$, respectively. In particular the bitangents of the envelopes Σ_i of the $rc_i(\sigma)$ satisfy the three following properties (1-) the bitangents but the source (resp. sink) of Σ_1 (resp. Σ_3) are in H, (2-) there are no bitangents entering (resp. leaving) Σ_1 (resp. Σ_3), and (3-) there are no bitangents entering or leaving Σ_2. Conversely any chain Σ of bitangents and arcs that can be decomposed into three convex subchains $\Sigma_1\Sigma_2\Sigma_3$ that satisfy the three properties above is the envelope of the right boundary of a face whose source and sink are the source and the sink of Σ, respectively. □

Remark 1. *A similar theorem can be stated in the case where the source b of σ is in H or is a hull-bitangent. For example if σ is the 2-face associated with the tail of b its left boundary is reduced to the sequence $b, e_{tail}(b), \pi(b)$ where $e_{tail}(b)$ is the open curve of rays r with origin the tail of b and positive directions with respect to b, and its right boundary $rc(\sigma)$ is decomposed into four convex subchains $rc_1(b)rc_2(b)rc_3(b)rc_4(b)$ where $rc_4(b) = \pi(b)$ is the sink of σ. The chains rc_1, rc_2 and rc_3 satisfy the three properties mentioned in the above theorem. We will use freely this extension without more details.*

2.3 The minimum element, right-to-left, flip and minimal elements properties

For technical reasons it is now preferable to work in the cover space $\mathbb{W}_H$ of $\mathbb{V}_H$ induced by the covering map $p : \mathbb{F}_H \times \mathbb{R} \to \mathbb{F}_H \times \mathbb{S}^1$ that associates with $x = (o(x), \theta(x))$ the ray $p(x)$ with origin o and direction $(\cos\theta, \sin\theta)$; the real $\theta(x)$ is called the angle of x; we denote by π the operator that increases the angle of a ray by π. So $\mathbb{W}_H$ is the quotient space of $\mathbb{F}_H \times \mathbb{R}$ by the equivalence relation $\sim_H^1$ defined by : $x \sim_H^1 x'$ if $\theta(x) = \theta(x')$ and $p(x) \sim_H p(x')$. Elements of $\mathbb{F}_H \times \mathbb{R}, \mathbb{W}$ are still called rays. $X_0(H), X_1(H)$ and $X_2(H)$ denote the sets of 0-,1- and 2-cells of $\mathbb{W}_H$. We identify the vertices of $\mathbb{V}_H$ with the vertices of $\mathbb{W}_H$ with angles in the interval $[0, 2\pi[$. We denote by $\prec_H$ the partial order defined on the 1-skeleton of $\mathbb{W}_H$ as the transitive closure of the relation defined on pair of rays tangent to a same disk by $x \prec_H x'$ if $\theta(x) < \theta(x')$. It is no hard to see that two crossing bitangents in $X_0(H)$ are comparable with respect to $\prec_H$. The design and correctness of the GFA is based on the three following theorems stated and proved in [39] in the case where H is the empty set; up to some minor modifications the proofs are also valid in the case $H \neq \emptyset$ and are not repeated here.

Theo. & Def. 4 (Minimum Element Property). *The map that associates with $b \in X_0(H)$ the minimum bitangent $\phi(b; H)$ in the set of $t \in X_0(H)$ such that t crosses b and $b \prec_H t$ is well defined and one-to-one. Furthermore if b is a non-hull bitangent then $\phi(b; H) = \phi_{sink}(b; H)$ ($\pi(b)$ otherwise), and if $H_1 \subset H_2$ then $\phi(b; H_1) \preceq_{H_2} \phi(b; H_2)$.* □

Recall that a filter I of a poset $(P, \preceq)$ is a subset of P such that if $x \in I$ and $x \preceq y$, then $y \in I$. The sets $I(\theta)$ of bitangents with angles $\geq \theta$ are examples of filters of the poset $(X_0(H), \prec_H)$. Each filter I of $(X_0(H), \prec_H)$ is associated with a greedy pseudo-triangulation $G(I) \subset I$ and a planar sub-complex $S(I)$ of $X(H)$ — the sweep structure S mentioned in the introduction. The pseudo-triangulation

$$G(I) = \{b_1, b_2, \ldots, b_{3n-3}\} \subset I \tag{7}$$

is defined as follows: (1-) b_1 is minimal in I, and (2-) b_{i+1} is minimal in the subset of bitangents in I crossing none of the bitangents $b_1, b_2, \ldots, b_i$; since crossing bitangents are comparable, $G(I)$ is well-defined and is a supset of $\min_{\prec_H} I$. The planar sub-complex $S(I)$ is defined as the set of 1- and 2-cells σ of $X(H)$ such that $\mathrm{sink}(\sigma) \in I$ but not $\mathrm{sour}(\sigma)$ i.e.,

$$S(I) = \{\, \sigma \in X_{\geq 1}(H) \mid \mathrm{sink}(\sigma) \in I,\, \mathrm{sour}(\sigma) \notin I \,\}; \tag{8}$$

$S(I)$ contains one edge $e \subset \gamma$ and its incident faces per curve $\gamma \in \{\gamma_i, \gamma_p\}$, and for each face $\sigma \in S(I)$ an edge $\sigma^+ \in lc(\sigma)$ and an edge $\sigma^- \in rc(\sigma)$; for σ a face in $S(I)$ the bitangent $b := \mathrm{sink}(\sigma)$ is said to be right (resp. left) minimal in I if $b = \mathrm{sink}(\sigma^-)$ (resp. $b = \mathrm{sink}(\sigma^+)$).

Theorem 5 (Flip, Right-to-left and Minimal Elements Properties). *Let I be a filter of $X_0(H)$, let $b \in G(I)$ and let $t \in I$ crossing b. Then (1-) $G(I) = I \setminus \phi(I; H) \subset sink(S(I))$, (2-) $b \prec_H t$, and (3-) b is minimal in I iff. b is right and left minimal in I.* □

The bitangents $\varepsilon(b)$ with $b \in G(I)$ induce a subdivision of free space whose bounded faces are pseudo-triangles. For $b \in G(I)$ we denote by $R(b, I)$ the unique chain in the one-skeleton of $\mathbb{W}_{G(I)}$ whose envelope is the boundary of the pseudo-triangle T incident upon $\varepsilon(b)$ and—locally—to the right of $\varepsilon(b)$ and whose source is the minimal bitangent a in $G(I)$ with $\varepsilon(a)$ in the boundary of T. Note that since $G(I)$ is a pseudo-triangulation there is no bitangent in $\mathbb{W}_{G(I)}$ leaving or entering the chain $R(b, I)$. By a slight abuse of language $R(b, I)$ is called a pseudo-triangle in $G(I)$. From the flip property we deduce that if b is right minimal in I then $\phi(b; H)$ leaves $R(b, I)$; in that case $R(b, I)$ is independent of I and is simply denoted $R(b; H)$; walking along the chain $R(b; H)$, starting from b, we find successively the convex subchains R^1, R^2, R^3, and R^4 of $R(b; H)$. We introduce the left version $L(b, I) := \pi^{-1} R\pi(b, I)$ of $R(b, I)$ and we omit the mention of H when $H = \emptyset$.

Remark 2. *An interesting consequence of the Flip Property is that the set of subsets of pairwise non-crossing bitangents ordered by inclusion is an abstract polytope [28, 8]. This abstract polytope is realizable in some $\mathbb{R}^d$ as a perturbation of the cone of expansive motions of some point set in the plane [45].*

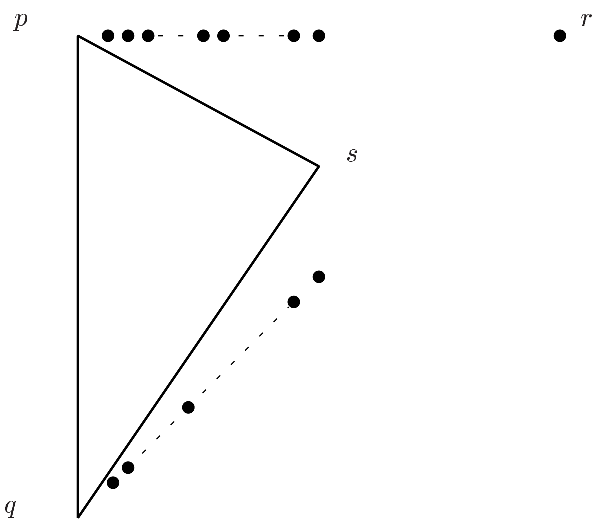

Figure 11: In the configuration $\{(i,0),(i,i-n)\}_{1\leq i\leq n/2-1}$ augmented with the point $r=(n,0)$ and the triangle joining the three points $p=(0,0), q=(0,-n)$ and $s=(n/2-1,-n/4-1/2)$ we have $\sum_b |R^{2_1}(b)| = \Omega(n^2)$ but $|B| = O(n)$.

2.4 The Greedy Flip Algorithm (GFA)

The GFA maintains a representation of the sweep structure $S(I)$ and its associated greedy pseudo-triangulation $G(I)$ when the filter I of $(X_0,\prec)$ describes a maximal chain of filters in the interval $[I_0, I_\pi]$, — here I_x is the filter of bitangents with angle $\geq x$. We use a linked structure to represent the atoms $S(I)$ and $G(I)$ and their incidence relations. Furthermore an atom of $S(I)$ points to its corresponding bitangent in $G(I)$. Finally we access to the linked structure via the set of 'minimal' faces of $S(I)$, i.e. the faces whose corresponding bitangents in $G(I)$ are minimal in the filter I. It should be clear that the update of the data structure when a minimal bitangent b in I is deleted from I can be done in constant time provided we have access to the arcs of G that $b' := \phi(b)$ leaves and enters. We explain now how these two arcs can be computed in constant amortized time using only the predicate χ_2. A crucial consequence of the Right-to-Left Property is that a bitangent (in the range $[b,\pi(b)]$) entering or leaving $R(b)$ can only enter or leave some specific subchain(s) of $R(b)$; in particular b' leaves either the chain $R^2(b)$ or the first atom of $R^3(b)$. This motivates the following definitions.

Theo. & Def. 6. *Let $R^{lea}(b)$ be the chain $R^2(b)$ augmented with the first atom of $R^3(b)$ if this atom is an arc and let $R^{ent}(b)$ be the chain $R^3(b)R^4(b)$. We claim that (1-) $R^4(b)$ is an arc, (2-) any bitangent that enters $R(b)$ enters $R^{ent}(b)$, (3-) any bitangent that leaves $R(b)$ leaves $R^{lea}(b)$ or $R^4(b)$; in particular there is no bitangent entering or leaving the chain $R^1(b)$ and the bitangent $\phi(b)$ leaves the chain $R^{lea}(b)$.*

The tail of $\phi(b)$ cuts the chain $R^{lea}(b)$ into two subchains denoted $R^{lea_1}(b)$ and $R^{lea_2}(b)$ where $R^{lea_1}(b)$ comes before $R^{lea_2}(b)$ when walking along $R^{lea}(b)$. Similarly the tail of $\phi(b)$ cuts the chain $R^i(b)$ where $i=2,3$ into two subchains $R^{i_1}(b)$ and $R^{i_2}(b)$ — this last one is empty if $\phi(b)$ does not leave $R^i(b)$. □

So since b' leaves $R^{\text{lea}}(b)$ and enters $L^{\text{ent}}(b)$ — the left version of $R^{\text{lea}}(b)$ that is formally defined as the chain $\pi^{-1}R^{\text{lea}}\pi(b)$ — a natural idea is to merge according to the χ_2 relations the increasing sequences of bitangents of the chains $R^{\text{lea}}(b)$ and $L^{\text{ent}}(b)$ until we detect the bitangent b'. This leads to the following piece of code where $\text{bit}(r,l)$ is the bitangent leaving (resp. entering) the disk supporting the arc r (resp. l) and where r^+ is the sink of r.

Function FLIP2(b : BITANGENT) : PAIR OF ARCS
/* returns the arcs of $G(I)$ that $\phi(b)$ leaves and enters */
1 $C \leftarrow R^{\text{lea}}(b)$; $C' \leftarrow L^{\text{ent}}(b)$;
2 χ_2-WALK(C, C');

3 **return** $(\mathrm{Top}(C), \mathrm{Top}(C'))$

Procedure χ_2-WALK(C, C': LIST OF ARCS)
0 **while** $C \neq \emptyset$ and $C' \neq \emptyset$
1 $r \leftarrow \mathrm{Top}(C)$; $l \leftarrow \mathrm{Top}(C')$; $\phi \leftarrow \mathrm{bit}(r, l)$;
2 **if** (ϕ leaves r and enters l) **exit**;
3 **if** $\chi_2(r^+, l^+)$ $C \leftarrow \mathrm{Pop}(C)$ **else** $C' \leftarrow \mathrm{Pop}(C')$;

The complexity of FLIP2 is clearly proportional to the sum of the sizes of the chains $R^{2_1}(b)$ and $L^{2_1}(b)(:= \pi^{-1}R^{2_1}\pi(b))$. Since a bitangent is incident to a constant (6) number of faces we can delete from $R^{2_1}(b)$, for the purpose of the complexity analysis, its first bitangent (if any) and the bitangents of $\sigma_{\mathrm{sour}}(b)$. Let us call $R'(b)$ the resulting set of bitangents. As illustrated in Figure 11 the average value of the size of $R'(b)$ on the whole set B of free bitangents is not constant. However, thanks to the sum of squares theorem, a better bound is obtained if we restrict our attention to the set of bitangents whose images by ϕ are right-xx bitangents.

Theo. & Def. 7. *Let $R'(b)$ be the subset of bitangents of the chain $R^{2_1}(b)$ not in the boundary of the 2-cell $\sigma_{sour}(b)$ and from which we delete the minimal element. Then*

$$\sum_{b:\phi(b)\in B_{rx}} |R'(b)| = O(|B|) \tag{9}$$

where B_{rx} is the set of bitangents that leave a disk on their right. □

Thanks to this bound it is now not hard to modify FLIP2 to obtain an optimal algorithm : instead of running χ_2-WALK on the pair $(R^{\mathrm{lea}}, L^{\mathrm{ent}})$ we run it in tandem or synchronously on the four pairs (R^2, L^2), (aR^3, L^2), (R^2, aL^3) and (aR^3, aL^3), where $aR^3(b)$ is the list reduced to the first atom of R^3 or the empty list depending on whether this atom is an arc or not.

Remark 3. *It also follows from our analysis the weaker bound $\sum_{b\in B} |R'(b)| = O(n^2)$ which shows that the procedure* FLIP2 *is worst-case optimal i.e., optimal when the visibility complex has quadratic complexity.*

Remark 4. *Since the lists $aR^3(b)$ and $aL^3(b)$ have zero or one element the χ_2 test at line 3 of the procedure χ_2-*WALK *is useless when running the procedure on the three pairs (aR^3, L^2), (R^2, aL^3) and (aR^3, aL^3). In these three cases the procedure* FLIP2 *is easily modified to use only the χ_1 predicate. This observation will be used in the design of our flip method based only on χ_1.*

Remark 5. *The above considerations are also valid in the case H non empty. In that case $R^{lea}(b)$ is defined as the chain $R^2(b; H)$ augmented with the maximal prefix subchain of $R^3(b; H)$ that contain only bitangents in H and $R'(b; H)$ is defined as the set of bitangents not in H of $R^{2_1}(b; H)$ that are not in the boundary of $\sigma_{sour}(b)$. Thus the set B_H is computable in optimal time $O(\#B_H + n \log n)$ under the assumption that the canonical chains induced by H have constant complexities. In particular the classical visibility graph of a polygon with holes with a total of n vertices is computable in optimal $O(k + n \log n)$ time and linear space.*

3 Technical developments

3.1 The left-turn or counterclockwise predicate for disks

Before going further in our analysis of the greedy flip algorithm we identify properties of the counter-clockwise or left-turn predicate on a collection of 3 disks that we will use in the next subsection to prove the correction of the flip method based only on the left-turn predicate. These properties can also be

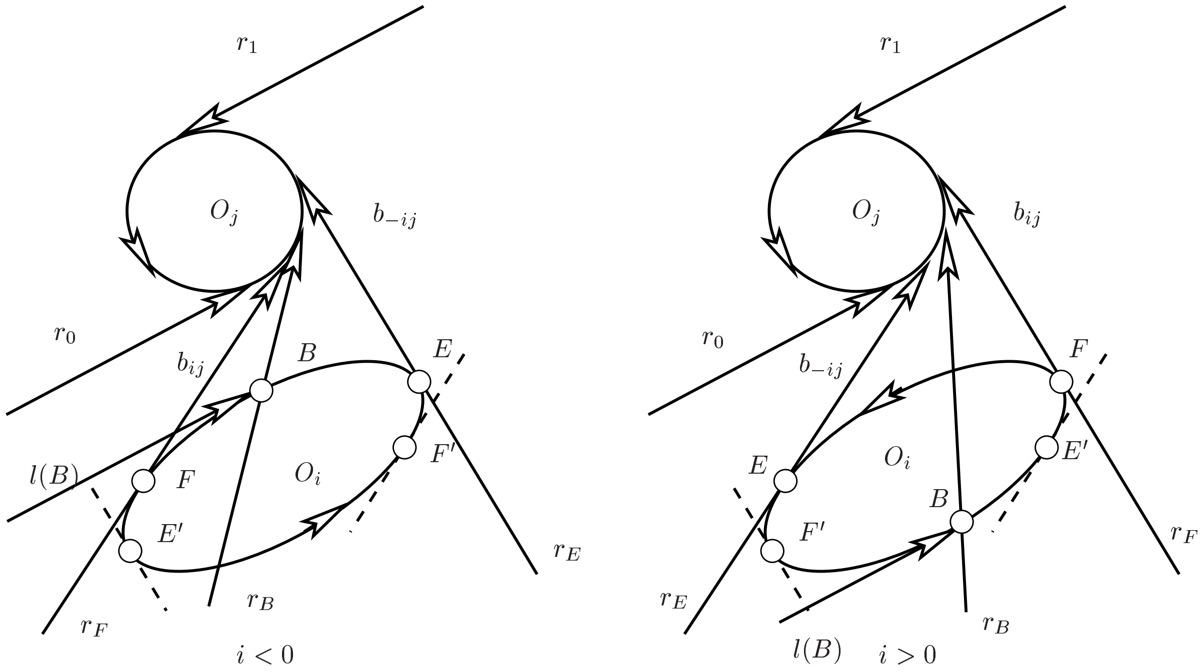

Figure 12: Illustrate Lemma 2.

considered as a first step toward a comprehensive study of the left-turn predicate in the spirit of the Knuth monograph on axioms and hulls [26]. A sequence of bitangents $t_1 t_2 \dots t_p$ such that $\chi_1(t_i, t_{i+1})$, and $\chi_1(t_p, \pi(t_1))$ (equivalent to $\chi_1(t_1, t_p)$) is called a χ_1-sequence. Let O_{+i} (resp. O_{-i}) be the positive (resp. negative) disk O_i. The bitangent that leaves the signed disk O_i and enters the signed disk O_j is denoted b_{ij}. Note that $\pi(b_{ij}) = b_{-j-i}$. We start with a simple geometric lemma whose proof is left to the reader.

Lemma 2. *Let O_i and O_j be two signed disks oriented counterclockwise, let B be a point on the boundary of O_i, and $l(B)$ the half-line that enters O_i at point B. We denote by r_0 (resp. r_1) the half-line that enters O_j and whose direction is the (resp. reverse of the) direction of $l(B)$. Let E (resp. F) be the tail of the bitangent b_{-ij} (resp. b_{ij}). For X a point we denote by r_X the half-line through X that enters O_j and for X on the boundary of O_i we denote by X' the point of tangency of the tangent to O_i whose direction is the reverse of the direction of the tangent to O_i through X. We claim that the half-lines r_0, r_B and r_1 enter in this order the disk O_j iff. B belongs to the arc EF iff. the half-line $l(B)$ lies on the left side of the half-line r_B. Furthermore for if O_i is a negative (resp. positive) disk then*

1. *the half-line r_X pierces O_i iff. r_F, r_X, r_E (resp. r_E, r_X, r_F) enter in this order the disk O_j*
2. *the half-line r_0 pierces O_i iff. B belongs to the arc FE' (resp. $E'F$).*
3. *the half-line r_1 pierces O_i iff. B belongs to the arc $F'E$ (resp. EF')*

In particular if $B \in EF$ and O_i is a negative disk then r_0 and r_1 do not pierce O_i. □

Theorem 8. *Consider three disjoint signed disks O_i, O_j, and O_k. Assume that the bitangents b_{ki}, b_{kj} and b_{-ij} are free and that $\chi_1(b_{ki}, b_{kj})$. We claim that the sequences (1-) $b_{ki}b_{kj}b_{-ij}$, (2-) $b_{ki}b_{ij}b_{-ij}$ if O_i is a negative disk, and (3-) $b_{kj}b_{-ij}b_{-i,-j}$ if O_j is a positive disk are χ_1-sequences.*

Proof. Since an axial symmetry reverses the χ_1-relations and the signs of the disks Claim (1-) is valid for the triple (i, j, k) iff. it is valid for the triple $(-j, -i, -k)$. Therefore it is sufficient to consider the

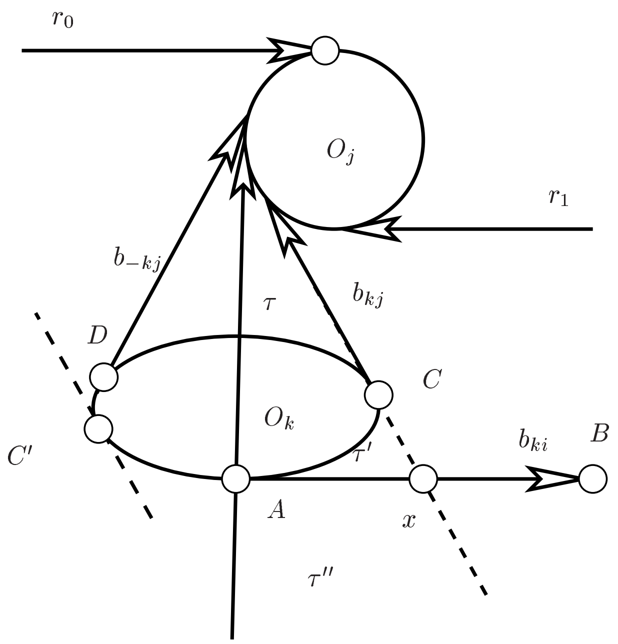

Figure 13: Illustrate the proof of Theorem 8.

case $k > 0$. We use the notations introduced in the statement of Lemma 2 and/or Figure 12 plus the notations of Figure 13, — τ is the region sweept by an half-line that enters O_j and rotates from r_A to r_C; this region is decomposed into three subregions τ, τ', τ'' as indicated in Figure 13 —, and we write that $r < r' < r''$ to state that the half-lines r, r', and r'' enter in this (circular) order the (counterclockwise oriented) disk O_j. Claim (1) is therefore equivalent to say that $r_0 < r_C < r_E < r_1$ which follows easily from the following observations.

1. $r_0 < r_A < r_C < r_1$: Indeed since $\chi_1(b_{ki}, b_{kj})$ holds A belongs to the arc $C'C$ and consequently $A \notin DC'$ (here we use $k > 0$). Therefore according to Lemma 2 applied to the disks O_k, O_j and the point A the half-lines r_0, r_A, and r_1 enter in this order O_j and r_1 does not pierce O_k; this proves our first claim.

2. $r_A < r_B < r_1$: indeed r_A is the supporting half-line of the bitangent $b_{lj} (= b_{Aj})$ where l is the negative disk $[AB]$ and r_1 does not cross O_l. Here we use the second claim of Lemma 2 applied to the disks O_k, O_j and the point A.

3. $B \in EF$ according to Lemma 2 since $r_0 < r_B < r_1$ (from observations 1. and 2.)

4. $E \notin \tau \cup \tau' \cup \tau''$ because b_{Ej} is free and $[EB] \subset O_i$ which is disjoint from O_j, O_k and b_{kj}.

5. $r_A < r_C < r_E$: equivalent to the previous observation

6. if $i < 0$ then $r_0 < r_B < r_E < r_1$. Indeed since $B \in EF$ (observation 3.) then according to Lemma 2 the half-lines r_0 and r_1 do not pierce O_i.

7. if $i > 0$ then $r_C < r_E < r_B$. Indeed if it is not the case then r_C crosses O_i and consequently since b_{kj} is free there is a point of O_i in τ' and, consequently, $E \in \tau'$: impossible.
   Similarly $r_A < r_C < r_B$ (otherwise $E \in \tau'$).

We turn now to the proof of Claim 2. According to Claim 1 the point B belongs to the arc EE' and consequently does not belong to the arc EF'. It follows according to Lemma 2 that r_1 does not pierce O_i and consequently that b_{ki} enters the arc EF. This proves Claim 2. The proof of Claim 3 is similar. $\square$

3.2 Dual order

The dual partial order $\prec_*$ is defined on the same set of bitangents X_0 by $b \prec_* b'$ if and only if $b' \prec b$. Any operator $\Phi : X \to \mathcal{P}(X)$ defined with respect to the partial order $\prec$ has a dual version defined with respect to the dual partial order $\prec_*$; this dual version is denoted Φ_* and satisfies similar properties. In particular $\phi_*(b)$ is the $\prec_*$-minimal bitangent in the set of bitangents $t \in X_0$ that cross b and that $\prec_*$-succeed b (Minimum Element Property), and $G_*(J) = J \setminus \phi_*(J)$ where $G_*(J)$ is the greedy (with respect to $\prec_*$) pseudo-triangulation associated with the filter J of $(X_0, \prec_*)$ (Flip Property), and $R_*(b)$ and $L_*(b)$ are the dual greedy pseudo-triangles associated with the bitangent b. One can easily check that $\phi_* = \phi^{-1}$ and that $\phi_* G(I) = G_*(I_*)$, — here I is a filter of $(X_0, \prec)$ and I_* is its "dual filter" i.e., the filter of $(X_0, \prec_*)$ defined by $I_* = X_0 \setminus I$.

0-cusp edge $a' = \phi_{\text{forw}}(a)$ $a = \phi_{\text{back}_*}(a')$ 	$R(a') = R(\phi(a), I \setminus a)$ $R^1(a) = aeR^1(a')$ $R^2(a') = R^{2_1}(a)$ $R^3(a') = R^{3_1}(a)\phi(a)L^{2_2}(a)$ $R^4(a') = L^{3_2}(a)\pi(e)$	$R_*(a) = R_*(\phi_*(a'), I_* \setminus a')$ $R^1_*(a') = a'eR^1_*(a)$ $R^2_*(a) = R^{2_1}_*(a')$ $R^3_*(a) = R^{3_1}_*(a')\phi_*(a')L^{2_2}_*(a')$ $R^4_*(a) = L^{3_2}_*(a')\pi(e)$
sink cusp edge $a' = \phi_{\text{righ}}(a)$ $a = \phi_{\text{left}_*}(a')$ 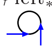	$L(a') = R(\phi(a), I \setminus a)$ $L^1(a') = R^{2_1}(a)$ $L^2(a') = R^{3_1}(a)\phi(a)L^{2_2}(a)$ $L^3(a') = L^{3_2}(a)\pi(e)$ is an arc $L^4(a') = \emptyset, \qquad R^1(a) = ae$	$R_*(a) = L_*(\phi_*(a'), I_* \setminus a'),\ a' \notin \gamma_+$ $R^1_*(a) = L^{2_1}_*(a')$ $R^2_*(a) = L^{3_1}_*(a')\phi_*(a')R^{2_2}_*(a')$ $R^3_*(a) = R^{3_2}_*(a')\pi(e)$ is an arc $R^4_*(a) = \emptyset, \qquad L^1_*(a') = a'$
source cusp edge $a' = \phi_{\text{left}}(a)$ $a = \phi_{\text{righ}_*}(a')$ 	$R(a') = L(\phi(a), I \setminus a),\ a \notin \gamma_+$ $R^1(a') = L^{2_1}(a), L^1(a) = a$ $R^2(a') = L^{3_1}(a)\phi(a)R^{2_2}(a)$ $R^3(a') = R^{3_2}(a)$ $R^4(a') = R^4(a)\pi(e)$	$L_*(a) = R_*(\phi_*(a'), I_* \setminus a')$ $L^1_*(a) = R^{2_1}_*(a'), R^1_*(a') = a'e$ $L^2_*(a) = R^{3_1}_*(a')\phi_*(a')L^{2_2}_*(a')$ $L^3_*(a) = L^{3_2}_*(a')$ $L^4_*(a) = L^4_*(a')\pi_*(e)$
2-cusp edge $a' = \phi(a)$ $a = \phi_*(a')$ 	$L(a') = L(\phi(a))$ $L^1(a') = a'R^{2_2}(a)$ $L^2(a') = R^{3_2}(a)$ $L^3(a') = R^4(a)\pi(e)$ $L^4(a') = \emptyset$	$L_*(a) = L_*(\phi_*(a'))$ $L^1_*(a) = aR^{2_2}_*(a')$ $L^2_*(a) = R^{3_2}_*(a')$ $L^3_*(a) = R^4_*(a')\pi_*(e),$ $L^4_*(a) = \emptyset$

Table 2: Relations between the greedy pseudotriangles of the source a and the sink a' of a positive edge e.

Let I be a filter with $a \in \min_\prec I$ and let e be an edge in X_1 with source a and sink a'. According to the flip property the pseudo-triangle $R(a')$ or $L(a')$ coincides with $R(\phi(a), I \setminus a)$ or with $L(\phi(a), I \setminus a)$ which in turn can be expressed simply in terms of the chains $R^i(a)$, $R^{i_1}(a)$, $R^{i_2}(a)$, $\pi(e)$ and their left versions — recall that R^{i_1}, R^{i_2} are the prefix and suffix subchains (possibly reduced to the empty chains) of the chain $R^i(a)$ cut at the tail of $\phi(a)$ for $i = 2, 3$. Similar relations hold in the dual setting. See Table 2. We use these relations to prove by induction three theorems which will be very useful in our subsequent analysis. The first theorem describes the greedy pseudo-triangles in terms of the ϕ operators.

Theorem 9 (Orbits Theorem). *Let b be a free bitangent not on γ_+ and let $e_{j0}b_{j1}e_{j1}, \ldots, b_{jk_j}e_{jk_j}$ with $k_j \geq 0$ be the chain $R^j(b)$ where b_{jk} stands for a bitangent and e_{jk} for an arc possibly empty. We claim (under the assumption that b_{31} and b_{21} are well-defined) that $b_{31} = \phi_{back}(b), b_{21} = \phi_{righ}(b)$, and that the*

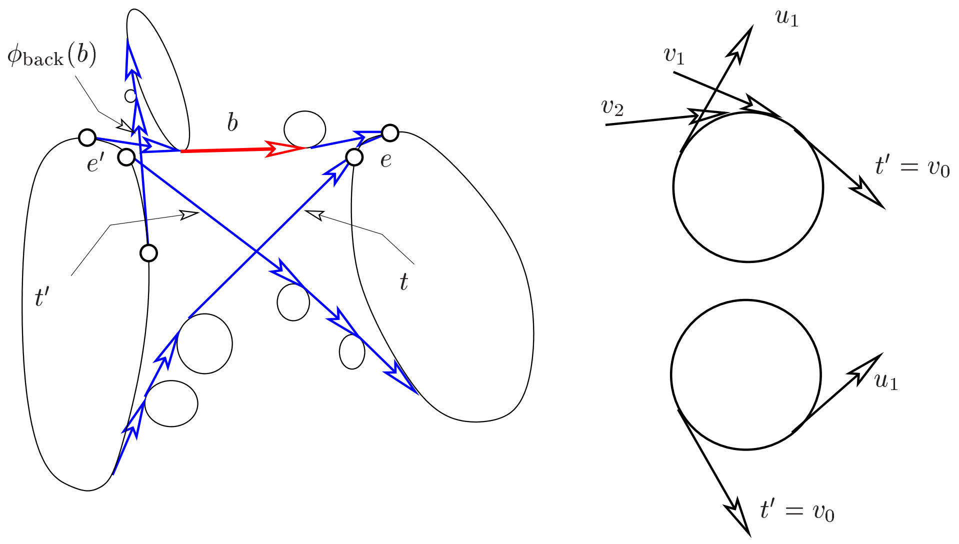

Figure 14: Illustration of Theorems 9 and 10.

chain $e_{20}b_{21}$ *is the envelope of the chain* $lc_3(\sigma_{righ}(b))$*. Futhermore* $b_{1,j+1} = \phi_{forw}(b_{1j})$*,* $b_{2,i+1} = \phi_{back}(b_{2i})$ *and* $b_{3,i+1} = \phi_{forw}(b_{3i})$*.*

Proof. We first show that $b_{31} = \phi_{\text{back}}(b)$. ($b_{31}$ is well-defined iff. b is a left-xx bitangent.) Let $b_1 = (\phi_{\text{back}})_*(b)$, $\sigma := \sigma_{\text{sour}}(b_1)$, and let $b_1e_1b_2e_2\dots$ be the chain $rc_1(\sigma)$. Note that b appears in this chain. The edge e_1 is a positive 0-cusp edge and for $i \geq 2$ the edge e_i is a positive source cusp edge. Therefore if $b_1 \notin \gamma_+$ then, according to Table 2, $R^3(b_2) = R^{3_1}(b_1)\phi(b_1)L^{2_2}(b_1)$ and $R^3(b_{i+1}) = R^{3_2}(b_i)$ for $i \geq 2$. It follows that $\min_\prec R^3(b_i) = \phi(b_1)(= \phi_{\text{back}}(b))$ for any $i \geq 2$, and in particular $b_{31} := \min_\prec R^3(b) = \phi_{\text{back}}(b)$. Now if $b_1 \in \gamma_+$ then b_2 is the successor of b_1 on γ_+, $R^3(b_3) = \pi(b_2)$ and $R^3(b_{i+1}) = R^{3_2}(b_i)$ for any $i \geq 3$. Since $\phi_{\text{back}}(b) = \pi(b_2)$ we get the result. Now we show that $b_{21} = \phi_{\text{righ}}(b)$. Let $b' = \min_\prec R^2(b)$ and let e be the first atom of $R^2(b)$ if this atom is an arc. According to the right-to-left property there is no bitangent entering or leaving the chain $R^1(b)$ and no bitangent entering the chain $R^2(b) = eb'\cdots$. Therefore the chain $R^1(b)eb'$ (which contains one cusp point) is a suffix suchain of the left boundary of the face whose sink is b'; therefore $b' = \phi_{\text{righ}}(b)$. Finally the last part of the theorem is just a reformulation of Theorem 6 in terms of the ϕ operators. □

The second theorem links $R(b)$ and $R(t')$ where t' is the source of the 2-cell $\sigma_{\text{righ}}(b)$.

Theorem 10 (Orbits Theorem Continued). *Let* σ *be a face of the visibility complex* $\mathbb{V}$ *with source* t' *and sink* t *not on* γ_+*. Let* b *be a bitangent of the chain* $lc_2(\sigma)$ *and let* H *be the set of bitangents of the chain* $lc_2(\sigma)$ *smaller than* b*. We claim that* $R^{lea}(b)$ *is a prefix subchain of* $R^{lea}(b;H)$ *and that*

$$R^{lea}(b;H) = etR^{lea_2}(t')\pi(e') \tag{10}$$

where $t'e'$ *and* et *are the envelopes of the chains* $lc_1(\sigma)$ *and* $lc_3(\sigma)$*, respectively.*

Proof. Let $u_1u_2\dots b$ be the increasing sequence of bitangents of the chain $lc_2(\sigma)$ smaller than b. Since the edge with source u_i and sink u_{i+1} is a positive 0-cusp edge, the chain $R^{\text{lea}}(u_{i+1})$ is a prefix subchain of $R^{\text{lea}}(u_i)$ (see Table 2). It is therefore sufficient to show the Lemma in the case $b = u_1$. Let now $v_0e_0v_1e_1\dots v_pe_p$ be the chain $lc_1(\sigma)$ where $v_0 = t'$ and where e_i is the edge with source v_i and sink v_{i+1} where $v_{p+1} = u_1$. We successively treat the cases $p \neq 0$ and $p = 0$. For $p \neq 0$ the edge e_0 is a negative 0-cusp edge, the edges $e_1, e_2, \dots e_{p-1}$ are negative source cusp edges, and the edge e_p is a negative 2-cusp edge. Therefore, according to the negative version of Table 2, $L^3(v_1) = L^{3_1}(v_0)\phi(v_0)R^{2_2}(v_0)$ and

$L^4(v_1) = R^{3_2}(v_0)\pi(e_0)$ which can be rewritten as

$$L^3(v_1)L^4(v_1) = L^{3_1}(v_0)\phi(v_0)R^{\mathrm{lea}_2}(v_0)\pi(e_0);$$

similarly for $i = 1, 2, \ldots, p-1$

$$\begin{aligned} L^3(v_{i+1})L^4(v_{i+1}) &= L^{3_2}(v_i)L^4(v_i)\pi(e_i) \\ &= (L^{3_1}(v_i))^{-1}L^3(v_i)L^4(v_i)\pi(e_i), \end{aligned}$$

and

$$\begin{aligned} R^{\mathrm{lea}}(u_1) &= L^{3_2}(v_p)L^4(v_p)\pi(e_p) \\ &= (L^{3_1}(v_p))^{-1}L^3(v_p)L^4(v_p)\pi(e_p); \end{aligned}$$

from which we deduce that

$$\begin{aligned} R^{\mathrm{lea}}(u_1) &= (L^{3_1}(v_p))^{-1}\ldots(L^{3_1}(v_i))^{-1}\ldots(L^{3_1}(v_1))^{-1} \\ &\quad L^{3_1}(v_0)\phi(v_0)R^{\mathrm{lea}_2}(v_0)\pi(e_0)\pi(e_1)\ldots\pi(e_p) \\ &= e''\phi(v_0)R^{\mathrm{lea}_2}(v_0)\pi(e') \end{aligned}$$

where e'' is a suffix subarc of the arc $L^{3_1}(v_0)$ and where e' is the concatenation of the edges $e_0, \ldots, e_p$ i.e., the envelope of $lc_1(\sigma)$. This proves that $R^2(u_1)$ is not reduced to an arc and consequently according to the Orbits Theorem $e'' = e$. This finishes the proof in the case $p \neq 0$. If $p = 0$, then the edge e_0 with source t' and sink u_1 is a positive source cusp edge or a negative sink cusp edge. In the first case e' is the empty arc and, according to Table 2, $R^2(u_1) = L^{3_1}(v_0)\phi(v_0)R^{2_2}(v_0)$ and $R^3(u_1) = R^{3_2}(v_0)$. In the second case $e_0 = e'$, $R^2(u_1) = L^{3_1}(v_0)\phi(v_0)R^{2_2}(v_0)$, and $R^3(u_1) = R^{3_2}(v_0)\pi(e_0)$. In both cases

$$R^{\mathrm{lea}}(u_1) = L^{3_1}(v_0)\phi(v_0)R^{\mathrm{lea}_2}(v_0)\pi(e')$$

and we conclude as in the case $p \neq 0$. □

The third and last theorem of this subsection provides a duality relation between the greedy pseudo-triangles. This duality relation plays a crucial role in the proof of the sum of squares theorem and is used to prove that some pseudo-triangles are adjacent, and consequently share a common bitangent (cf. Lemma 1).

Theorem 11 (Duality Theorem). *$a \in R^{2_1}(b)$ iff $b \in L^{2_1}_*(a)$.*

Proof. Let $a \in [b, \pi(b)]$. We claim that
(0-) if a' and b are not comparable then $\phi(a') \notin R(b)$,
(1-) if a belongs to $R(b)$ then $\phi_*(a) \prec b$,
(2-) if a leaves or enters $R(b)$ then $b \preceq \phi_*(a)$,
(3-) if a belongs to $R^1(b)$ then b belongs to $R^1_*(a)$,
(4-) if a leaves or belongs to $R^{\mathrm{lea}}(b)$ then b belongs to or leaves $L^{\mathrm{lea}}_*(a)$,
(5-) if a belongs to or enters $R^{\mathrm{ent}}(b)$ then b belongs to or leaves $R^{\mathrm{lea}}_*(a)$,
(6-) if a leaves $R^4(b)$ then b leaves $L^4_*(a)$
from which our theorem follows easily.

Claim (0): Let J be the filter of bitangents t such that $b \preceq t$ or $a' \preceq t$. Since b is minimal in J the pseudotriangle $R(b) = R(b, J) \subseteq G(J)$; since a' is minimal in J the bitangent $\phi(a') \notin G(J)$. It follows that $\phi(a') \notin R(b)$.

Claim (1): According to Claim (0) the bitangent b and $\phi_*(a)$ are comparable; since $a \in R(b)$ it follows that $\phi_*(a)$ does not belong to the filter of bitangents greater than b. Consequently $\phi_*(a) \prec b$.

Claim (2): If a leaves or enters $R(b)$ then a pierces a bitangent b' of $R(b)$ from right-to-left; according to the right-to-left property it follows that $b' \preceq \phi_*(a')$, and consequently $b \preceq \phi_*(a')$ since $b \preceq b'$.

Claims (3,4,5,6) are proved by induction of the linearly ordered set S of bitangents a that belong to or leave or enter the chain $R(b)$. The case $a = b$ is trivial. Suppose the result is true for a bitangent a, and let a' be the successor of a in S. The bitangents a and a' are the source and the sink of an edge e of the visibility complex whose type, thanks to Theorem 6, is restricted.

If a' belongs to $R^1(b)$ then so does a and the edge e is 0-cusp positive edge. By induction $b \in R^1_*(a)$ and, according to Table 2, $R^1_*(a') = a'eR^1_*(a)$. Thus $b \in R^1_*(a')$ which proves the result for a'.

Suppose now that a' belongs to or leaves $R^{\text{lea}}(b)$. The edge e can be a:

(i) negative source cusp or positive sink cusp edge in which case a is the last bitangent of $R^1(b)$ and a' is the first bitangent of $L^2(b)$; by induction $b \in R^1_*(a)$ and, according to Table 2, $R^1_*(a) = L^{2_1}_*(a')$ in both cases. Therefore $b \in L^{2_1}_*(a') \subseteq L^{\text{lea}}_*(a')$.

(ii) negative 2-cusp edge in which case a is the last bitangent of $R^1(b)$ and a' leaves the first atom of $L^2(b)$; by induction $b \in R^1_*(a)$ and, according to Table 2, $R^1_*(a) = aL^{2_2}_*(a')$. Therefore $b = a$, from which we deduce that b leaves $L^{\text{lea}}_*(a')$ since $a = \phi_*(a')$, or $b \in L^{2_2}_*(a')$ which is a subchain of $L^{\text{lea}}_*(a')$.

(iii) positive source cusp or a negative sink cusp edge in which case a belongs to or leaves $R^{\text{lea}}(b)$, $a' \notin R(b)$ (and thus $b \preceq \phi_*(a')$ according to Claim 3); by induction b belongs to or leaves $L^{\text{lea}}_*(a)$ and, according to Table 2, $L^2_*(a) = R^{3_1}_*(a')\phi_*(a')L^{2_2}_*(a')$, $L^3_*(a) = L^{3_2}_*(a')\pi(e)$ (positive) or $L^3_*(a) = L^{3_2}_*(a')$ (negative). Since $b \preceq \phi_*(a')$, b does not belong to nor leaves $R^{3_1}_*(a')$. Therefore $b = \phi_*(a')$ or $b \in L^{2_2}_*(a')$ or b leaves $L^{2_2}_*(a')$ or b leaves $L^{3_2}_*(a')$. In all cases b leaves or belongs to $L^{\text{lea}}_*(a')$.

(iv) negative 0-cusp edge in which case a belongs to or leaves $R^2(b)$, $a' \in R^2(b)$; by induction b belongs to or leaves $L^{\text{lea}}_*(a)$ and, according to Table 2, $L^2_*(a) = L^{2_1}_*(a')$ and $L^3_*(a) = L^{3_1}_*(a')\phi_*(a')R^{2_2}(a')$; and Claim (3) gives $\phi_*(a') \prec b$ from which we deduce that b does not leaves nor belongs to $R^{2_2}_*(a')$. It remains that $b = \phi_*(a')$ or b leaves or belongs to $L^{2_1}_*(a')$ or b leaves $L^{3_1}_*(a')$. In all case b leaves or belongs to $L^{\text{lea}}_*(a')$.

(v) positive 0-cusp edge in which case a enters $R^{\text{lea}}(b)$ and a' leaves $R^3(b)$; by induction b belongs to or leaves $R^{\text{lea}}_*(a)$ and, according to Table 2, $R^3_*(a) = R^{3_1}_*(a')\phi_*(a')L^{2_2}_*(a')$ and $R^4_*(a) = L^{3_2}_*(a')\pi(e)$; Claim (3) gives $\phi_*(a') \preceq_* b$ from which we deduce that b does not leave nor belongs to $R^{2_2}_*(a')$. It remains $b = \phi_*(a')$ or b leaves or belongs to $L^{2_2}_*(a')$ or leaves $L^{3_2}_*(a')$ which is an arc. In all cases b leaves or belongs to $L^{\text{lea}}_*(a')$.

We assume now that a' enters or belongs to $R^{\text{ent}}(b)$. We have to show that b belongs to or leaves $R^{\text{lea}}_*(a') = R^3_*(a')R^4_*(a')$. The edge e can be a:

(i) positive 2-cusp edge; in that case a' enters the first atom of $R^3(b)$, and a leaves the first atom of $R^3(b)$ or is the last bitangent of $R^2(b)$ (in both case a belongs to or leaves $R^{\text{lea}}(b)$); by induction b belongs to or leaves $L^{\text{lea}}_*(a)$ and, according to Table 2, $L^2_*(a) = R^{3_2}_*(a')$ and $L^3_*(a) = R^4_*(a')\pi(e)$. Therefore b leaves or belongs to $R^{\text{lea}}_*(a')$.

(ii) positive sink cusp; in that case a enters or belongs to $R^{\text{ent}}(b)$ and $a' \notin R(b)$; by induction b leaves or belongs to $R^{\text{lea}}_*(a)$ and, according to Table 2, $R^3_*(a) = R^{3_2}_*(a')$ and $R^4_*(a) = R^4(a')\pi(e)$. Therefore b leaves or belongs to $R^{\text{lea}}_*(a')$.

(iii) positive 0-cusp edge; in that case a enters or belongs to $R^3(b)$ and a' belongs to $R^3(b)$; by induction b leaves or belongs to $R^{\text{lea}}_*(a)$ and, according to Table 2, $R^3_*(a) = R^{3_1}_*(a')\phi_*(a')L^{2_2}(a')$ and $R^4_*(a) = L^{3_2}(a')\pi(e)$. Since $a' \in R(b)$ one has $b \prec_* \phi_*(a')$ therefore b does not leave nor belongs to $L^{2_2}(a')$; $R^3(b)$ is nonempty and therefore b is left-xx and consequently b does not leave $R^4_*(a)$ Therefore $b = \phi_*(a')$ or b leaves or belongs to $R^{3_1}_*(a')$. Thus b leaves or belongs to $R^{\text{lea}}_*(a')$.

(iv) negative 0-cusp edge; in that case a leaves $R^4(b)$ and $a' \notin R(b)$. by induction b leaves $L^4_*(a)$ and, according to Table 2, $L^4_*(a) = R^{3_2}_*(a')\pi(e)$. Therefore b leaves $R^{\text{lea}}_*(a')$.

(v) negative sink cusp; in that case a enters or belongs to $R^{\text{ent}}(b)$ et $a' \notin R(b)$. By induction b leaves or belongs to $R^{\text{ent}}_*(a)$ and, according to Table 2, $R^3_*(a) = R^{3_2}_*(a')\pi(e)$ et $R^4_*(a) = \emptyset$. Therefore b leaves or belongs to $R^{3_2}_*(a')$ and consequently b leaves or belongs to $R^{\text{lea}}_*(a')$.

Assume now that a' leaves $R^4(b)$. The edge e can be a:

(i) negative sink cusp edge; in that case a leaves $R^4(b)$; by induction b leaves $L^4_*(a)$ and, according to

Table 2, $L_*^4(a) = L_*^4(a')\pi(e)$. Therefore b leaves $L_*^4(a')$.

(ii) negative 2-cusp edge; in that case a enters or belongs to $R^{\text{ent}}(b)$. By induction b leaves or belongs to $R_*^{\text{ent}}(a)$ and, according to Table 2, $R_*^3(a) = L_*^4(a')\pi(e)$ et $R_*^4(a) = \emptyset$. Therefore b leaves $L_*^4(a')$. □

0-cusp edge a is an obstacle	$R^1(a;H) = aeR^1(a';H)$ $R^2(a';H) = R^2(a;H)$ $R^3(a';H) = R^3(a;H)$ $R^4(a';H) = R^4(a;H)\pi(a)\pi(e)$
sink cusp edge a is an obstacle	$L^1(a';H) = R^2(a;H)$ $L^2(a';H) = R^3(a;H)$ $L^3(a';H) = R^4(a;H)\pi(a)\pi(e)$ $L^4(a';H) = \emptyset, R^1(a;H) = ae$
source cusp edge a is an obstacle	$R^1(a';H) = L^2(a;H), L^1(a;H) = a$ $R^2(a';H) = L^3(a;H)$ $R^3(a';H) = L^4(a;H)\pi(a)$ $R^4(a';H) = \pi(e)$

Table 3: Relations between two greedy pseudotriangles of the source a and the sink a' of a positive edge e in the case where a is a bitangent in H, so $\phi(a;H) = \pi(a)$.

Remark 6. *Table 2 is also valid in the case H non-empty under the assumption that $a \notin H$ and can be easily completed in case $a \in H$ as indicated in Table 3. Therefore very similar results can be proved in the case H non-empty; they are freely used at few places in the sequel without more details.*

3.3 The ϕ_R operator

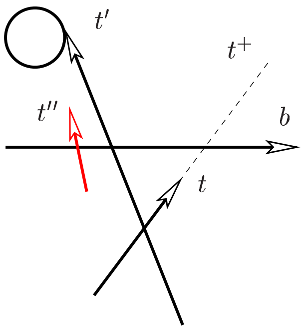

Figure 15: Illustrate the proof of Theorem and Definition 12.

This subsection is devoted to the study of the ϕ_R operator. This operator is not only crucial to the formulation of the sum of squares theorem but also to the flip method based only on the left-turn predicate. We start by several characterizations of ϕ_R and some technicalities before proving that $\phi_R(b)$ can be computed efficiently in the dynamic setting of the GFA using only the left-turn predicate.

Theo. & Def. 12. *For $b \in X_0$ not on the convex hull we define $\phi_R(b)$ to be the maximal bitangent in the range $[b, \pi(b)]$ crossing b and leaving $R(b)$. Similarly we define $\phi_L(b)$ (the left version of $\phi_R(b)$) to* *be the maximal bitangent in the range $[b, \pi(b)]$ crossing b and entering $L(b)$. We claim that $\phi_R(b)$ is the bitangent in the range $[b, \pi(b)]$ crossing b whose crossing point with b is the closest to the tail of b and that*

$$\phi_R(b) = \phi(R(b), \pi L_*(b)), \tag{11}$$

that is $\phi_R(b)$ is the bitangent leaving $R(b)$ and entering $\pi L_*(b)$.

Proof. Let b, t, t' be three bitangents such that t' crosses t from right to left and b crosses t' and the supporting half-line t^+ of t with origin the tail of t from left to right. Then we claim that there exists a bitangent t'' crossing b whose crossing point with b is closer to the tail of b than the crossing points of t' and t^+ with b. Indeed let T be a maximal subset of pairwise non crossing bitangents that cross neither t nor t'; the bitangents in T subdivide free space into pseudotriangles and one pseudoquadrangle Q whose diagonals are t and t'; walking in clockwise order around the boundary of Q starting at the tail of t and ending at the head of t' we found three convex chains C_1, C_2, C_3 in this order where C_1 (resp. C_3) is the empty chain if t (resp. t') leaves (resp. enters) a disk on its left; clearly b intersects the chain C_2 and, consequently, intersects a bitangent t'' in a point which is closer to the tail of b than the crossing points of t and t' with b (see Figure 15). Therefore the bitangent b' in the range $[b, \pi(b)]$ crossing b whose crossing point with b is the closest to the tail of b is well defined and crosses no bitangent crossing b. We show that b' leaves $R^{\mathrm{lea}}(b)$ (and consequently is the maximal one). Let t' crossing b from right to left, that does not leave $R^{\mathrm{lea}}(b)$ and whose crossing point with b is closer to the tail of b than the crossing point of $\phi(b)$ with b. We show that t' is not b'. If t' crosses $\phi(b)$, we are done since b' does not crosses $\phi(b)$. So assume now that t' does not cross $\phi(b)$. In that case $\phi(b)$ leaves the chain $R^2(b)$ and t' crosses from right-to-left a bitangent $t \in R^{2_2}(b)$ (here we use the greedyness of $R(b)$). Moreover, as the angle of t is such that $\Theta(\phi(b)) < \Theta(t) < \Theta(t')$, b intersects the half-line t^+. Therefore, according to the claim at the begining of the proof, t' is not b'. We now prove the last part of the theorem. The left dual version of the first part states that $\phi_{L_*}(b)$, the largest (for $\prec_*$) bitangent in the range $[b, \pi_*(b)]$ that leaves $L_*^{\mathrm{lea}}(b)$, is the bitangent crossing b whose crossing point with b is the closest to the tail of b. Therefore $\phi_{L_*}(b) = \pi_* \phi_R(b)$, and consequently $\phi_R(b)$ enters $\pi L_*(b)$. Uniquiness follows Lemma 1. □

Let e be a positive 0-cusp edge with source b and sink b_R. In the next theorem we show that $\phi_R(b_R)$ belongs to the chain $rc_2(b)$ and that $\phi_R(b_R)$ splits the chain $rc(b)$ into two subchains that are subchains of $R(b)$ and $\pi R_*(b)$, respectively. This result is at the core of the flip method using only the left-turn predicate.

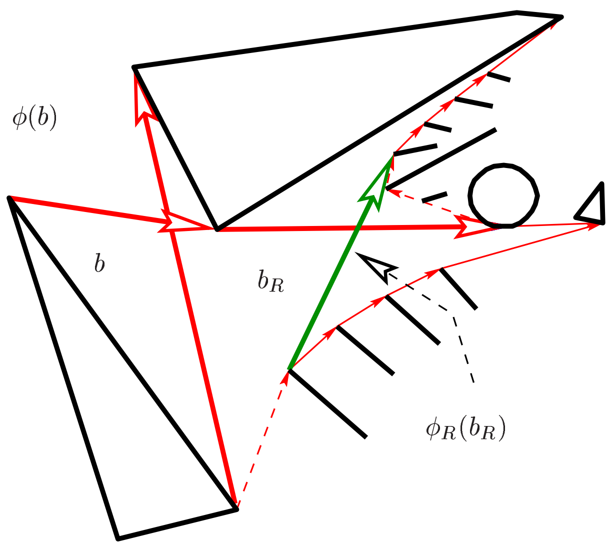

Figure 16: The bitangent $\phi_R(b_R)$ is a *bridge* between $R^{\mathrm{lea}}(b)$ and $R_*\pi(b_R)$.

Theorem 13 (Bridge Theorem). *Let e be a positive 0-cusp edge with source b and sink b_R, let $b_R S(b)$ be the prefix subchain of $R_*\pi(b_R)$ (for $\prec$) cut at the head of $\phi_R(b_R)$, and let $T(b)$ be the suffix subchain of $R^{lea}(b_R)$ cut at the tail of $\phi_R(b_R)$. We claim that $R^{lea}(b_R)$ is a prefix subchain of $R^{lea}(b)$ and that the envelope of $rc(b)$ is the chain*

$$beS(b)\phi_R(b_R)T(b)\phi(b). \tag{12}$$

Furthermore the chain $S(b)$ is a suffix subchain of $\pi R_*(b)$.

Proof. The first part of the Theorem follows immediately from Table 2. There is exactly one cusp point on the chain $beS(b)\phi_R(b_R)$, namely the head of $\phi_R(b_R)$ (if this bitangent is xx-left) or the second cusp point of $R_*\pi(b_R)$. Similarly the chain $\phi_R(b_R)T(b)\phi(b)$ also has only one cusp point. Let E denote the chain defined by the right side of the equality 12 and let $E = E_1E_2E_3$ be its decomposition by these two cusp points. Since the source and sink of E are respectively equal to b and $\phi(b)$, to prove $E = rc(b)$ it is, according to Theorem 3, sufficient to prove that:
(1) there are no bitangents entering E_1,
(2) there are no bitangents tangent to E_2 and
(3) there are no bitangents emanating from E_3.
These three claims follow immediately from the fact that $\phi_R(b_R)$ is the maximal bitangent entering $R_*\pi(b_R)$ and the maximal bitangent leaving $R^{\text{lea}}(b_R)$ (Theorem 12).

We now prove the last part of the Theorem. Let I be a filter such that b is minimal in I. Then b_R belongs to $G(I)$ and consequently $\phi_*(b_R)$ belongs to $G_*(I_*)$. The chain $S(b)$ is a subchain of $\pi R_*^{\text{ent}_2}(b_R)$; we show that $R_*^{\text{ent}_2}(b_R)$ is a subchain $R_*(\phi_*(b_R), I_*)$. To prove this result we pick a bitangent t leaving $R_*^{\text{ent}_2}(b_R)$ (if it exists) and we show that $t \notin G_*(I_*)$. It should be clear that the bitangent t crosses a bitangent $a \in R_*^{2_1}(b_R)$. We prove that $a \in G_*(I_*)$ thus proving $t \notin G_*(I_*)$. According to Table 2, $R_*^2(b) = R_*^{2_1}(b_R)$, and consequently $a \in R_*^2(b)$ from which we deduce that $a \prec b$. Since b is minimal in I we have $a \notin I$. Furthermore, since a belongs to $R_*(b)$ we have $b \prec \phi(a)$ and consequently $\phi(a) \in I$. This achieves to prove that a belongs to $G_*(I_*)$. □

The case where $\phi(b) = \phi_L(b)$ (or $\phi_R(b)$) requires a special treatment in our new flip method. In the next theorem we identify this situation.

Theorem 14 (Equation $\phi_L(b) = \phi(b)$). *The five following assertions are equivalent:*

1. *$\phi(b) = \phi_L(b)$.*
2. *$\phi(b)$ does not cross $\phi_{righ_*}(b)$.*
3. *$\phi_L(b)$ does not cross $\phi_{righ}(b)$.*
4. *$\phi(b)$ leaves the first atom of $R^2(b)$ or leaves the first atom of $\pi R_*^2(b)$*
5. *$\phi_L(b)$ leaves the first atom of $R^2(b)$ or leaves the first atom of $\pi R_*^2(b)$.*

Furthermore if $\phi_L(b)$ leaves the first atom of $\pi R_*^2(b)$ then so does $\phi_R(b)$ and $\phi_R(b) = \pi\phi_*(b)$.

Proof. Let $\sigma = \sigma_{\text{righ}}(b)$, $t = \phi_{\text{righ}}(b)$, $t' = \phi_{\text{righ}_*}(b)$, $et = \epsilon(lc_3(\sigma))$, $t'e' = \epsilon(lc_1(\sigma))$ and let $\Delta = R(b; H)$ and $\Delta_* = R_*(b; H_*)$ where H and H_* are the sets of bitangents of $lc_2(\sigma)$ respectively smaller and greater than b. According to Theorem 10 the chain $R^{\text{lea}}(b)$ is a prefix subchain of Δ^{lea} and $\Delta^{\text{lea}} = etR^{\text{lea}_2}(t')\pi(e')$, from which we deduce that (1-) $\phi(b) = \phi(\Delta, L(b))$ and $\phi_R(b) = \phi(\Delta, \pi L_*(b))$, (2-) $\phi(b)$ and t are non crossing, and (3-0 $\phi(b)$ and t' are non crossing iff. (Claim 4') $\phi(b)$ leaves e or $\pi(e')$. Similarly according to the dual version of Theorem 10 the chain $R_*^{\text{ent}}(b)$ is a prefix subchain of Δ_*^{ent} and $\Delta_*^{\text{ent}} = e't'R_*^{\text{ent}_2}(t)\pi_*(e)$ from which we deduce that (1-) $\phi_L(b)$ and t' are non crossing, (2-) $\phi_L(b)$ and t are non crossing iff. (Claim 5') $\phi_L(b)$ leaves e or $\pi(e')$, and (3-) $\phi_L(b) = \phi(\pi\Delta_*, L(b))$ and $\pi\phi_*(b) = \phi(\pi\Delta_*, \pi L_*(b))$. It follows that Claims (1), (2), (3), (4') and (5') are equivalent. Since e and $\pi(e')$ contain the first atom of $R^2(b)$ and $\pi R_*^2(b)$ respectively, Claim (4) implies Claim (4'). Conversely Claim (4') and Claim (1) imply Claim (4). We prove now the last part of the Theorem. If $\phi(b)$ leaves $\pi(e')$ so does $\phi_R(b)$ since $\phi(b) \preceq \phi_R(b)$. Therefore $\phi_R(b)$ leaves $\pi\Delta_*$ and the identity $\pi\phi_*(b) = \phi(\pi\Delta_*), \pi L_*(b))$ implies $\phi_R(b) = \pi\phi_*(b)$. □

Now we show that $\phi_R(b)$ can be computed efficiently with the χ_1 predicate in the dynamic setting of the GFA. We first examine the case where the edge e is a 1- or 2-cusp edge.

Theorem 15. *Let e be a positive sink cusp edge or a negative source cusp edge with source b and sink b_R. Then $\phi_L(b_R)$ coincides either with $\phi_L(b)$ or with $\pi\phi_*(b_R)$ depending on whether $\phi_L(b) \neq \phi(b)$ or not. Let e be a negative 2-cusp edge then $\phi_R(b_R) = \pi(b)$.*

Proof. Suppose first that e is a positive sink cusp edge or a negative source cusp edge. In both cases, according to Table 2,

$$L^{\text{ent}}(b_R) = R^{3_1}(b)\phi(b)L^{\text{ent}_2}(b)[\pi(e)].$$

where in the above formula the arc $\pi(e)$ is added when e is positive. If $\phi(b) \neq \phi_L(b)$ then $\phi_L(b)$ is the maximal bitangent entering $L^{\text{ent}_2}(b)$, and therefore is the maximal bitangent entering $L^{\text{ent}}(b_R)$. Theorem 12 yields $\phi_L(b_R) = \phi_L(b)$. Assume now that $\phi_L(b) = \phi(b)$. Then there are no bitangents entering $L^{\text{ent}_2}(b)$ and $\phi_L(b_R)$ enters $R^{3_1}(b)$ which is the first atom of $L^2(b_R)$. Since $\phi_L(b_R) = \pi\phi_{R_*}(b_R)$, the bitangent $\phi_{R_*}(b_R)$ leaves the first atom of $\pi_* L^2(b_R)$ and consequently, according to left dual version of Theorem 14, $\phi_{R_*}(b_R) = \phi_*(b_R)$ i.e., $\phi_L(b_R) = \pi\phi_*(b_R)$. Suppose now that e is a negative two cusp edge. According to Table 2 the arc $R^3(b_R)$ is equal to $R^4(b)\pi(e)$. Since $\pi(b)$ is the sink of the chain $R^4(b)$ and since there are no bitangents leaving $\pi(e)$, $\pi(b)$ is the maximal bitangent leaving $R^{\text{lea}}(b)$. We conclude using Theorem 12. □

We now examine the case where the sweep structure sweeps a bitangent b which is the source of a positive 0-cusp edge e.

Theorem 16. *Let e be a positive 0-cusp edge with source b and sink b_R and let I be a filter with $b \in \min_\prec I$. Then $\phi_R(b_R)$ is computable as the bitangent that leaves $R^{lea}(b)$ and enters $\pi R_*^{lea}(b)$ using only the predicate χ_1 in time $|rc(b)| + |R^{2_1}(b)|$ if $\phi_R(b_R)$ leaves a disk on its right, and time $|rc(b)|$ otherwise.*

Proof. We use the following function.

Function PHI-R$(b, b_R$: BITANGENT$)$: PAIR OF ARCS
Comments/* returns the arcs of $R(b)$ and $R_*\pi(b_R)$
that $\phi_R(b_R)$ leaves and enters, respectively */
$C \leftarrow R^{\text{lea}}(b)$; $C' \leftarrow \pi R_*^{\text{lea}}(b)$ in reverse order;
χ_1-WALK(C, C');
return $(\text{Top}(C), \text{Top}(C'))$

Procedure χ_1-WALK$(C, C'$: LIST OF ARCS$)$
while $C \neq \emptyset$ and $C' \neq \emptyset$
$r \leftarrow \text{Top}(C)$; $l \leftarrow \text{Top}(C')$; $t \leftarrow \text{bit}(r, l)$;
if $\chi_1(r^+, t)$ $C \leftarrow \text{Pop}(C)$
elseif $\chi_1(l^+, t)$ $C' \leftarrow \text{Pop}(C')$
else exit;

Its complexity is clearly the sum of the sizes of the chains $R^{2_1}(b)$ and $S(b)$ where $S(b)$ is the chain introduced in Theorem 13. This shows that the complexity of χ_1-WALK is at most $|rc(b)| + |R^{2_1}(b)|$. Using a tandem walk (as explained in the previous section) we can reduce this complexity to $|rc(b)|$ in case where $\phi_R(b_R)$ leaves a disk on its left. It remains to prove the correction of the method. Call r_g and l_g the arcs of $R^{\text{lea}}(b)$ and $R_*^{\text{ent}}\pi(b_R)$ that $\phi := \phi_R(b_R)$ leaves and enters, respectively. The correction of χ_1-WALK is a straightforward consequence of the following four claims:

1. If $l \prec l_g$ and $r = r_g$ then $\chi_1(t, r_g^+)$ and $\chi_1(l^+, t)$.

2. If $l = l_g$ and $r \prec r_g$ then $\chi_1(r^+, t)$.

3. If $l \prec l_g, r \prec r_g$ and $\chi_1(t, r^+)$ then $\chi_1(l^+, t)$.

4. $\chi_1(\phi, l_g^+)$ and $\chi_1(\phi, r_g^+)$.

Let O_i and O_j be the signed disks supporting the arcs $r \preceq r_g$ and $l \preceq l_g$. We denote by p (resp. q) the contact point of r^+ (resp. l^+) with O_i (resp. O_j).

Proof of Claim 1. $r = r_g$ **and** $l \prec l_g$**.** We first prove the relation $\chi_1(t, r_g^+)$. Let x be the intersection point of the supporting lines of l^+ and r_g^+. We observe that (1) $\chi_2(l^+, r_g^+)$ holds, (2) b_{xj} is free and $\vec{l^+} = \vec{b}_{xj}$ i.e. x, q appear in this order along the supporting line of l^+, and (3) either (Case 1) b_{xi} is free and $\vec{r_g^+} = \vec{b}_{xi}$ with O_i positive or (Case 2) b_{ix} is free and $\vec{r_g^+} = \vec{b}_{ix}$. In the first case one has $\chi_1(b_{xj}, b_{xi})$. Therefore according to Theorem 8 Claim (3-) applied to the disks O_j, O_i and the point x the sequence $b_{ij}b_{xi}b_{-ji}$ is a χ_1-sequence and, consequently, $\chi_1(b_{ij}, b_{xi})$ i.e., $\chi_1(t, r_g^+)$. In the second case one has $\chi_1(b_{xj}, b_{ix})$ or, equivalently, $\chi_1(b_{x-i}, b_{xj})$. Therefore according to Theorem 8 Claim (1-) applied to the disks O_{-i}, O_j and the point x the sequence $b_{x-i}b_{xj}b_{ij}$ is a χ_1-sequence and, consequently, $\chi_1(b_{ij}, b_{ix})$ i.e., $\chi_1(t, r_g^+)$. We turn now to the the proof of the relation $\chi_1(l^+, t)$. We redefine x to be the intersection point of the supporting lines of l^+ and ϕ and we observe that (1-) b_{xj} is free and $\vec{l^+} = \vec{b}_{xj}$, (2-) b_{ix} is free and $\vec{\phi} = \vec{b}_{ix}$, and (3-) $\chi_1(b_{xj}, b_{ix})$ since $b_R \prec l^+ \prec \phi \prec \pi(b_R)$. Therefore according to Theorem 8 Claim (1-) applied to the disks O_{-i}, O_j and the point x the sequence $b_{x-i}b_{xj}b_{ij}$ is a χ_1-sequence and, consequently, $\chi_1(b_{xj}, b_{ij})$, i.e, $\chi_1(l^+, t)$.

Proof of Claim 2. $r \prec r_g$ **and** $l = l_g$**.** Let x be the intersection point of the supporting lines of ϕ and r^+. We observe that (1-) $\chi_2(r^+, \phi)$ holds, (2-) b_{xj} is free and $\vec{\phi} = \vec{b}_{xj}$, (3-) O_i is a negative disk, and (4-) either (Case 1) b_{xi} is free and $\vec{r^+} = \vec{b}_{xi}$ or (Case 2) b_{ix} is free and $\vec{r^+} = \vec{b}_{ix}$. In the first case one has $\chi_1(b_{xi}, b_{xj})$. Therefore according to Theorem 8 Claim (2-) the sequence $b_{xi}b_{ij}b_{-ij}$ is a χ_1-sequence and, consequently, $\chi_1(b_{xi}, b_{ij})$ i.e., $\chi_1(r^+, t)$. In the second case $\chi_1(b_{xj}, b_{x-i})$; therefore according to Theorem 8 Claim (3-) applied to the disks O_j, O_{-i} and the point x the sequence $b_{-ij}b_{x-i}b_{-j-i}$ is a χ_1-sequence and, consequently, $\chi_1(b_{x-i}, b_{-j-i})$ i.e., $\chi_1(r^+, t)$.

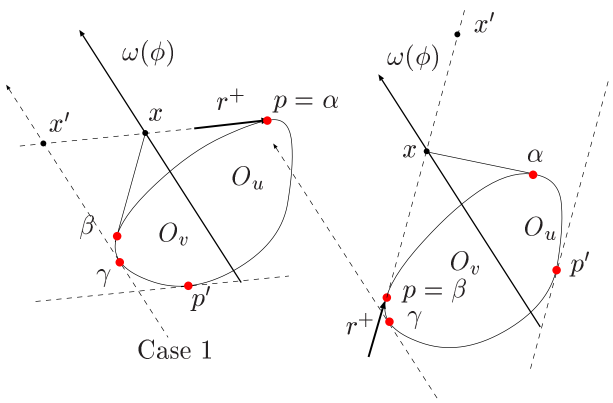

Figure 17: The two cases of the proof of Claim 3.

Proof of Claim 3. $l \prec l_g$ **and** $r \prec r_g$**.** Let y (resp. x) be the intersection point of the supporting lines of l^+ (resp. r^+) and ϕ. We observe that (1-) O_i is a negative disk, (2-) x and y appear in this order along the supporting line of ϕ and the line segment $[xy]$ is free, (3-) $\chi_2(l^+, \phi)$ and $\chi_2(r^+, \phi)$ hold, (4-) b_{yj} is free and $\vec{l^+} = \vec{b}_{yj}$ i.e., y, q appear in this order along the supporting line of l^+, and (5-) either (Case 1) b_{xi} is free and $\vec{r^+} = \vec{b}_{xi}$ or (Case 2) b_{ix} is free and $\vec{r^+} = \vec{b}_{ix}$. The supporting line of ϕ splits the convex-hull of x and O_i into two negative disks denoted O_u and O_v where O_u lies on the right of the supporting line of ϕ (note that O_u and O_v might be reduced to the point x.) Let α and β be the contact point of the bitangents b_{xi} and b_{ix} with the disk O_i. We note that $t(= b_{ij})$ leaves the disk O_u, O_v, or the arc $\beta\alpha$. See Figure 17. We have $\vec{\phi} = \vec{b}_{uy}$ and $\chi_1(b_{yj}, b_{uy})$. Therefore according to Theorem 8 Claim (1-) applied to the disks O_{-u}, O_j and the point y the sequence $b_{y-u}b_{yj}b_{uj}$ is a χ_1-sequence and, consequently,

$$\chi_1(b_{yj}, b_{uj}). \tag{13}$$

We have $\vec{\phi} = \vec{b}_{-vy}$ and $\chi_1(b_{yj}, b_{-vy})$ i.e., $\chi_1(b_{yv}, b_{yj})$. Therefore according to Theorem 8 Claim (2-) applied to the disks O_v, O_j and the point y the sequence $b_{yv}b_{vj}b_{-vy}$ is a χ_1-sequence and consequently $\chi_1(b_{yv}, b_{vj})$ which can be rewritten as $\chi_2(b_{vj}, \phi)$. This last relation proves that b_{vj} leaves the arc γx of O_v where γ is the tangency point of the supporting line to O_v with direction $\vec{\phi}$. Since by hypothesis $\chi_1(t, r^+)$ the bitangent t leaves the arc pp' of the disk O_i where p' is the tangency point of the supporting line to O_i with direction $-\vec{r^+}$.

Case 1. In that case $p = \alpha$ and the points p, p', γ, x appear in this order along the boundary of O_i. Since the arcs pp' and γx are disjoint it follows that $b_{ij} = b_{uj}$ and we are done thanks to Equation 13.

Case 2. In that case $p = \beta$, γ appears on the boundary of O_i, and the point $\beta, \alpha, p', \gamma$ appears in this order along the boundary of O_i. Therefore $b_{ij} = b_{uj}$ and we are done thanks to equation 13 or b_{ij} leaves the arc $\beta\alpha$. In this latter case we apply Theorem 8 to the disks O_i, O_j and O_k where O_k is the line segment $[x, y]$ oriented counterclockwise.

Proof of Claim 4. $r = r_g$ **and** $l = l_g$**.** This is clear since $b_R \preceq r_g^- \prec \phi \prec r_g^+ \prec \pi(b_R)$, and $b_R \preceq l_g^- \prec \phi \prec l_g^+ \preceq \pi(b_R)$. □

3.4 Proof of the sum of squares theorem

We now turn to the proof of the sum of squares theorem. The proof is split into several lemmas. But first we introduce some notations and reformulate the duality theorem. Let b be a non-hull bitangent and recall that $R'(b)$ denotes the set of bitangents $t \neq \phi_{\text{righ}}(b)$ of the chain $R^{2_1}(b)$ not in $\sigma_{\text{sour}}(b)$ (Theo. & Def. 7). We introduce also the set $r(b)$ of bitangents of the chain $rc_2(\sigma_{\text{righ}}(b))$ as well as its left version $l(b)$ i.e., the set of bitangents of the chain $lc_2(\sigma_{\text{left}}(b))$; clearly $a \in r(b)$ iff. $b \in l(a)$. The following theorem is a reformulation/variation of the duality theorem.

Theorem 17 (Duality Theorem bis). *Let $Z_R(b)$ be the set of bitangents in $R'(b), r(b)$ or $R'_*(b)$, the dual version of $R'(b)$. Then $a \in Z_R(b)$ iff. $b \in Z_L(a)$ (the left version of Z_R).* □

The first step of the proof of the sum of squares theorem is to show, among others more technical things, that if a and b are non crossing and $\phi(a, b) = \phi(b, a)$ then $a \in Z_R(b)$ or $a \in Z_L(b)$. To this end we are going to show that the bitangents of $Z_R(b)$ are bitangents of the right boundary of some face of some visibility complex. Let $H = \{\phi_R(b), \phi_L(b)\}$. We assume that $\phi_R(b) \neq \phi_L(b)$ and we introduce the 2-face $\sigma(b)$ of $\mathbb{V}_H$ defined by the directed line segment $[x_1, x_2]$ where x_1 and x_2 are the crossing points of the bitangent b with the bitangents $\phi_R(b)$ and $\phi_L(b)$, respectively. We study the right boundary of $\sigma(b)$. If $\phi_R(b)$ and $\phi_L(b)$ leave the same disk then the right boundary of the face $\sigma(b)$ is reduced to the curve γ_p where p is the tail of $\phi_R(b)$ or the tail of $\phi_L(b)$ depending on whether $\phi_R(b)$ leaves a disk on its right or its left. So assume now that $\phi_R(b)$ and $\phi_L(b)$ do not leave the same disk. In particular this implies that the minimal bitangent of $R^2(b)$ is well-defined, and therefore coincides with $\phi_{\text{righ}}(b)$. Let E be the subchain of $R^{\text{lea}}(b)$ with source $\phi_{\text{righ}}(b)$ and sink $\phi(b)$ when $\phi_{\text{righ}}(b) \prec \phi(b)$ i.e., $\phi(b) \neq \phi_L(b)$ otherwise E is not defined. We define a chain C as follows:

1. if $\phi_L(b)$ and $\phi_R(b)$ enter the same disk, then $\phi(b) = \phi_R(b)$ and $\phi_*(b) = \pi_*\phi_L(b)$ and we define
$$C = \begin{cases} E\epsilon & \text{if } \phi_R(b) \text{ and } \phi_L(b) \text{ enter in this order the first atom of } L^2(b) \\ E & \text{if } \phi_R(b) \text{ and } \phi_L(b) \text{ enter in this order the first atom of } \pi L_*^2(b) \end{cases}$$
where ϵ is the subarc of $L^2(b)$ with source $\phi_R(b)$ and sink $\phi_L(b)$.

2. if $\phi_R(b)$ and $\phi_L(b)$ do not enter the same disk, then the minimal bitangent of $L^2(b)$ is well-defined, and consequently is $\phi_{\text{left}}(b)$, and we define
$$C = \begin{cases} \emptyset & \text{if } \phi(b) = \phi_L(b) \text{ (Case 1)} \\ E\epsilon & \text{if } \phi(b) = \phi_R(b) \text{ (Case 2)} \\ E & \text{otherwise (Case 3)} \end{cases}$$

where ϵ is the subarc of $L^2(b)$ with source $\phi(b)$ and sink $\phi_{\rm left}(b)$. In Case 1 $\phi_L(b)$ leaves the first atom of $R^2(b)$, and in Case 2 $\phi_R(b)$ enters the first atom of $L^2(b)$ (but not $\phi_L(b)$).

Let now D be the envelope of the right boundary of the face $\sigma := \sigma_{\rm righ}(b)$ and let C_* be the dual version of the chain C. Since the sink $t := \phi_{\rm righ}(b)$ of D is the source of C and since the source t' of D is the sink of C_* we can introduce the envelope $\Sigma(b)$ of the concatenation of the chains $\overline{C}_*$ (i.e., the chain C_* in reverse order), D and C.

Lemma 3. *$\phi_R(b)$ crosses no bitangent of the chain $\Sigma(b)$.* □

Proof. Since $\phi_R(b)$ leaves $R^{\rm lea}(b)$ it is clear that $\phi_R(b)$ crosses no bitangent of C (and similarly no bitangent of C_*). We show now that $\phi_R(b)$ crosses no bitangent of D. According to Theorem 10 the bitangent $\phi_R(b)$ leaves $e, \pi(e')$ or $R^{{\rm lea}_2}(t')$. In the first two cases $\phi_R(b)$ is a bitangent of $lc(\sigma)$, and consequently crosses no bitangent of $rc(\sigma)$. We now examine the third case i.e., $\phi_R(b)$ leaves the chain $R^{{\rm lea}_2}(t')$. Let e'' be the edge with source t' that t' enters. Assume first that e'' is a 1-cusp edge; in that case the chain $rc_2(\sigma)$ is a subchain of $R^2(t')$ and we are done. Assume now that e'' is a (positive) 0-cusp edge and let t'_R be the sink of e''; in that case, according to the Bridge Theorem, $\phi_R(b)$ leaves the pseudo-triangle on the left of $\phi_R(t'_R)$ in the quadrangle obtained by merging $R(t')$ and $R_*\pi(t'_R)$; since this pseudo-triangle contains the envelope of $rc(\sigma)$ (Bridge Theorem) we are done. □

Lemma 4. *The chain $\Sigma(b)$ is the envelope of the right chain of the 2-face $\sigma(b)$.*

Proof. We only prove the lemma in the case where nor $\phi(b)$ nor $\pi\phi_*(b)$ coincides with $\phi_L(b)$ or $\phi_R(b)$. The other cases can be treated similarly. In that case $\Sigma(b)$ contains exactly two cusp points that split $\Sigma(b)$ into the three convex chains

$$\Sigma_1(b) = \overline{e' R_*^{3_1}(b)\phi_*(b)}, \tag{14}$$

$$\Sigma_2(b) = \overline{(e_{20_*}(b)t')^{-1} R_*^{2_1}(b)} r(b)(e_{20}(b)t)^{-1} R^{2_1}(b) \tag{15}$$

$$\Sigma_3(b) = eR^{3_1}(b)\phi(b) \tag{16}$$

where $t'e'$, $r(b)$ and et are the envelopes of the chains $rc_1(\sigma)$, $rc_2(\sigma)$ and $rc_3(\sigma)$; therefore
(1-) there are no bitangents in $\mathbb{V}_H$ entering or leaving $\Sigma_2(b)$
(2-) there are no bitangents leaving $\Sigma_3(b)$ and $\Sigma_3(b)$ is an arc concatenated with a bitangent in $\mathbb{V}_H \setminus H$.
(3-) there are no bitangents entering $\Sigma_1(b)$ and $\Sigma_1(b)$ is the concatenation of an arc with a bitangent in $\mathbb{V}_H \setminus H$.
Therefore, according to Theorem 3 and since the bitangents of the chain $\Sigma(b)$ are in $\mathbb{V}_H$, the chain $\Sigma(b)$ is the envelope of the right boundary of the face of $\mathbb{V}_H$ with sink $\phi(b)$ and source $\phi_*(b)$. This proves in particular that $\phi(b)$ and $\phi_*(b)$ are crossing. We show now that $\phi(b)$ is the sink of the face $\sigma(b)$. Let x be the intersection point of $\phi(b)$ with b and let $r(\theta)$ be the ray in $\mathbb{V}_H$ with origin x and angle θ. We observe that the backward view of the ray $r(\theta)$ does not change when θ ranges over the semi-open interval $[\Theta(b), \Theta(\phi(b)))$; consequently $\phi(b)$ belongs to the boundary of the face $\sigma(b)$ and consequently since $\phi(b)$ and b are crossing $\phi(b)$ is the sink of the face $\sigma(b)$. A dual argument shows that $\phi_*(b)$ is the source of $\sigma(b)$. □

It should be clear at this point that the bitangents in $Z_R(b)$ are bitangents of the chain $\Sigma_2(b)$. The next step is to describe the greedy pseudo-triangles $R(a,b)$ and $L(a,b)$ of the bitangents a of the chain $\Sigma_2(b)$ and to evaluate $\phi(a,b)$.

Lemma 5. *Let $a \in \Sigma_2(b)$. Then $R(a,b) = R(a)$ and $\phi(a,b) = \phi(a; \{\phi_L(b)\})$.*

Proof. Let C_1 and C_2 (resp. D_1 and D_2) be the backward and forward views of the rays in the face $\sigma_{\rm righ}(a)$ (resp. $\sigma_{\rm left}(a)$) of $\mathbb{V}_H$ and let C_3 (resp. C_4) be backward (resp. forward) view of the rays in $\sigma_{\rm back}(a)$ (resp. $\sigma_{\rm forw}(a)$) in case a is a left-xx (resp. xx-left) bitangent. One can check that these chains are distincts, and consequently since $\phi_R(b)$ (resp. $\phi_L(b)$) belongs to D_1 (resp. D_2) the chains C_1, C_2, C_3, C_4 do not contain $\phi_R(b)$ or $\phi_L(b)$. This proves that the bitangents of $R(a)$ belong to $\mathbb{V}_H$, and consequently that $R(a,b) = R(a)$. We now prove that $\phi(a,b) = \phi(a;\{\phi_L(b)\})$. Since $\phi_L(b) \in H(b)$ one has $\phi(a;\{\phi_L(b)\}) \preceq \phi(a,b)$ (Theorem 4). Let t be a bitangent that leaves $R(a)$ and pierces a. If $t \prec \phi(a,b)$ then t pierces the forward view of the rays in the face $\sigma(b)$ i.e., pierces $\phi_L(b)$; therefore $\phi(a;\{\phi_L(b)\}) \not\preceq \phi(a,b)$. This proves that $\phi(a,b) = \phi(a;\{\phi_L(b)\}$. □

Lemma 6. *Let $a \in \Sigma_2(b)$, let*

$$C(b) = b_1 e_1 b_2 e_2 \dots e_p b_{p+1} = C'(b) b_{p+1} \tag{17}$$

with $p \geq 0$ be the subchain of $L^{ent}(b)$ with source $b_1 = \phi(b)$ and sink $b_{p+1} = \phi_L(b)$, and let e_0 be the first atom of the chain $\Sigma_3(b)$ if this atom is an arc. Assume that a is a xx-right bitangent. Then $\phi_L(b)$ and $\phi_{forw}(a,b)$ are comparable and one of the three following cases holds

1. *$\phi_{forw}(a,b) \prec \phi_L(b)$ and $L^{ent}(a,b)$ is a prefix subchain of $e_0C'(b)$.*
2. *$\phi_L(b) = \phi_{forw}(a,b)$ and $L^{ent}(a,b) = e_0C(b)$.*
3. *$\phi_L(b) \prec \phi_{forw}(a,b)$ and $L^{ent}(a,b) = e_0C(b)e_{p+1}$ where e_{p+1} is the suffix subarc of the first atom of $L^3(a)$ with source $\phi_L(b)$ and sink $\phi_{forw}(a)(= \phi_{forw}(a,b))$.*

Furthermore if a is an xx-left bitangent then $L^{ent}(a,b) = e_0C(b)e_{p+1}$ where e_{p+1} is the suffix subarc of the first atom of $L^3(a)$ with source $\phi_L(b)$ and sink $\pi(a)$.

Proof. According to the Orbit Theorem (1-) $L^1(a,b)$ is the suffix subchain of $\Sigma_2(b)$ with source a, and (2-) the minimal bitangent of $L^2(a,b)$ is the minimal bitangent of $\Sigma_3(b)$ i.e., $\phi(b)$ (cf. definition of $\Sigma(b)$) or $L^2(a,b)$ is reduced to an arc. Now according to the Orbit Theorem $b_{i+1} = \phi_{\text{forw}}(b_i)$ for $i < p$; since for $t \in L^2(b)$ with $b_1 \prec t$ one has $t \in X_0(H(b))$ it follows that $\phi_{\text{forw}}(b_i, b) = b_{i+1}$ for $i < p$ and either b_p is xx-right or b_p is xx-left and $\phi_{\text{forw}}(b_p, b) = b_{p+1}$. This proves our lemma. □

Lemma 7. *Let $a \in \Sigma_2(b)$. Then the bitangents $\phi(a,b), \phi(b)$ and $\phi_L(b)$ are comparable and distinct, and one of the three following cases holds:*

1. *$\phi(a,b) \prec \phi(b)$ and $\phi(a,b) = \phi_R(a)$ is a xx-left bitangent and $\phi(b)$ is left-xx*
2. *$\phi_L(b) \prec \phi(a,b)$ and $\phi(a,b) = \phi(a)$ is xx-right and $\phi_L(b)$ is right-xx*
3. *$\phi(b) \prec \phi(a,b) \prec \phi_L(b)$ and if $a \in Z_R(b)$ then $\phi(a,b) = \phi(b,a)$.*

Proof. According to Lemma 5, $R(a,b) = R(a)$. Assume first that $\phi(a,b) \prec \phi(b)$. In that case $\phi(a,b)$ enters the (positive) arc e_0. It follows that $\phi(a,b)$ is the maximal bitangent of X_0 leaving $R^{\text{lea}}(a,b)$. Since $R(a,b) = R(a)$, $\phi(a,b)$ is the maximal bitangent of X_0 leaving $R(a)$, and consequently $\phi(a,b) = \phi_R(a)$. Assume now that $\phi_L(b) \prec \phi(a,b)$. In that case $\phi(a,b)$ enters the (negative) arc e_{p+1} which belongs to $L^3(a)$. Since $R(a,b) = R(a)$, the bitangent $\phi(a,b)$ leaves $R(a)$ and enters $L(a)$ i.e., $\phi(a,b) = \phi(a)$. Finally assume that $\phi(b) \prec \phi(a,b) \prec \phi_L(b)$ and that $a \in Z_R(b)$. According to Lemma 6, the bitangent $\phi(a,b)$ enters $L^{\text{ent}}(b)$ and crosses b. According to the Duality Theorem $b \in Z_L(a)$, and consequently b belongs to $X_0(H(a))$. The bitangent $\phi(b,a)$ is therefore well-defined. We prove successively that $\phi(b,a) \preceq \phi(a,b)$ and $\phi(a,b) \preceq \phi(b,a)$. Since b belongs to $Z_L(a)$, b belongs to $\Sigma_2(a)$, and consequently, thanks to the left version of Lemma 5, $\phi(b,a) = \phi(b;\{\phi_R(a)\})$. The bitangent $\phi(a,b)$ crosses b but $\phi_R(a)$ doesn't, therefore $\phi(a,b) \neq \phi_R(a)$. Since the bitangents $\phi(a,b)$ and $\phi_R(a)$ leave $R(a)$ they do not cross and the minimum element property in $X_0(\phi_R(a))$ yields $\phi(b;\{\phi_R(a)\}) \preceq \phi(a,b)$ from which we deduce $\phi(b,a) \preceq \phi(a,b)$. Since $\phi(a,b)$ and $\phi(b,a)$ both enter $L(a,b)$ the opposite inequality will follow from the fact that $\phi(a,b)$ crosses no bitangent of $R^{2_1}(b,a)$. The knowledge of this chain is provided by the left version of Lemma 6 which we can apply because $\phi_R(a)$ and $\phi_L(a)$ do not enter the same disk. This last claim is a consequence of $\phi_R(a) \neq \phi(a)$ which in turn is a consequence of the existence of at least two bitangents leaving $R(a)$ and crossing a, namely $\phi(a,b)$ and $\phi_R(a)$. A bitangent $t \in R^{2_1}(b,a)$ is therefore equal to $\phi(a)$ or belongs to $R^{2_2}(a)$. Since $\phi(a,b)$ leaves $R(a)$, $\phi(a,b)$ crosses no such bitangent t. □

We are now ready to prove what we announced in the introduction of this subsection.

Lemma 8. *Let a, b non crossing with $\phi(a,b) = \phi(b,a)$. Then*

1. *$\phi_L(b) \neq \phi_R(b)$ and $\phi(a,b)$ pierces a and b.*
2. *$a \in Z_R(b)$ or $a \in Z_L(b)$ (the left version of Z_R)*
3. *If $a \in Z_R(b)$ then $\phi(b) \neq \phi_L(b)$.*
4. *$b \prec \phi(a)$ and $a \prec \phi(b)$.*

Proof. Proof of Claim 1. The bitangent $\phi(a,b)$ crosses a and $\phi(b,a)$ crosses b so the second part of the claim is clear. The existence of $\phi(a,b)$ implies that a does not intersect $\phi_R(b)$ and therefore $\phi(a,b) \neq \phi_R(b)$. Since $\phi(a,b) = \phi(b,a)$ it follows that $\phi(b,a) \neq \phi_R(b)$ and therefore there are at least two bitangents intersecting b. This proves that $\phi_L(b) \neq \phi_R(b)$.

Proof of Claim 2. Without loss of generality we can assume that $\phi(a,b)$ crosses a and b in that order. We prove successively that a belongs to $\Sigma_2(b)$ and $a \in Z_R(b)$. Let x be the crossing point of b with $\phi(a,b)$ and let y_1, y and y_2 denote the crossing points of a with $\phi_R(a), \phi(a,b)$ and $\phi_L(a)$, respectively.

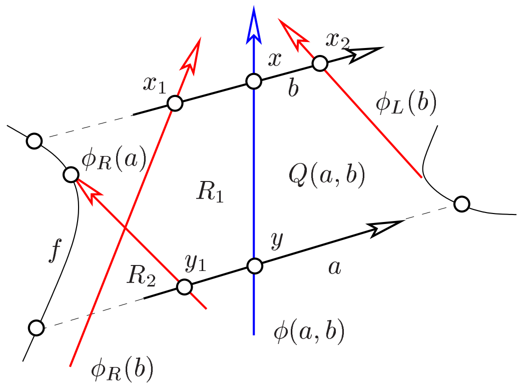

Figure 18: Illustration of the proof of the second claim of Lemma 8.

With the help of the face $\sigma_{\text{sour}}(b,a) \in \mathbb{V}(H(a))$ with source b and sink $\phi(b,a)$ we define a pseudo-quandrangle $Q(b,a)$ as follows. Let s and s' be the two sub-segments of $\text{seg}_{H(a)}(b)$ and $\text{seg}_{H(a)}(\phi(b,a))$ with target x and let e be the sub-chain of the backward view of $\sigma_{\text{sour}}(b,a)$ defined by the source points of s and s'. The interior of the pseudo-triangle $\delta := ss'e$ is a subset of the face $\sigma_{\text{sour}}(b,a)$ therefore it is a subset of free space. Since the segment $[y_1, y]$ crosses s' from left to right, it follows that y_1 belongs to e or the segment $[y_1, y]$ crosses s. We show that the latter case does not hold. Indeed otherwise $[y_1, y]$ would cross $s \setminus [x_1, x]$ (because a and b do not cross) and consequently $\phi_R(b)$ would cross one of the two segments $[y, x]$ or $[y_1, y]$. The two cases are impossible because $[y, x] \subset \phi(a,b)$, $[y_1, y] \subset a$ and $\phi_R(b)$ crosses neither of the bitangents $\phi(a,b)$ or a. This achieves to prove $y_1 \in e$. We define $Q(b,a)$ to be the pseudo-quadrangle defined by the three segments $s, [y,x], [y_1, y]$ and the subchain of e defined by y_1 and the source point of s. Since $Q(b,a) \subset \delta$, the interior of $Q(b,a)$ is a subset of free space.
The bitangent $\phi_R(b)$ crosses s but does not cross $\phi(b,a)$ nor $[y_1, y]$, therefore $\phi_R(b)$ leaves or crosses the chain e. In both cases $\phi_R(b)$ partitions $Q(b,a)$ in two regions; let R_1 be the one containing y. Let t be the sub-segment of $\text{seg}(a)$ with target y_1, t' be the sub-segment of $\text{seg}(\phi_R(a))$ with source y_1 and let f be the sub-arc of the left backward view of a in $\mathbb{V}$ defined by the source and target points of t and t' respectively. The pseudo-triangle $\tau := tft'$ is a subset of free space. If $\phi_R(b)$ leaves e, we define R_2 as being equal to τ. Otherwise $\phi_R(b)$ crosses $\phi_R(a)$ and therefore partitions the interior of τ in two regions and we define R_2 to be the one containing y_1. Merging the two adjacent (along a sub-segment of $\phi_R(a)$) regions R_1 and R_2 yields a pseudo-quadrangle Q.
The construction operated above to build $Q(b,a)$ using the face $\sigma_{\text{sour}}(b,a)$ and its backward view can be

performed on the face $\sigma_{\rm sour}(a,b)$ and its forward view to yield the interior free pseudo-quadrangle $Q(a,b)$. Merging the two pseudo-quadrangles Q and $Q(a,b)$ adjacent along the segment $[y,x]$, defines an interior free pseudo-quadrangle having $[x_1,x_2]$ and $\mathrm{seg}_{H(b)}(a)$ as two sides. This proves that a belongs to $\Sigma_2(b)$. To prove that a belongs to $Z_R(b)$ we use the description of the chain $\Sigma_2(b)$ given in Lemma 4 and prove that $a \notin rc_2(\sigma_{\rm sour}(b))$. This last conclusion together with its dual version will allow us to conclude. Let t be a bitangent of $rc_2(\sigma_{\rm sour}(b))$ with $t \neq \phi_{\rm righ}(b)$, we prove that $t \neq a$. Note that $b = \phi_{{\rm left}_*}(t)$. Moreover, according to the Duality Theorem, b belongs to $L_*^{21}(t)$ and therefore $\phi_*(t)$ does not leave the the first atom of $L_*^2(t)$. This last conclusion together with the left dual version of Theorem 14 prove that $\phi_L(t)$ crosses b (first case) or $\phi_*(t)$ leaves the first atom of $\pi_* L^2(t)$ (second case). In the first case $t \neq a$ because $\phi_L(a)$ does not cross b. In the second case assume $t \in X_0(H(b))$ otherwise it is clear that $t \neq a$. The bitangent $\pi\phi_*(t)$ is smaller than $\phi_{\rm left}(t)$ $(= \phi(b))$ and is the maximal bitangent crossing t and greater than t. Consequently $\phi(t,b) \preceq \pi\phi_*(t)$ and therefore $\phi(t,b) \prec \phi(b)$. This last inequality implies that $\phi(t,b)$ does not cross b and achieves to prove $t \neq a$.

Proof of Claim 3. It is sufficient to show that there are at least two bitangents entering $L(b)$ and crossing b. The first of these two bitangents is $\phi_L(b)$ and does not cross a. Since $b \in Z_L(a)$, b belongs to $\Sigma_2(a)$, and, according to the left version of Lemma 5, $L(b,a) = L(b)$. Consequently the bitangent $\phi(b,a)$ crosses a and b and enters $L(b)$. $\phi(a,b)$ is the second bitangent we were looking for.

Proof of Claim 4. Assume $a \in Z_R(b)$. The preceeding claim states that $\phi_L(b) \neq \phi(b)$ therefore the sink of the chain Σ introduced in Lemma 4 is equal to $\phi(b)$. This proves $a \prec \phi(b)$. The left version of this result yields $a \prec \phi(b)$ for $a \in Z_L(b)$. To prove $b \prec \phi(a)$ simply swap the roles of a and b. □

So far we have identify the set of pairs (a,b) such that $\phi(a,b) = \phi(b,a)$. In the next lemma we study the equation $\phi(a,b) = \phi(a',b')$. Let $\mathcal{G}$ denote the set of pairs $(a,b) \in X_0 \times X_0$ such that $\phi(a,b) = \phi(b,a)$ and $\phi_*(a,b) = \phi_*(b,a)$. We partition $\mathcal{G}$ into two subsets $\mathcal{G}_R$ and $\mathcal{G}_L$ where $\mathcal{G}_R$, $\mathcal{G}_L$ contain the pairs $(a,b) \in \mathcal{G}$ such that $a \in Z_R(b)$, $a \in Z_L(b)$, respectively.

Lemma 9. *Let $(a,b),(a',b') \in \mathcal{G}_R$ where b,b' are xx-left bitangents and where a,a' are xx-right bitangents. Assume that $\phi(a,b) = \phi(a',b')$ and that $b' \not\preceq b$. Then $b \prec b'$ and $\pi(b) \prec \phi_{forw}(a',b')$.*

Proof. We first prove that

$$\phi_L(b) \prec \pi(b) \prec \phi_L(b') \prec \pi(b'). \tag{18}$$

Let $t_0 = \phi(b,a)$. Since t_0 enters $L(b,a)$ and since $L(b,a) = L(b)$ (see Lemma 7) we can introduce the suffix subchain $C(b)$ of $L(b)$ with source t_0. Let $t_0, t_1, \ldots$ be the sequence of bitangents defined by $t_{i+1} = \phi_{\rm forw}(t_i)$ for $i \geq 0$. According to the Orbits Theorem

$$C(b) = t_0\epsilon_0 t_1 \epsilon_1 \ldots t_p \epsilon$$

where p is the first index such that t_p is an xx-right bitangent and where ϵ is the arc with source t_p and sink $\pi(b)$ (here we use the fact that b is an xx-left bitangent). Similarly we introduce the suffix subchain $C(b')$ of $L(b')$ with source $t'_0 = \phi(b',a')$ and write

$$C(b') = t'_0\epsilon'_0 t'_1 \epsilon'_1 \ldots t'_{p'} \epsilon'$$

where $t'_{i+1} = \phi_{\rm forw}(t'_i)$ and where ϵ' is the arc with source $t'_{p'}$ and sink $\pi(b')$. Since by hypothesis $t_0 = t'_0$ and $b' \not\preceq b$ it follows that

$$C(b') = C(b)\pi(e)$$

where e is an arc with source b and sink b'. Since $\phi_L(b)$ and $\phi_L(b')$ are the maximal bitangents entering $C(b)$ and $C(b')$, respectively, it follows that $\phi_L(b) = \phi_L(b')$ or $\phi_L(b')$ enters $\pi(e)$ which proves Claim (18). It remains to prove that $\phi_L(b) \neq \phi_L(b')$. If $\phi_L(b) = \phi_L(b')$ then according to Lemma 7 one has $\phi(a;\{\phi_L(b)\}) = \phi(a';\{\phi_L(b)\})$, and consequently $a = a'$ since $\phi(.;\{\phi_L(b)\}$ is one-to-one. But $a = a'$ implies that $\phi(b,a) = \phi(b',a')$, and consequently $b = b'$ since $\phi(.,a)$ is one-to-one. We also note that it follows from our analysis that $\phi(a,b), \phi_L(b), \phi_L(b')$ are pairwise non crossing. We prove now that

$$\pi(b) \prec \phi_{\rm forw}(a',b').$$

According to Lemma 8 Claim 4 the bitangent $\phi_{\text{forw}}(a', b')$ enters $L(b')$ or $\phi_L(b') \prec \phi_{\text{forw}}(a', b')$. In the latter case we are done since we have seen that $\pi(b) \prec \phi_L(b')$. Assume now that $\phi_{\text{forw}}(a', b')$ enters $L(b')$, and consequently enters $C(b')$. We claim that

$$\phi_L(b) \prec \phi_{\text{forw}}(a', b') \tag{19}$$

from which we deduce that $\phi_{\text{forw}}(a', b')$ enters $\pi(e)$, and consequently that $\pi(b) \prec \phi_{\text{forw}}(a', b')$. We prove successively that

1. $\phi_L(b')$ and a are non crossing.
2. $\phi_L(b)$ and a' are crossing.
3. $\phi_L(b) \prec \pi(a')$.
4. $\phi_L(a', b') \prec \phi_{\text{forw}}(a', b')$.
5. $\phi_L(a', b') = \pi\phi_*(a'; T)$ where $T = L^2(a', b')$.
6. $\phi_L(b) \preceq \pi\phi_*(a'; T)$.

from which Claim (19) follows.

Claim 1. : We observe that

$$\phi_*(a) \prec \pi_*\phi_L(b') \prec \phi(a; \phi_L(b)) \tag{20}$$

— the first inequality follows Claim (18) and Lemma 8 Claim 4 applied to ϕ_*, and the second inequality follows from the fact that $\phi_L(b')$ and $\phi(a, b)$ $(= \phi(a; \{\phi_L(b)\}))$ enter $L^{\text{ent}}(b')$ with $\phi(a, b) \prec \phi_L(b')$ —, from which we deduce according to the Minimum Element Property that $\phi_L(b')$ and a are non crossing.

Claim 2. : If $\phi_L(b)$ and a' are non crossing we can write that $\phi(a; G) = \phi(a'; G)$ where $G = \{\phi_L(b), \phi_L(b')\}$, and consequently that $a = a'$ since $\phi(\cdot; G)$ is one-to-one.

Claim 3. : This is a consequence of $\phi_L(b) \prec \phi_L(b')$, $\phi_L(b') \preceq \pi\phi_*(b')$ and $\pi\phi_*(b') \prec \pi(a')$ (Lemma 8 Claim 4 applied to ϕ_*).

Claim 4. : We observe that $\phi_L(a', b')$ enters $L^{\text{ent}}(a', b')$ and that $\phi_{\text{forw}}(a', b')$ is the sink of $L^{\text{ent}}(a', b')$ (see the Orbits Theorem with a' is an xx-right bitangent).

Claim 5. : is clear.

Claim 6. : Affiliation of Claims 2. and 3., $\phi_L(b) \in X_0(T)$ and the minimum element property in $X_0(T)$. □

Proof of the Sum of Squares Theorem. Let now I be a filter of $(X_0, \prec)$ and let $\mathcal{G}(I), \mathcal{G}_R(I), \mathcal{G}_L(I)$ be the sets of pairs $(a, b) \in \mathcal{G}, \mathcal{G}_R, \mathcal{G}_L$ respectively, such that $b \in I \setminus \pi(I)$. The sum of squares Theorem can be restated as $|\mathcal{G}(I)| = O(|B|)$ and it is clearly sufficient to prove that $|\mathcal{G}_R(I)| = O(|B|)$. Let $\mathcal{G}_R^+(I)$ (resp. $\mathcal{G}_R^-(I)$) be the set of $(a, b) \in \mathcal{G}_R(I)$ such that $\phi(a, b) \notin \pi(I)$ (resp. $\phi_*(a, b) \in I$). We claim that

1. $\mathcal{G}_R(I) = \mathcal{G}_R^+(I) \cup \mathcal{G}_R^-(I)$
2. $\mathcal{G}_R^-(I) = (\mathcal{G}_R^+)_*(J)$ where $J = X_0 \setminus \pi(I)$
3. $|\mathcal{G}_R^+(I)| = O(|B|)$

from which the theorem follows easily.

Claim 1. follows $\phi(a, b) \preceq \pi\phi_*(a, b)$

Claim 2. $(\mathcal{G}_R^+)_*(J) = \{(a, b) \in (\mathcal{G}_R)_*(J) \mid \phi_*(a, b) \in I\}$. The result follows from $(\mathcal{G}_R)_*(J) = \mathcal{G}_R(I)$.

Claim 3. Since $Z_R(b)$ contains at most one xx-left bitangent the number of pairs (a, b) in $\mathcal{G}_R(I)$ such that a is an xx-left bitangent is a $O(|B|)$. Similar arguments allow us to restrict our attention to the case where b is an xx-left bitangent and a is an xx-right bitangent. Let $E(b)$ denote the set of bitangents a such that $(a, b) \in \mathcal{G}_R^+(I)$. It is sufficient to prove that if the pairs (a, b) and (a', b') satisfy the assumptions of Lemma 9 then $a' = \min_\prec E(b')$. Let a'' be a bitangent of $Z_R(b')$ with $a'' \prec a'$ (otherwise we are done). We claim that

$$\phi(a', b') \prec \phi_{\text{forw}}(a', b') \preceq \phi(a'', b') \tag{21}$$

— indeed two consecutive bitangents t, t' of $Z_R(b')$ verify $t' = \phi_{\text{forw}_*}(t, b')$ and according to the Orbits Theorem $\phi(t, b') \prec \phi_{\text{forw}}(t, b')$ — from which we deduce using Lemma 9 that $\pi(b) \prec \phi(a'', b')$, and consequently $a'' \notin E(b')$. This proves that $a' = \min_\prec E(b')$. □

3.5 Pairs of bitangents that satisfy the conditions of the Sum of Squares Theorem

We identify now a set of pairs (a, b) that satisfy the conditions of the sum of squares theorem. As we shall see in the next section this set is closely related to the the multiset of bitangents visited during the course of the flip method based on the left-turn predicate. We denote by b_R the successor of b in the pseudo-triangle $R(b)$.

Theorem 18. *Let (a, b) be a pair of bitangents that satisfies one of the three following conditions: Condition (1-) $a \in R'(b)$ and $\phi_R(b)$ is a right-xx bitangent, Condition (2-) $a \in r(b)$, $\phi_R(b)$ and $\phi_{R_*}(a)$ are right-xx bitangents, and $\phi_{righ}(b)$ is an xx-left bitangent (or $\phi_{left}(a)$ is a right-xx bitangent), and Condition (3-) $b \in L'(a)$ and $\phi_L(a)$ is an xx-left bitangent. Then $(a, b) \in \mathcal{G}_R$.*

Proof. Assume that the pair (a, b) satisfies Condition (1-) We claim that
(i) $\phi_R(b)$ and $\phi(b)$ are right-xx,
(ii) $\phi_{L_*}(b) \neq \phi_*(b)$ and $\phi_L(b) \neq \phi(b)$, and
(iii) $\phi(a) \preceq \phi_L(a) \prec \phi_L(b) \preceq \pi\phi_*(b)$
from which we deduce, according to Lemma 7 and its dual version, that $\phi(a, b) = \phi(b, a)$ and $\phi_*(a, b) = \phi_*(b, a)$ i.e., $(a, b) \in \mathcal{G}_R$. Point (i) is clear since by assumption $\phi_R(b)$ is right-xx. Point (ii) is consequence of (i), $a \in R'(b)$ and Theorem 14. To prove point (iii) we use the following Lemma.

Lemma 10. *Let $t' = \phi_{back}(t)$ where t' enters on its right the disk that t leaves (on its right). Then $\phi_L(t') \prec \phi(t)$ (because $\phi(t) = \min_\prec L^3(t')$ and $\phi_L(t')$ enters $L^{ent}(t')$).*

According to this lemma $\phi_L(a) \prec \phi(\phi_{\text{righ}}(b))$. Furthermore since $\phi_L(b)$ crosses $\phi_{\text{righ}}(b)$ it follows that $\phi(\phi_{\text{righ}}(b)) \preceq \phi_L(b)$ (Minimum Element Property). Therefore $\phi_L(a) \prec \phi_L(b)$. Since $\phi(a) \preceq \phi_L(a)$ is clear we are done.

Assume now that a verifies Condition (2-). We claims that
(i) $\phi_R(b)$ and $\phi(b)$ are right-xx
(ii) $\phi_{R_*}(a)$ is right-xx and $\phi(a)$ is xx-left,
(iii) $\phi_{L_*}(b) \neq \phi_*(b)$ and $\phi_L(b) \neq \phi(b)$, from which we deduce, according to Lemma 7 that $\phi(a, b) = \phi(b, a)$ and $\phi_*(a, b) = \phi_*(b, a)$ i.e., $(a, b) \in \mathcal{G}_R$.

The proof under Condition (3-) is similar using the left version of Lemma 7 and its dual version. □

Theorem 19. *Let $Y_R(b)$ be the set of bitangents a such that one of three following conditions is satisfied (Condition 1) $a \in R'(b)$ and $\phi_R(b_R)$ is a right-xx bitangent, (Condition 2) $a \in r(b)$, $a_L \in L^1(a)$, $b_R \in R^1(b)$ and $\phi_R(b_R)$ is a right-xx bitangent, $\phi_L(a_L)$ is an xx-left bitangent and $\phi_{righ}(b)$ is a right-left bitangent, (Condition 3) $a \in R'_*(b)$ and $\phi_L(a_L)$ is an xx-left bitangent. Let a and b be such that $a \in Y_R(b)$ (iff. $b \in Y_L(a)$). Assume that $a \neq \min_\prec Y_R(b)$ and $b \neq \min_\prec Y_L(a)$. Then $(a, b) \in \mathcal{G}$.*

Proof. Let (a, b) be a pair of bitangents that satisfy the assumptions of the theorem. We claim that

1. $a \in R'(b)$ or $\phi_L(a)$ is a xx-left bitangent
2. $b \in L'(a)$ or $\phi_R(b)$ is a right-xx bitangent

from which we deduce — since the two conditions $a \in R'(b)$ and $b \in L'(a)$ (which is equivalent to $a \in R'_*(b)$ by the Duality Theorem) cannot occur simultaneously — that the pair (a, b) satisfies the assumptions of Theorem 18 and, consequently, that $(a, b) \in \mathcal{G}$. We now prove Claim (1) (Claim (2) is the left version of Claim (1)). Assume that $Y_R(b) \setminus R'(b)$ contains at least two bitangents – otherwise we are done – and let $c = \min_\prec Y_R(b) \setminus R'(b) = \min_\prec Y_R(b)$ and c_L be its successor in $L^1(c)$. Let e be an elementary edge belonging to $rc_2(\sigma_{\text{righ}}(b))$ or $R^2_*(b)$ and let t and t' be its source and sink respectively. Assume that $\phi_L(t)$ is an xx-left bitangent; we prove that $\phi_L(t')$ is also an xx-left bitangent. The edge e is either a negative regular edge (case (i)) or a negative source cusp edge (case (ii)). In case (i), Table 2 tells us that $\phi_L(t')$ enters $L^{\text{ent}}(t)$ and consequently $\phi_L(t')$ is an xx-left bitangent. In case (ii), Table 2 tells us that $\phi_L(t')$ enters the first atom of $L^2(t')$ or enters $L^{\text{ent}}(t)$, consequently $\phi_L(t')$ is an xx-left bitangent. Since c belongs to $Y_R(b)$ the bitangent $\phi_L(c_L)$ is an xx-left bitangent. It follows that all the bitangents $\phi_L(a)$ are xx-left bitangents for $a \in Y_R(b) \setminus R'(b)$ and $a \neq \min_\prec Y_R(b)$. □

We are now in position to justify Theorem 7. Assume that $R'(b)$ is nonempty and that $\phi(b)$ is right-xx. Then $\phi_R(b_R)$ right-xx. It follows, according to the above theorem, that the set of pairs (a, b) such that $a \in R'(b)$ and $\phi(b)$ right-xx is a subset of $\mathcal{G}$, and consequently, thanks to the sum of squares theorem, its size is a big-O of $|B|$.

4 Optimal flip using the left-turn predicate

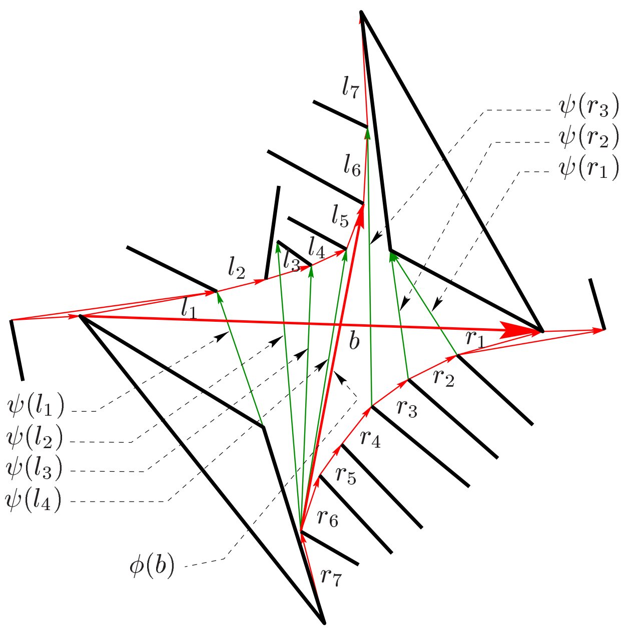

Figure 19: Flipping b using only χ_1. Here $r_u = r_3$, $l_5 = l_q$, $r_6 = r_p$ and $\psi(r_6) = \psi(l_5) = \phi(b)$.

We maintain as described previously not only $G(I)$ and $S(I)$ but also the dual pseudo-triangulation $G_*(I_*)$; furthermore for each bitangent $t \in G(I)$ which is right (resp. left) minimal we maintain a pointer to the arc of $G(I)$ that $\phi_R(t)$ leaves (resp. $\phi_L(t)$ enters) and a pointer to the arc of $G_*(I_*)$ that $\phi_R(t)$ enters (resp. $\phi_L(t)$ leaves). Let b be minimal in the filter I, and let b_R, b_L be its successors in $R(b), L(b)$, respectively. Thanks to Theorem 14 we can detect in constant time if $\phi(b)$ coincides with $\phi_R(b)$ or $\phi_L(b)$; in that case we are done. We assume now that $\phi_R(b)$ and $\phi_L(b)$ are both $\neq \phi(b)$, and consequently $\phi(b)$ does not leave the first atom of $R^2(b)$ nor enter the first atom of $L^2(b)$. If b and b_R are separated by a cusp then b_R is left minimal in $I' := I \setminus b$ and $\phi_L(b_R) = \phi_L(b)$ (cf. Theorem 15). If b and b_R are not separated by a cusp then b_R is right minimal and we compute $\phi_R(b_R)$, as explained in the previous section, to restore the invariant of our algorithm (cf. Theorem 16). A similar computation is done with b_L. It remains to compute $\phi(b)$. We set $r_1 = b_R$ if b and b_R are separated by a cusp otherwise we set $r_1 = \phi_R(b_R)$. If r_1 is a left-xx bitangent then $\phi(b)$ is also a left-xx bitangent and therefore leaves the first atom (which is an arc) of $R^3(b)$. Similarly if l_1 is an xx-right bitangent then $\phi(b)$ is an xx-right bitangent and enters the first atom (which is an arc) of $L^3(b)$. In both cases the method FLIP2 is easily adapted to work only with the χ_1 predicate (cf. Remark 4). Assume now that r_1 and l_1 are right-xx and xx-left bitangents, respectively. Let $l_2 \prec \ldots \prec l_{\max}$ be the increasing sequence of bitangents in $L^2(b)$ greater than l_1. We denote by l_q the last element of this sequence which is smaller than $\phi(b)$. Note that $l_1, \ldots, l_q \in lc(b)$ according to the Bridge Theorem. We introduce similar notations for $R^2(b)$ i.e., $r_1, r_2, \ldots, r_p, \ldots, r_{\max}$. According to the Duality Theorem the pseudo-triangles $R(b)$ and $R_*(l_i)$, for $i \leq q$, are incident to the same bitangent (b) and they don't (locally) overlap. It follows (cf. Lemma 1) that there is a unique bitangent $\psi(l_i)$ that leaves $R(b)$ and enters $\pi R_*(l_i)$. The operators ψ, ϕ_R, and ϕ are related as described in the following theorem.

Theorem 20. *There is an index q' with $1 \leq q' \leq q$ such that $\psi(l_i) = \phi_R(l_{i+1}) \neq \phi(b)$ for $1 \leq i < q'$ and $\psi(l_i) = \phi(b)$ for $q' \leq i \leq q$. More precisely:*

1. *For $1 \leq i < q'$ the bitangent $\psi(l_i)$ enters the chain $\pi R_*^{lea}(\phi_*(l_{i+1}), I'_*)$ where $I' = I \setminus b$.*
2. *For $q' \leq i < q$ the bitangent $\pi\phi_*(l_{i+1})$ (resp. $\psi(l_i)$) leaves (resp. enters) the first atom of $L^3(b)$.*
3. *$\psi(l_q)$ enters the arc l_q^+ which is a subarc of the last atom of $\pi R_*^{lea}(\phi_*(l_{q+1}), I_*)$.*

Proof. According to Theorem 10 and since l_i belongs to the chain $lc_2(b)$ the chain $R^{\text{lea}}(l_i)$ is a prefix subchain of the chain $e\phi(b)R^{\text{lea}_2}(b)e'$ where e is a suffix subarc of $L^{3_1}(b)$. Moreover $\phi_R(l_{i+1}) \prec \phi(l_i) \preceq \phi_R(l_i)$. If b_L belongs to $L^2(b)$ the arc e' is empty or is an elementary arc; otherwise $\pi(b_L)$ leaves e' and intersects l_1. It follows that $\phi_R(l_i) \prec \pi(b_L)$ for $i \geq 2$, and therefore $\phi_R(l_i)$ does not leave e' for $2 \leq i \leq q$. If $\phi_R(l_i)$ leaves $R^{\text{lea}}(b)$ for all the indices $2 \leq i \leq q$ we set $q' = q$, otherwise let $q' < q$ be such that $\phi_R(l_i)$ leaves $R^{\text{lea}}(b)$ for $2 \leq i \leq q'$ and $\phi_R(l_i)$ leaves e for $q' < i \leq q$.

Let $1 \leq i < q'$. The bitangent $\phi_R(l_{i+1})$ leaves $R^{\text{lea}_2}(b)$ and therefore crosses b. To prove that $\phi_R(l_{i+1}) = \psi(l_i)$ it remains to prove that $\phi_R(l_{i+1})$ enters $\pi R_*^{\text{lea}}(l_i)$. According to the dual version of Theorem 14 and since $b = \phi_{\text{righ}_*}(l_{i+1})$ and $\phi_R(l_{i+j})$ crosses b, $\phi_R(l_{i+1}) \neq \pi\phi_*(l_{i+1})$, and consequently $\phi_R(l_{i+1})$ enters $\pi L_*^{\text{lea}_2}(l_{i+1})$. The edge e_i with source l_i and sink l_{i+1} is an elementary positive 0-cusp edge, so according to Table 2

$$R_*^{\text{lea}}(l_i) = R_*^{3_1}(l_i)\phi_*(l_{i+1})L_*^{\text{lea}_2}(l_{i+1})\pi(e_i).$$

This equality proves that $\phi_R(l_{i+1})$ enters $\pi R_*^{\text{lea}}(l_i)$. To complete the proof of claim 1 it remains to show that $\phi_R(l_{i+1})$ enters $\pi R_*^{\text{lea}}(\phi_*(l_{i+1}), I'_*)$. This will be a consequence of the fact that $C(b) = \phi_*(l_{i+1})L_*^{\text{lea}_2}(l_{i+1})$ is a chain of $G_*(I'_*)$. Since $l_{i+1} \in I'$, we have $\phi_*(l_{i+1}) \in G_*(I'_*)$ and the Orbits Theorem tells us that the bitangents of the chain $C(b)$ also belong to $G_*(I'_*)$. We now show that there are no bitangent in $G_*(I'_*)$ leaving the chain $C(b)$. Since $b \in G_*(I'_*)$ and $b = \phi_{\text{righ}_*}(l_{i+1})$, the bitangents of $R_*^2(l_{i+1})$ belong to $G_*(I'_*)$. Since a bitangent leaving $C(b)$ crosses a bitangent of $R_*^2(l_{i+1})$, $G_*(I'_*)$ contains no bitangent leaving $C(b)$. This achieves to prove that $C(b)$ is a chain of $G_*(I'_*)$; the result follows.

Assume $q' \neq q$ and let $q' \leq i < q$. The bitangent $\phi_R(l_{i+1})$ leaves the arc e. Since e is a suffix subarc of $L^{3_1}(b)$, $\phi_R(l_{i+1})$ leaves $L^{3_1}(b)$ and $\phi(b)$ enters the first atom of $L^3(b)$. Moreover, $\phi_R(l_{i+1})$ leaving the first atom of $R^2(l_{i+1})$, the dual version of Theorem 14 states that $\phi_R(l_{i+1})$ $(= \pi\phi_{L_*}(l_{i+1}))$ is equal to $\pi\phi_*(l_{i+1})$. To finish the proof of claim 2 it remains to prove that $\phi(b)$ enters $\pi L_*^{\text{lea}}(l_i)$. Since $\phi_*(l_{i+1})$ $(= \min_{\prec_*} R_*^3(l_i)$ by the Orbits Theorem) is xx-left, $R_*^2(l_i)$ is not reduced to an arc and therefore contains b $(= \phi_{\text{righ}_*}(l_i))$. Let $u = \max_{\prec_*} R_*^2(l_i)$; the bitangent u is xx-left and the Orbits Theorem allows us to claim that $\phi(b) \prec \phi(u) \prec \pi(u)$. Since $\pi\phi_*(l_{i+1})$ $(= \phi_R(l_{i+1}))$ leaves $L^{3_1}(b)$ we have $\pi\phi_*(l_{i+1}) \prec \phi(b)$; it follows that $\phi(b)$ enters the subarc of $\pi R_*^{3_1}(l_i)$ whose source and sink for $\prec_*$ is $\pi(u)$ and $\pi\phi_*(l_{i+1})$ respectively. Consequently $\phi(b)$ enters $\pi R_*^{\text{lea}}(l_i)$.

We now prove claim 3. It is clear that $\phi(b)$ enters l_q^+. As $l_{q+1} = \phi_{\text{forw}}(l_q)$ (Orbit Theorem) the arc l_q^+ is a subarc of $\pi R_*^4(l_q)$ thus $\phi(b)$ enters $\pi R_*^{\text{lea}}(l_q)$ and $\phi(b) = \psi(l_q)$. We now prove that $L_*^{\text{lea}_2}(l_{q+1})l_q^+$ is a chain of the pseudotriangle $G_*(I_*)$. The result will follow because the source (for $\prec_*$) of this chain is $\phi_*(l_{q+1})$. We already know, thanks to the Orbits Theorem, that the bitangents of this chain belong to $G_*(I_*)$ because $\phi_*(l_{q+1}) \in G_*(I_*)$. We must show that there are no bitangent in $G_*(I_*)$ leaving it. To this end we exhibit a chain $C(b)$ consisting of bitangents in $G_*(I_*)$ intersected by any bitangent leaving $L_*^{\text{lea}_2}(l_{q+1})l_q^+$. The source (for $\prec_*$) of this chain will be $\phi_*(l_q)$; we show that this bitangent enters $R_*^{\text{ent}}(l_{q+1})$. This is a consequence of the three simple following remarks:
(i) $b \prec_* \phi_*(l_q) \prec_* \phi_*(l_{q+1})$ (Orbits Theorem),
(ii) $b = \min_{\prec_*} R_*^2(l_q)$ (Orbits Theorem),
(iii) b enters $R_*^{\text{ent}}(l_{q+1})$ (Duality Theorem).
The chain $C(b)$ which we were looking for is the suffix (for $\prec_*$) subchain of $R_*^{\text{ent}}(l_{q+1})$ with source $\phi_*(l_q)$. □

We come back to the computation of $\phi(b)$. A first idea is to compute successively the bitangents $\psi(l_i)$ for $i = 1, 2, \ldots, q'$ with the procedure PHI-R . Unfortunately the repeated use of this procedure is in general to costly. To cut the budget we first compute the bitangents $\psi(r_1), \psi(r_2), \ldots, \psi(r_u)$ — with the procedure PHI-L which is the left version of PHI-R — where u is the first index such either $\psi(r_u) = \phi(b)$ or $\psi(r_u)$ is an xx-left bitangent. In this later case we can now safely compute the bitangents $\psi(l_1), \psi(l_2), \ldots, \psi(l_q)$ with the function PHI-R-NEW(b, l_{i+1}) which differs from PHI-R only in the initialization of the variables C and C':

Function PHI-R-NEW$(b, l$: BIT.): PAIR OF ARCS
1 $C \leftarrow [r_u, \ldots, r_{\max}]$; $C' \leftarrow \pi R_*^{\text{lea}}(\phi_*(l), I_*)c$;
2 χ_1-WALK(C, C');
3 **return** $(\text{Top}(C), \text{Top}(C'))$;

here c is a suffix subarc of the first atom of $L^3(b)$ added to the last atom of $\pi R_*^{\text{lea}}(\phi_*(b'), I_*)$ in the case where this last atom is itself a subarc of the first atom of $L^3(b)$ (added to treat case (2-) of Theorem 20). We call this procedure FLIP1. We show now that its complexity is constant amortized.

Let $1 \leq i < q'$; we have $\psi(l_i) = \phi_R(l_{i+1})$. Let $r_u, \ldots, r_m$ be the sequence of bitangents in C visited during the procedure PHI-R-NEW(b, l_{i+1}). As $\psi(l_i)$ leaves $R^{\text{lea}}(b)$ and crosses b it should be clear that $p \leq m$. According to Theorem 16, the number of bitangents in the chain C' visited by the procedure $Y_R(l_i)$ is $|rc(l_i)|$. When $\psi(l_i)$ is a left-xx bitangent the procedure PHI-R-NEW(b, l_{i+1}) can be too costly. Running tandem walks as in Theorem 16 allows us to compute $\psi(l_i)$ in time $|rc(l_i)|$ when $\psi(l_i)$ is a left-xx bitangent. Assume now that $\psi(l_i)$ is a right-xx bitangent. We show that the bitangents $r_u, \ldots, r_m$ all belong to $Y_R(l_i)$ except maybe for r_p. This will prove that the complexity of PHI-R-NEW(b, l_{i+1}) is at most $1+|Y_R(l_i)|+|rc(l_i)|$ when $\psi(l_i)$ is a right-xx bitangent. If $u \leq p$ then the bitangents $r_u, \ldots, r_p$ belong to $rc(b)$ whereas l_i belongs to $lc(b)$. According to the definition of $r(l_i)$, this implies that the bitangents $r_u, \ldots, r_p$ belong to $r(l_i)$. If $p \neq u$, let j be an index such $u \leq j \leq p-1$; as $\phi_R(l_{i+1})$ $(= \psi(l_i))$ is a right-xx bitangent, we have $r_j \in Y_R(l_i)$ by definition of $Y_R(l_i)$. Assume that $p \neq m$ otherwise we are done. According to Theorem 9 we have $r_{j+1} = \phi_{\text{back}}(r_j)$ for $p \leq j \leq m-1$. Furthermore $r_p = \phi_{\text{righ}}(l_i)$ which implies that r_p is the minimal bitangent of $R^2(l_i)$ (see Theorem 9). Using the fact that the bitangents of $R^2(l_i)$ are in the orbit of r_p under the operator ϕ_{back} (Theorem 9), we obtain $r_j \in R'(l_i)$ for $p+1 \leq j \leq m$. As $\phi_R(l_{i+1})$ is a right-xx bitangent, the bitangents $r_{p+1}, \ldots, r_m$ belong to $Y_R(l_i)$. The complexity of the procedure FLIP1 is therefore at most the sum of the following quantities:

1. $|Y_R(b)| + |rc(b)|$ to compute $\phi_R(b_R)$; this is the cost of the procedure PHI-R(b, b_R).

2. $|Y_L(b)| + |lc(b)|$ to compute $\phi_L(b_L)$; this is the cost of the procedure PHI-L(b, b_L).

3. $\sum_{1 \leq i \leq u-1} |lc(r_i)|$ to compute all the $\psi(r_i)$ such that $\psi(r_i)$ is an xx-right bitangent; this is the procedure PHI-L-NEW(b, r_{i+1}).

4. $|lc(r_u)| + q + |L'(r_u)|$ to compute $\psi(r_u)$; this is the procedure PHI-L-NEW(b, r_{u+1}).

5. $\sum_{1\leq i\leq q'-1} 1+|Y_R(l_i)|+|rc(l_i)|$ to compute all the $\psi(l_i)$; this is the procedure PHI-R-NEW(b, l_{i+1}) for $i = 1, \ldots, q'-1$

6. $p-u+1$ to compute $\psi(l_{q'}) = \phi(b)$.

For a bitangent t, there is at most one bitangent b such that $t \in rc(b)$. The same holds for $lc(b)$. Moreover the sum $\sum_b |rc(b)|$ is a $O(k)$. Theorem 19 connects $Y_R(b)$ to $\mathcal{G}$ so with the help of Theorem 2 we can claim that $|Y_R(b)|$ is amortized constant. It follows that the cost of FLIP1 is amortized constant. This closes the proof of Theorem 1.

5 Conclusion and further research

ONE WAY TO ADVANCE the science of computational geometry is to make a comprehensive study of primitive operations that are used in many different algorithms.
—D. E. Knuth, Axioms and Hulls, LNCS 606, 1992.

We have presented a new method to compute in optimal output sensitive time and linear space the visibility graph or visibility complex of a collection of disks of constant complexity. The method is built on top of the greedy flip algorithm of the second author of this paper and G. Vegter [39]; this is the first optimal algorithm for general disks (circles, ellipses, etc) that uses only simple data structures and only the counterclockwise or left-turn predicate for disks, once the disks are sorted with respect to some canonical direction. A weaker version of the method has been implemented and experimental results are discussed and reported in [3]. We believe that tuning a combinatorial model of the left-turn predicate on triplets of disks in the plane—as oriented matroids [7] or Knuth CCC-systems [26] are for the left-turn predicate on triplets of points in the plane— is an interesting line of research for the future from which our presentation might benefit as well as the implementation of our algorithms.

References

[1] J. Abello and K. Kumar. Visibility graphs and oriented matroids. In R. Tamassia and I. G. Tollis, editors, *Graph Drawing (Proc. GD '94)*, volume 894 of *Lecture Notes Comput. Sci.*, pages 147–158. Springer-Verlag, 1995.

[2] P. K. Agarwal, J. Basch, L. J. Guibas, J. Hershberger, and L. Zhang. Deformable free space tiling for kinetic collision detection. In B. Donald, K. Lynch, and D. Rus, editors, *New Directions in Algorithmic and Computational Robotics*, pages 83–96. A. K. Peters, 2001.

[3] P. Angelier. *Implémentation du Greedy Flip Algorithm.* PhD thesis, Ecole Normale Supérieure (Paris), 2001. To appear fall 2001.

[4] B. Aronov, J. Matoušek, and M. Sharir. On the sum of squares of cell complexities in hyperplane arrangements. *J. Combin. Theory Ser. A*, 65:311–321, 1994.

[5] T. Asano, S. K. Ghosh, and T. C. Shermer. Visibility in the plane. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 829–876. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.

[6] T. Asano, L. J. Guibas, and T. Tokuyama. Walking on an arrangement topologically. *Internat. J. Comput. Geom. Appl.*, 4:123–151, 1994.

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

[8] H. Brönnimann, L. Kettner, M. Pocchiola, and J. Snoeyink. Counting and enumerating pseudotriangulations with the greedy flip algorithm. In preparation, 2001.

[9] B. Chazelle, H. Edelsbrunner, M. Grigni, L. J. Guibas, J. Hershberger, M. Sharir, and J. Snoeyink. Ray shooting in polygons using geodesic triangulations. *Algorithmica*, 12:54–68, 1994.

[10] B. Chazelle, L. J. Guibas, and D. T. Lee. The power of geometric duality. *BIT*, 25:76–90, 1985.

[11] F. Cho and D. Forsyth. Interactive ray tracing with the visibility complex. *Computers and Graphics (Special Issue on Visibility - Techniques and Applications)*, 23(5):703–717, 1999.

[12] R. Connelly, E. D. Demaine, and G. Rote. Straightening polygonal arcs and convexifying polygonal cycles. In *Proc. 41th Annu. IEEE Sympos. Found. Comput. Sci.*, pages 432–442, 2000.

[13] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, Germany, 2nd edition, 2000.

[14] F. Durand, G. Drettakis, and C. Puech. Fast and accurate hierarchical radiosity using global visibility. *ACM Transactions on Graphics*, 18(2):128–170, 1999.

[15] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*, volume 10 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, Heidelberg, West Germany, 1987.

[16] H. Edelsbrunner and L. J. Guibas. Topologically sweeping an arrangement. *J. Comput. Syst. Sci.*, 38:165–194, 1989. Corrigendum in 42 (1991), 249–251.

[17] H. Edelsbrunner, J. O'Rourke, and R. Seidel. Constructing arrangements of lines and hyperplanes with applications. *SIAM J. Comput.*, 15:341–363, 1986.

[18] J. E. Goodman. Pseudoline arrangements. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 5, pages 83–110. CRC Press LLC, Boca Raton, FL, 1997.

[19] M. T. Goodrich and R. Tamassia. Dynamic ray shooting and shortest paths in planar subdivisions via balanced geodesic triangulations. *J. Algorithms*, 23:51–73, 1997.

[20] M. Hagedoorn, M. Overmars, and R. C. Veltkamp. A robust affine invariant similarity measure based on visibility. In *Abstracts 16th European Workshop Comput. Geom.*, pages 112–116. Ben-Gurion University of the Negev, 2000.

[21] D. Halperin. Arrangements. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 21, pages 389–412. CRC Press LLC, Boca Raton, FL, 1997.

[22] J. M. Henle and E. M. Kleinsberg. *Infinitesimal Calculus*. MIT Press, 1980.

[23] L. Kettner, D. Kirkpatrick, A. Mantler, J. Snoeyink, B. Speckmann, and F. Takeuchi. Tight degree bounds for pseudotriangulations of points. *Comp. Geom. Theory Appl.*, 2001. To appear in special issue.

[24] D. Kirkpatrick, J. Snoeyink, and B. Speckmann. Kinetic collision detection for simple polygons. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 322–329, 2000.

[25] D. Kirkpatrick and B. Speckmann. Separation sensitive kinetic separation for convex polygons. In *Proc. Japan Conf Disc. Comp. Geom.*, LNCS. Springer Verlag, 2000. To appear.

[26] D. E. Knuth. *Axioms and Hulls*, volume 606 of *Lecture Notes Comput. Sci.* Springer-Verlag, Heidelberg, Germany, 1992.

[27] J.-C. Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, Boston, 1991.

[28] P. MacMullen. Modern developments in regular polytopes. In T. Bisztriczky, P. McMullen, R. Schneider, and A. I. Weiss, editors, *Polytopes: Abstract, Convex and Computational*, pages 97–124. Kluwer Academic Publischers, 1994.

[29] J. S. B. Mitchell. Shortest paths and networks. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 24, pages 445–466. CRC Press LLC, Boca Raton, FL, 1997.

[30] J. S. B. Mitchell. Geometric shortest paths and network optimization. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 633–701. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.

[31] N. Nilsson. A mobile automaton: An application of artificial intelligence techniques. In *Proc. IJCAI*, pages 509–520, 1969.

[32] J. O'Rourke. Visibility. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 25, pages 467–480. CRC Press LLC, Boca Raton, FL, 1997.

[33] J. O'Rourke. Computational geometry column 39. *SIGACT News*, 31(3):47–49, 2000.

[34] J. O'Rourke and I. Streinu. Vertex-edge pseudo-visibility graphs: Characterization and recognition. In *Proc. 13th Annu. ACM Sympos. Comput. Geom.*, pages 119–128, 1997.

[35] J. O'Rourke and I. Streinu. The vertex–edge visibility graph of a polygon. *Comput. Geom. Theory Appl.*, 10:105–120, 1998.

[36] M. Pocchiola and G. Vegter. Order types and visibility types of configurations of disjoint convex plane sets (extended abstract). Technical Report 94-4, Labo. Inf. Ens, Jan. 1994.

[37] M. Pocchiola and G. Vegter. Minimal tangent visibility graphs. *Comput. Geom. Theory Appl.*, 6:303–314, 1996.

[38] M. Pocchiola and G. Vegter. Pseudo-triangulations: Theory and applications. In *Proc. 12th Annu. ACM Sympos. Comput. Geom.*, pages 291–300, 1996.

[39] M. Pocchiola and G. Vegter. Topologically sweeping visibility complexes via pseudo-triangulations. *Discrete Comput. Geom.*, 16:419–453, Dec. 1996.

[40] M. Pocchiola and G. Vegter. The visibility complex. *Internat. J. Comput. Geom. Appl.*, 6(3):279–308, 1996.

[41] M. Pocchiola and G. Vegter. On polygonal covers. In B. Chazelle, J. Goodman, and R. Pollack, editors, *Advances in Discrete and Computational Geometry*, volume 223 of *Contemporary Mathematics*, pages 257–268. AMS, Providence, 1999.

[42] D. Randall, G. Rote, F. Santos, and J. Snoeyink. Counting triangulations and pseudo-triangulations of wheels. In *Proc. 13th CCCG*, Univ. of Waterloo, Ont, 2001.

[43] S. Rivière. *Calculs de Visibilité dans un Environnement 2D.* PhD thesis, Universit'e Joseph Fourier, Grenoble, France, 1997.

[44] S. Rivière, R. Orti, F. Durand, and C. Puech. Using the visibility complex for radiosity computation. In *ACM Workshop Appl. Comput. Geom.*, May 1996.

[45] G. Rote, I. Streinu, and F. Santos. The expansion cone and the polytope of minimum pseudotriangulations. In preparation, 2001.

[46] I. Streinu. Stretchability of star-like pseudo-visibility graphs. In *Proc. 15th Annu. ACM Sympos. Comput. Geom.*, pages 274–280, 1999.

[47] I. Streinu. A combinatorial approach to planar non-colliding robot arm motion planning. In *Proc. 41th Annu. IEEE Sympos. Found. Comput. Sci.*, pages 443–453, 2000.