A sum of squares theorem for visibility complexes *

Extended Abstract

Pierre Angelier

Michel Pocchiola †

ABSTRACT

We present a new and simpler method to implement in constant amortized time the flip operation of the so-called 'Greedy Flip Algorithm', an optimal algorithm to compute the visibility complex of a collection of pairwise disjoint bounded convex sets of constant complexity. The method relies on a 'sum of squares' like theorem for visibility complexes stated and proved in this paper. (The sum of squares theorem for an 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

This paper is concerned with visibility among convex objects in the plane. 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 cell complex whose underlying topological space is the quotient of the space of rays by the equivalence relation whose classes are the sets of rays supported by the same line and such that the line segment joining the origins of the rays lies in free space, the complement of the objects [15]. The one-skeleton of the visibility complex is combinatorially equivalent to the (tangent) visibility graph of the objects; in particular its vertices correspond to the free bitangents of the objects. Several data structures based on the visibility complex have been designed to solve 2D geometric problems like planning the motion of a rod admist obstacles, computing the visible objects from a query point or from a (moving) query ray, computing ray-traced images at interactive rates or computing exact form-factors (see [15, 14, 13, 5, 17] and the references cited therein); visibility complexes of simple polygons have been elegantly characterized in the framework of oriented matroid theory as a class of bipartite graphs, closing partially the characterization/recognition problem of visibility graphs of simple polygons (see [10, 12]). The related notions of pseudotriangulation and greedy pseudotriangulation, introduced in [15, 14] to design efficient visibility complexes algorithms, have recently been used for computing depth orders in the plane [16], for collision detection of moving deformable objects in the plane (see [1, 11]), or yet for planning expansive motions of polygonal chains (see [6, 18] where rigidity properties of pseudotriangulations are established). A challenging question is to extend these notions and applications to 3D environments (see [7] for a first step.)

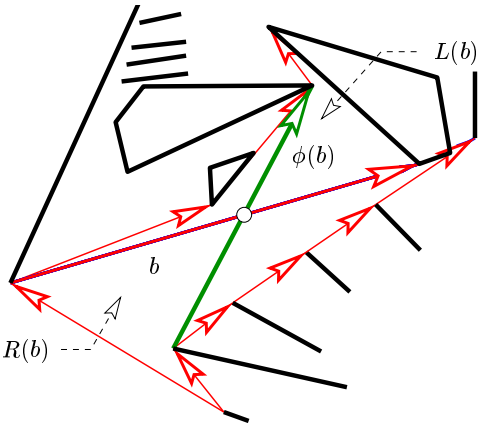

Figure 1: Flipping the bitangent b to obtain $\phi(b)$.

In this paper we focus on the algorithmic of visibility complexes. 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 (line segments, circles, ellipses, etc) [14]. The GFA realizes a topological sweep of the visibility complex with respect to some partial order $\prec$ defined on the set of upward (with respect to some canonical direction) directed bitangents. This partial order is the transitive closure of the χ_1 relations defined by

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

i.e., $b \prec b'$ if there exists a sequence $b = b_1, \ldots, b_r = b'$

*A preliminary version of this work was presented at the Dagstuhl-Seminar 99 102.

†Département d'Informatique, Ecole Normale Supérieure, 45 rue d'Ulm 75230 Paris, France.
Pierre.Angelier@ens.fr, Michel.Pocchiola@ens.fr

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
SCG'01, June 3-5, 2001, Medford, Massachusetts, USA.
Copyright 2001 ACM 1-58113-357-X/01/0006 ...$5.00.

of upward directed bitangents such that $\chi_1(b_i, b_{i+1})$ for $i = 1, \ldots, r-1$. The sweep structure S, some subcomplex of the visibility complex, is associated with a pseudotriangulation $G(S)$ 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 plus the flip of the corresponding bitangent b in $G(S)$, i.e., the computation of the bitangent b' that leaves the pseudotriangle T_R of $G(S)$ incident upon b on its right and enters the pseudotriangle T_L of $G(S)$ incident upon b on its left; it is known that

1. b' is the bitangent $\phi(b)$ crossing b from right-to-left and whose angle with b is minimal.
2. $T_R = b \cdots$ (resp. $T_L = b \cdots$), considered as the $\prec$-increasing sequence of bitangents that appear in its boundary, is the lexicographic smallest pseudotriangle incident upon b on its right (resp. left). T_R and T_L are simply denoted $R(b)$ and $L(b)$, and are said *greedy*. (See Figure 1.)

Besides linked structures to represent the incidence relations of the atoms of S and $G(S)$ the method described in [14] needs auxiliary data structures to implement efficiently the flip operation: two parts of the boundary of each pseudotriangle 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 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); in [14] it is shown that the flip operation requires a constant number of enqueue and split operations; the dequeue/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. Some comments on the bit-complexity of the predicates χ_1, χ_2 and χ_3 are in order. 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

$$a_1 + a_2\sqrt{b_2} + a_3\sqrt{b_3} + a_4\sqrt{b_2 b_3} \qquad (2)$$

where the a_i, b_i are polynomials in the inputs whose degrees depend on the predicate. In the worst case the evaluation of χ_3 requires the evaluation of the sign of an irreducible polynomial of degree 16 while χ_2 requires the evaluation of an irreducible polynomial of degree only 6. Therefore the bit-complexity 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) : a substantial difference from a practical point of view.

Since the visibility complex depends only on the χ_1 relations (and not on χ_2 nor χ_3) it is natural to ask whether or not it can be computed using only χ_1. A related question 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 is to introduce in the picture

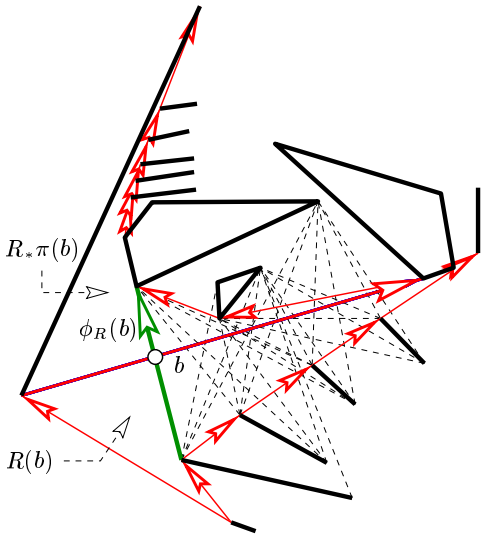

Figure 2: The bitangent $\phi_R(b)$.

1. The dual pseudotriangulation $G_*(S)$ 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.
2. The bitangent $\phi_R(b)$ 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 the bitangent that leaves $R(b)$ and enters the dual greedy pseudotriangle $R_*\pi(b)$ where π is the operator that changes the direction of a bitangent. See Figure 2.

THEOREM 1. *The flip operation of the GFA can be implemented in constant amortized time using only the predicate χ_1. 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 pseudotriangulations $G(S)$ and $G_*(S)$,*
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 $\phi_L(t) := \pi\phi_R\pi(t)$.*

The complexity analysis of our flip method is based on a kind of *sum of squares* theorem for visibility complexes. Recall that the 'sum of squares' theorem for an arrangement of lines in the plane 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 consequence of the linear

bound on the complexity of the so-called zone of a line in an arrangement of lines; see [4, 9] and [2] for an higher dimensional analogue). To state our sum of squares theorem we need to extend the definition of the ϕ operator: 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 b' crossing no bitangent in H, crossing b from right-to-left and whose angle with b is minimal. Note that $\phi(b; \emptyset) = \phi(b)$. Finally it is convenient to introduce the notation $\phi(a, b) := \phi(a; \{\phi_R(b), \phi_L(b)\})$. See Figure 3 for an illustration.

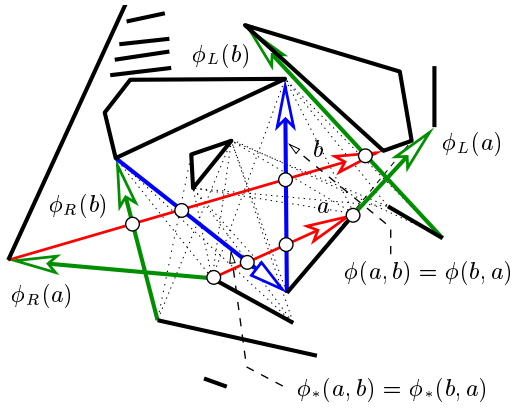

Figure 3: A pair (a, b) such that $\phi(a, b) = \phi(b, a)$ and $\phi_*(a, b) = \phi_*(b, a)$.

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 $\phi_* = \phi^{-1}$.)*

The paper is organized as follows. In Section 2 we recall the definition of the visibility complex of the scene composed of the disks plus the bitangents in H and we introduce notations for the incidence relations of its cells (the ϕ_{xxx} operators); then we relate the $\phi(b, H)$ operator to the visibility complex and we show that this operator satisfies the so-called Minimum Element, Right-to-Left and Flip Properties (Theorems 3,4,5) which were stated in [14] in the case where $H = \emptyset$. Finally we describe an optimal flip method that uses only the χ_2 predicate. 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). In section 3

1. we describe the greedy pseudotriangles in terms of the incidence relations of the visibility complex (Theorem 6) and establish a duality relation between the greedy pseudotriangles defined with respect to the partial order $\prec$ and the ones defined with respect to its dual partial order $\prec_*$ (Theorem 7);

2. we give several characterizations of the operator ϕ_R (Theorem 8), study the equation $\phi(b) = \phi_R(b)$ (Theorem 9), identify key properties of the predicate χ_1 (Theorem 10), and explain how $\phi_R(b)$ can be efficiently computed in the dynamic setting of the GFA using only the χ_1 predicate (Theorems 12, 13 and 11);

3. we prove the Sum of Squares Theorem and state a sufficient condition — used in our flip method — for a pair of (a, b) to satisfy the conditions of the Sum of Squares Theorem (Theorem 14).

In the fourth and last section we describe the optimal flip method based on the predicate χ_1 and analyze its complexity. Finally we mention that our technique, specialized to point obstacles (= circles with an infinitesimal radius), subsumes the topological sweep technique of Edelsbrunner and Guibas [8] and its extension developed in [3]. This point will be explained in a companion paper.

2. THE GREEDY FLIP ALGORITHM

The visibility complex. 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 with nonempty interiors and smooth boundaries (disks or obstacles for short); for the ease of the exposition we assume that the disks are in generic position; free space $\mathbb{F}$ is by definition the complement of the interiors of the disks. We recall that a bitangent is a closed directed line segment whose supporting line is tangent to two disks at its endpoints and that an arc is an open connected oriented piece of the boundary of a disk whose endpoints are endpoints of free bitangents. A *right-xx* (resp. *xx-right*) bitangent is a bitangent that leaves (enters) a disk on its right; similarly we introduce left-xx and xx-left bitangents. Let H be a subset of pairwise disjoint free bitangents and let H' be the set H augmented with the bitangents on the convex hull of the disks; $B_H(\supseteq H)$ denotes the set of free bitangents that cross no bitangent in 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 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 $\text{seg}_H(r)$ the (directed) line segment in $\mathbb{F}$ joining the origins $o(r')$ of the rays r' in the equivalence class of r; we denote by γ_{+i} (γ_{-i}) the curve of rays in $\mathbb{V}_H$ that leave tangentially the disk O_i on their left (right) and by γ_+ (γ_-) the curve of rays that leaves tangentially the convex hull of the disks on their left (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 $\text{seg}_H(r_i)$ 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 (vertices) are the intersection points of the

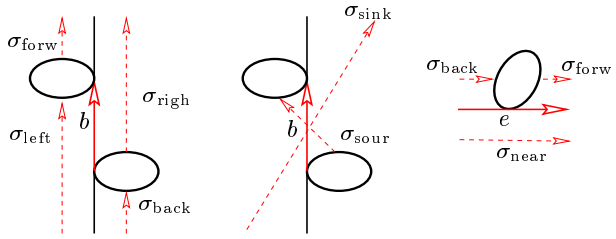

Figure 4: Topology around a vertex and an edge of the visibility complex.

curves γ_i. The map that associates with the vertex v the bitangent $\varepsilon(v) \in B_H$ included in the line segment $\text{seg}_H(v)$ is one-to-one.

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 (sink) cusp edge is a 1-cusp edge whose cusp point is realized by the source (sink) of the edge. An edge is said *positive (negative)* if it lies on some γ_{+i} (γ_{-i}).

3. its 2-cells (faces) are the connected components of $\mathbb{V}_H$ minus the vertices and edges; we denote by $\sigma_{\pm\text{hull}}$ the 2-cells of rays whose supporting lines let the disks on their left (right). The boundary of a 2-cell $\sigma \neq \sigma_{\pm\text{hull}}$ is composed of two chains — 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 (follows) an edge is the sink (source) of this edge — 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 sour(σ) the supporting line of a ray $r \in rc(\sigma)$ lies on the right side of the supporting line of the corresponding parallel ray $r' \in lc(\sigma)$. Walking along the chain $rc(\sigma)$ we find successively three convex subchains denoted $rc_{\text{back}}(\sigma)$, $rc_{\text{left}}(\sigma)$ and $rc_{\text{forw}}(\sigma)$.

A vertex v is incident to eight or six 2-cells depending on whether the bitangent $b := \varepsilon(v) \in H$ or not. These 2-cells are defined as follows : σ_{righ} (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 (left) side of the supporting line of v; similar constructions done with the segments $\text{seg}_{\text{back}}, \text{seg}_{\text{forw}}$ left after removing from $\text{seg}_H(v)$ the bitangent b, where seg_{back} comes before seg_{forw} along $\text{seg}_H(r)$ give rise to two new cells denoted σ_{back} and σ_{forw}; for $b \notin H$ the ray with origin the middle point of b and with direction a direction slightly larger (smaller) then the direction of v defines a 2-cell, denoted $\sigma_{\text{sour}}(\sigma_{\text{sink}})$ whose source (sink) is v. See Figure 4 for an illustration. The sink of the cell $\sigma_{\text{xxx}}(b; H)$ is denoted $\phi_{\text{xxx}}(b; H)$. One can check that the bitangents a of $rc_{\text{back}}(\sigma)$, $rc_{\text{left}}(\sigma)$ and $rc_{\text{forw}}(\sigma)$ are characterized by $\phi_{\text{back}}(a; H), \phi_{\text{left}}(a; H), \phi_{\text{forw}}(a; H) = \phi(b; H)$ where b is the source of σ. For $b \in H$ a similar construction gives rise to two 2-cells whose source (sink) is v. An edge e is incident to one or three 2-cells depending on whether e lies on some γ_p or some γ_i. If $e \subset \gamma_i$ these 2-cells are denoted $\sigma_{\text{near}}(e)$, $\sigma_{\text{forw}}(e)$, and $\sigma_{\text{back}}(e)$ as indicated in Figure 4. For $b \notin H$ we set $rc(b) = rc(\sigma_{\text{sour}}(b))$.

The minimum element, right-to-left and flip 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 $\sim^1_H$ defined by : $x \sim^1_H x'$ 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 whose proofs appear in [14] in the case where H is the empty set; up to some minor modifications the proofs are also valid in the case $H \neq \emptyset$.

THEOREM 3 (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*

$$\phi(b; H) = \begin{cases} \phi_{sink}(b; H) & \text{if } b \notin H' \\ \pi(b) & \text{otherwise,} \end{cases} \tag{3}$$

and if $H_1 \subset H_2$ then $\phi(b; H_1) \preceq \phi(b; H_2)$.

Each filter I of $(X_0(H), \prec_H)$ is associated with

1. its so-called (maximal) antichain $\hat{I}$ defined as the set of 1- and 2-cells such that $\text{sink}(\sigma) \in I$ but not $\text{sour}(\sigma)$. $\hat{I}$ contains one edge $e \subset \gamma$ and its incident faces per curve $\gamma \in \{\gamma_i, \gamma_p\}$, and for each face $\sigma \in \hat{I}$ an edge $\sigma^+ \in lc(\sigma)$ and an edge $\sigma^- \in rc(\sigma)$. For σ a face $\hat{I}$ the bitangent $b := \text{sink}(\sigma)$ is said to be right (left) minimal in I if $b = \text{sink}(\sigma^-)$ ($b = \text{sink}(\sigma^+)$).

2. its so-called greedy pseudotriangulation

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

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 $\min_{\prec_H} I$.

THEOREM 4 (RIGHT-TO-LEFT PROPERTY). *Let $b \in G(I)$ and $t \in I$ that crosses b. Then $b \prec_H t$.*

THEOREM 5 (FLIP PROPERTY).

$$G(I) = I \setminus \phi(I; H) \subset sink(\hat{I}). \tag{4}$$

Furthermore $b \in G(I)$ 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 pseudotriangles. For $b \in G(I)$ we denote by $R(b, I)$ the chain in the one-skeleton of $\mathbb{W}_{G(I)}$ whose envelope is the boundary of the pseudotriangle T incident upon $\varepsilon(b)$ and—locally—to the right of $\varepsilon(b)$, whose source is the minimal bitangent a in $G(I)$ with $\varepsilon(a)$ in the boundary of T, and whose sink is $\pi(a)$. By a slight abuse of language $R(b, I)$ is called a pseudotriangle 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 in clockwise order around the boundary of $R(b; H)$, starting from the tail of 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 $H = \emptyset$.

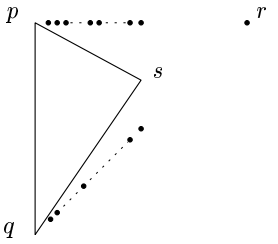

Figure 5: In the configuration $\{(i, 0), (i, i - n)\}_{1 \le i \le n/2-1}$ augmented with the point $r = (n, 0)$ and the triangle joining the three points $p = (0, 0), q = (0, -n)$ and $s = (n/2 - 1, -n/4 - 1/2)$ we have $\sum_b |R^{2_1}(b)| = \Omega(n^2)$ but $|B| = O(n)$.

The Greedy Flip Algorithm (GFA). The GFA maintains a representation of the antichain $\hat{I}$ and its associated greedy pseudotriangulation $G(I)$ 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 two linked structures to represent the atoms $\hat{I}$ and $G(I)$ and their incidence relations. Furthermore an atom of $\hat{I}$ points to its corresponding bitangent in $G(I)$. Finally we access to the linked structure via the set of 'minimal' faces of $\hat{I}$, 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. We observe, after [14], that b' leaves either the chain $R^2(b)$ or the first atom of $R^3(b)$. This motivates the introduction of the chain $R^{\text{lea}}(b)$ as the chain $R^2(b)$ augmented with the first atom of $R^3(b)$ if this atom is an arc. So since b' leaves $R^{\text{lea}}(b)$ and enters $L^{\text{ent}}(b)(:= \pi^{-1} R^{\text{lea}} \pi(b))$ 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 (entering) the disk supporting the arc r (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** $(\text{Top}(C), \text{Top}(C'))$

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

The complexity of FLIP2 is clearly proportional to the sum of the sizes of $R^{2_1}(b)$ and $L^{2_1}(b)(:= \pi^{-1} R^{2_1} \pi(b))$, where $R^{2_1}(b)$ is the prefix subchain of the chain $R^2(b)$ cut at $\phi(b)$. As illustrated in Figure 5 the average value of the size of $R^{2_1}(b)$ on the whole set B of free bitangents is not constant. However we will show in the next section that

$$\sum_{b:\phi(b) \in B_{rx}} |R^{2_1}(b)| = O(|B|) \tag{5}$$

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 $(R^{\text{lea}}, L^{\text{ent}})$ we run it in tandem 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.

3. TECHNICAL DEVELOPMENTS

We push one step further the study of the combinatorial properties of the visibility complexes. We introduce the dual partial order $\prec_*$, the ϕ_R operator, and we prove the Sum of Squares Theorem.

3.1 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 $\prec_*$) pseudotriangulation associated with the filter J of $(X_0, \prec_*)$ (Flip Property), and $R_*(b)$ and $L_*(b)$ are the dual greedy pseudotriangles 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 $I_* = X_0 \setminus I$.

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 pseudotriangle $R(a')$ 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 $R^i(a)$, $R^{i_j}(a)$, $\pi(e)$ and their left versions — here 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 $\phi(a)$ for $i = 2, 3$. See Table 1. Similar relations hold in the dual setting. We use these relations to prove by the two following theorems which will be very useful in our subsequent analysis.

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

Table 1: Relations between two greedy pseudotriangles of the bitangents a and a' where a and a' are the source and the sink of a positive edge e.

THEOREM 6 (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 that*

1. *$b_{11} = b$, $b_{21} = \phi_{righ}(b)$ and $b_{31} = \phi_{back}(b)$*
2. *$b_{1,j+1} = \phi_{forw}(b_{1j})$ and the e_{1i} are arcs in X_1.*
3. *$b_{2,i+1} = \phi_{back}(b_{2i})$ and $b_{3,i+1} = \phi_{forw}(b_{3i})$*
4. *$R^4(b) = e_{40}\pi(b)$ and $\phi(b)$ leaves $R^2(b)e_{30}$*

PROOF. Claim 1 : We first show that $b_{31} = \phi_{\text{back}}(b)$. Let $b_1 = (\phi_{\text{back}})_*(b)$, $\sigma := \sigma_{\text{sour}}(b_1)$, and let $b_1e_1b_2e_2\ldots$ be the chain $rc_{\text{forw}}(\sigma)$. 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 1, $R^3(b_2) = R^{3_1}(b_1)\phi(b_1)L^{2_2}(b_1)$ 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 $b' = \phi_{\text{righ}}(b)$.

Claims 2,3,4 and 5 are direct applications of the Right-to-Left Property. □

THEOREM 7 (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 $R^{\text{lea}}(b)$ then b belongs to or leaves $L_*^{\text{lea}}(a)$,
(5-) if a belongs to or enters $R^{\text{ent}}(b)$ then b belongs to or leaves $R_*^{\text{lea}}(a)$,
(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 bitangents 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 Minimum Element 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 or are tangent to $R^1(b)$ or $R^{\text{lea}}(b)$. 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. a and a' are the source and the sink of an edge e of the visibility complex. Assume for example that a' belongs to or leaves $R^{\text{lea}}(b)$ and that the edge e is a positive source cusp or a negative sink cusp 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); the induction hypothesis gives b belongs to or leaves $L_*^{\text{lea}}(a)$; according to the dual version of Table 1 $L_*^2(a) = R_*^{3_1}(a')\phi_*(a')L_*^{2_2}(a')$, $L_*^3(a) = L_*^{3_2}(a') \oplus e$ (pos. source cusp) or $L_*^3(a) = L_*^{3_2}(a')$ (neg. sink cusp). 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 $L_*^{\text{lea}}(a')$. The other cases are similar. □

3.2 The ϕ_R operator

THEOREM 8. *For $b \in X_0$ not on the convex hull we define $\phi_R(b)$ to be the maximal bitangent in the range $[b, \pi(b)]$ that leaves $R^{lea}(b)$. We claim that*

1. *$\phi_R(b)$ is the bitangent that leaves $R(b)$ and enters $R_*\pi(b)$,*
2. *$\phi_R(b)$ is the bitangent crossing b whose crossing point with b is the closest to the tail of b.*

Similarly we introduce the maximal bitangent $\phi_L(b)$ that enters $L^{ent}(b)$. Note that $\pi\phi_R = \phi_L\pi$.

THEOREM 9 (EQUATION $\phi_L(b) = \phi(b)$). *The three following assertions are equivalent:*

1. *$\phi(b) \neq \phi_L(b)$.*
2. *$\phi_L(b)$ crosses $\phi_{righ}(b)$.*

3. $\phi_L(b)$ does not leave the first atom of $R^2(b)$ or $\pi R^2_*(b)$.

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

Before going further in our analysis we state some properties of the χ_1 predicate on a collection of 2 disks and a point (= circle with an infinitesimal radius) that will be useful in the sequel. 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. A *signed* disk is a disk with an orientation of its boundary. O_{+i} (resp. O_{-i}) is the disk O_i oriented counterclockwise (resp. clockwise).

THEOREM 10. *Consider two disjoint disks O_1, O_2 and a point O_0, and for $i, j \in \{\pm 0, \pm 1, \pm 2\}$ let b_{ij} be the bitangent that leaves the signed disk O_i and enters the signed disk O_j. Assume that b_{0i}, b_{0j} are free and that $\chi_1(b_{0i}, b_{0j})$. Then the following sequences are χ_1-sequences:*

1. *$b_{0i} b_{0j} b_{-ij}$ for $(i, j) \neq (-1, +2)$.*
2. *$b_{0i} b_{ij} b_{-ij}$ for $i = -1, j = \pm 2$.*
3. *$b_{0j} b_{-ij} b_{-i-j}$ for $i = \pm 1, j = +2$.*

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 11. *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 neg. two cusp edge then $\phi_R(b_R) = \pi(b)$.*

We now examine the case where the sweep structure sweeps a bitangent b which is the source of a positive 0-cusp edge e. We need to relate the chain $rc(b)$ to the bitangent $\phi_R(b_R)$ where b_R is the sink of e.

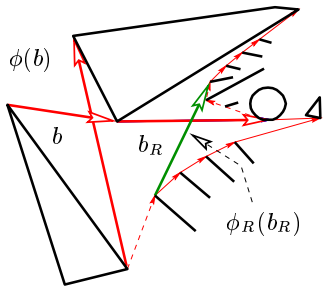

Figure 6: The bitangent $\phi_R(b_R)$ is a *bridge* between $R^{\mathbf{lea}}(b)$ and $R_*\pi(b_R)$.

THEOREM 12 (BRIDGE THEOREM). *Let e be a positive 0-cusp edge with source b and sink b_R where b is minimal in the filter I, let $b_R S(b)$ be the prefix subchain of $R_*\pi(b_R)$ (for $\prec$) with sink $\phi_R(b_R)$, and let $T(b)$ be the suffix subchain of $R^{lea}(b_R)$ with source $\phi_R(b_R)$ (its sink is $\phi(b)$). We claim that $R^{lea}(b_R)$ is a prefix subchain of $R^{lea}(b)$ and that*

$$rc(b) = beS(b)\phi_R(b_R)T(b)\phi(b). \tag{6}$$

Furthermore the chain $S(b)$ is a subchain of $R_*(\phi_*(b_R), I_*)$.

THEOREM 13. *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 $\pi R^{lea}_*(\phi_*(b_R), I_*)$ 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 */
1 $C \leftarrow R^{\text{lea}}(b)$; $C' \leftarrow R^{\text{ent}}_*\pi(b_R)$ in reverse order;
2 χ_1-WALK(C, C');
3 **return** (Top(C), Top(C'))

Procedure χ_1-WALK(C, C' : LIST OF ARCS)
0 **while** $C \neq \emptyset$ and $C' \neq \emptyset$
1 $r \leftarrow$ Top(C); $l \leftarrow$ Top(C'); $t \leftarrow$ bit(r, l);
2 **if** $\chi_1(r^+, t)$ $C \leftarrow$ Pop(C)
3 **elseif** $\chi_1(l^+, t)$ $C' \leftarrow$ Pop(C')
4 **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 12. 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 $\chi_1(\phi, r_g^+)$.

We first prove claim 1. Let O_i and O_j be the signed disks supporting the arcs r_g and l, and let O_0 be the intersection point of the supporting lines of l^+ and r_g^+. Using the notations of Theorem 10 we can write that $t = b_{ij}$, $\vec{l}^+ = \vec{b}_{0j}$, $\vec{r}_g^+ = \epsilon\vec{b}_{0,\epsilon i}$ where $\epsilon = \pm 1$ depending on the relative positions of r_g^+ and O_0 on the supporting line of r_g^+, and Claim 1 states that

$$\det(\vec{b}_{ij}, \epsilon\vec{b}_{0,\epsilon i}) > 0 \quad \text{and} \quad \det(\vec{b}_{0j}, \vec{b}_{ij}) > 0$$

under the assumption that $\det(\vec{b}_{0j}, \epsilon\vec{b}_{0,\epsilon i}) > 0$. The proof can be completed by a case analysis using Theorem 10. The proofs of claims (2-),(3-) and (4-), omitted from this version, are similar. □

3.3 Proof of the sum of squares theorem

We split the proof into several lemmas. Let $R'(b)$ be the set of bitangents $t \neq \phi_{\text{righ}}(b)$ of the chain $R^{2_1}(b)$ not in $\sigma_{\text{sour}}(b)$ and let $r(b) := rc_{\text{left}}(\sigma_{\text{righ}}(b))$. Let $Z_R(b)$ be the set of bitangents in $R'(b), r(b)$ or $R'_*(b)$.

LEMMA 1. *Let $a \in Z_R(b)$ and assume that $\phi_R(b)$ and $\phi_L(b)$ do not leave the same disk. Then $R(a, b) = R(a)$ and $\phi(a, b) = \phi(a; \{\phi_L(b)\})$. Furthermore if $b \in X_0(H(a))$ and $\phi_L(b) \neq \phi(b)$ then one the three following situations holds:*

1. $\phi(a,b) \prec \phi(b)$ *or, equivalently,* $\phi(a,b) = \phi_R(a)$. *In which case* $\phi(b)$ *leaves the disk that* $\phi(a,b)$ *enters.*

2. $\phi_L(b) \prec \phi(a,b)$ *or, equivalently,* $\phi(a,b) = \phi(a)$. *In which case* $\phi_L(b)$ *leaves the disk that* $\phi(a,b)$ *enters.*

3. $\phi(b) \prec \phi(a,b) \prec \phi_L(b)$ *or, equivalently,* $\phi(a,b) = \phi(b,a)$.

LEMMA 2. *Let* a, b *non crossing with* $\phi(a,b) = \phi(b,a)$. *Then:*

1. $\phi_R(b) \neq \phi_L(b)$ *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. $b \prec \phi(a)$ *and* $a \prec \phi(b)$.

4. *if* $a \in Z_R(b)$ *enters a disk on its right then* $\phi_{forw}(a,b)$ *enters* $L(b)$ *or* $\phi_L(b) \prec \phi_{forw}(a,b)$.

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 3. *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

$$\phi_L(b) \prec \pi(b) \prec \phi_L(b') \prec \pi(b'). \tag{7}$$

Let $t_0 = \phi(b,a)$. Since t_0 enters $L(b,a)$ and since $L(b,a) = L(b)$ (see Lemma 1) 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 $t_{i+1} = \phi_{\mathrm{forw}}(t_i)$ 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 $t'_{i+1} = \phi_{\mathrm{forw}}(t'_i)$ 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 (7). It remains to prove that $\phi_L(b) \neq \phi_L(b')$. If $\phi_L(b) = \phi_L(b')$ then according to Lemma 1 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

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

According to Lemma 2 Claim 4 the bitangent $\phi_{\mathrm{forw}}(a',b')$ enters $L(b')$ or $\phi_L(b') \prec \phi_{\mathrm{forw}}(a',b')$. In the latter case we are done since we have seen that $\pi(b) \prec \phi_L(b')$. Assume now that $\phi_{\mathrm{forw}}(a',b')$ enters $L(b')$, and consequently enters $C(b')$. We claim that

$$\phi_L(b) \prec \phi_{\mathrm{forw}}(a',b') \tag{8}$$

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

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 (8) follows.

Claim 1. : We observe that

$$\phi_*(a) \prec \pi_*\phi_L(b') \prec \phi(a;\phi_L(b)) \tag{9}$$

— the first inequality follows Claim (7) and Lemma 2 Claim 3 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 $L^{\mathrm{ent}}(b')$ 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 2 Claim 3 applied to ϕ_*).

Claim 4. : We observe that $\phi_L(a',b')$ enters $L^{\mathrm{ent}}(a',b')$ and that $\phi_{\mathrm{forw}}(a',b')$ is the sink of $L^{\mathrm{ent}}(a',b')$ (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 3 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

$$\phi(a',b') \prec \phi_{\text{forw}}(a',b') \preceq \phi(a'',b') \tag{10}$$

— 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 3 that $\pi(b) \prec \phi(a'',b')$, and consequently $a'' \notin E(b')$. This proves that $a' = \min_{\prec} E(b')$. □

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 sequel this set is closely related to the the multiset of bitangents visited during the course of our flip method. We denote by b_R the successor of b in the pseudotriangle $R(b)$.

Theorem 14. *Let $Y_R(b)$ be the set of bitangents a such that one of three following conditions is satisfied*

1. *$a \in R'(b)$ and $\phi_R(b_R)$ is a right-xx bitangent*
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*
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 $(a,b) \in \mathcal{G}$.

4. OPTIMAL FLIP USING χ_1

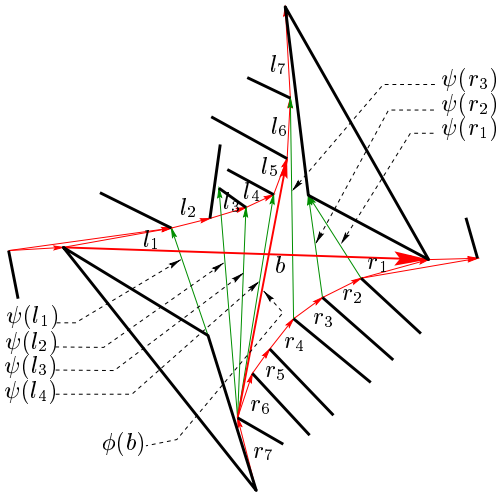

Figure 7: Flipping b using only χ_1. Here $r_u = r_3$, $l_5 = l_q$, $r_6 = r_p$ and $\psi(r_6) = \psi(l_5) = \phi(b)$.

We maintain as described previously not only $G(I)$ and $\hat{I}$ but also the dual pseudotriangulation $G_*(I_*)$; furthermore for each bitangent $t \in G(I)$ which is right (left) minimal we maintain a pointer to the arc of $G(I)$ that $\phi_R(t)$ leaves ($\phi_L(t)$ enters) and a pointer to the arc of $G_*(I_*)$ that $\phi_R(t)$ enters ($\phi_L(t)$ 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 9 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 (enter) the first atom of $R^2(b)$ ($L^2(b)$). If b and b_R are separated by a cusp then b_R is left minimal in $I' := I \setminus b$ and $\phi_L(b_R) = \phi_L(b)$ (Theorem 11). 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 (Theorem 13). 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. 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 pseudotriangles $R(b)$ and $R_*(l_i)$, for $i \leq q$, are incident to the same bitangent (b) and they don't (locally) overlap. It follows 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 whose proof is similar to the proof of Theorem 13.

Theorem 15. *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 $\pi R_*^{lea}(\phi_*(l_{i+1}), I_*)$.*
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 $\pi R_*^{lea}(\phi_*(l_{q+1}), I_*)$.*

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 $C \leftarrow [r_u, \ldots, r_{\max}]$; $C' \leftarrow \pi R_*^{\text{lea}}(\phi_*(l), I_*) \oplus c$;
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 15). 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 13, 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 13 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 6 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 6). Using the fact that the bitangents of $R^2(l_i)$ are in the orbit of r_p under the operator ϕ_{back} (Theorem 6), 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 14 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.

5. REFERENCES

[1] P. K. Agarwal, J. Basch, L. J. Guibas, J. Hershberger, and L. Zhang. Deformable free space tiling for kinetic collision detection. In *Proc. 4th Workshop Algorithmic Found. Robot.*, 2000. To appear.

[2] 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.

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

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

[5] 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.

[6] 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.

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

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

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

[10] S. K. Ghosh. On recognizing and characterizing visibility graphs of simple polygons. *Discrete Comput. Geom.*, 17:143–162, 1997.

[11] 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.

[12] 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.

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

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

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

[16] 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.

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

[18] 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.