A 'Sum of Squares' Theorem for Visibility Complexes

Pierre Angelier and Michel Pocchiola

G - LIENS - 00 - 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

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 graph/complex of a collection of pairwise disjoint bounded convex sets of constant complexity. The method uses only the incidence structure of the visibility complex and only the primitive predicate which states that the angle of a first bitangent is less than the angle of a second bitangent, both bitangents being tangent to the same convex. 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

Results of the paper. 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 [1]. 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 of 2D environments at interactive rates or computing exact form-factors in 2D environments (see [1, 2, 3, 4, 5] 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 [6, 7]). The related notions of pseudotriangulation and greedy pseudotriangulation, introduced in [1, 2] to design efficient visibility complexes algorithms, have recently been used for collision detection of moving deformable objects in the plane (see [8, 9]), or for planning non-colliding motions of linkages via monotone motions (see [10] where rigidity properties of greedy pseudotriangulations are established). A challenging question is to extend these notions and applications to 3D environments (see [11] for a first step.)

In this paper we focus on the algorithmic of visibility graphs/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) [2]. The GFA realizes a topological sweep of the visibility complex with respect to the partial order $\prec$ defined on the set of upward directed bitangents as the transitive closure of the χ_1 relations which are defined as follows

$$\chi_1(b_1, b_2) \iff \begin{cases} b_1 \text{ and } b_2 \text{ are tangent} \\ \text{to the same disk and} \\ \det(\vec{b_1}, \vec{b_2}) > 0. \end{cases} \quad (1)$$

The sweep structure S (some subcomplex of the vis-

*Département d' Informatique, Ecole normale supérieure, 45 rue d'Ulm 75230 Paris Cedex 05, France (pocchiola@ens.fr, angelier@ens.fr). A preleminary version of this work was presented at the Dagstuhl-Seminar 99 102.

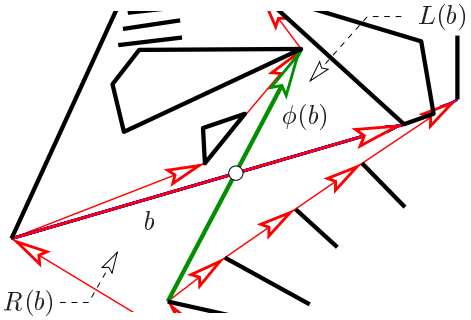

Figure 1: Flipping the bitangent b to obtain $\phi(b)$.

ibility complex) is associated with a greedy 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 some canonical direction. The sweep of a vertex of the visibility complex boils down to some local change in the sweep structure plus — and this is algorithmically the difficult part — the flip of the corresponding bitangent b in $G(S)$, i.e. the computation of the (sole) bitangent that leaves the greedy pseudotriangle $R(b)$ of $G(S)$ incident upon b on its right and enters the greedy pseudotriangle $L(b)$ of $G(S)$ incident upon b on its left (see Figure 1). Besides linked structures to represent the incidence relations of the atoms of S and $G(S)$ the method described in [2] 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 [2] 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.

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 can be computed using only χ_1. We answer affirmatively this question: thanks to an amortized analysis based on a kind of "sum of squares" theorem for visibility complexes stated and proved in this paper we show that the flip operation of the GFA can be implemented in constant amortized time using only the predicate χ_1, furthermore the only data structure used is the linked structure to represent the atoms (and their incidences) of the sweep structure S and its associated pseudotriangulation $G(S)$. These two features simplify the implementation of the GFA and should improve its practical performance. 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. Finally we mention that our technique, specialized to point obstacles (= circles with an infinitesimal radius), subsumes the topological sweep technique of Edelsbrunner and Guibas [12] and its extensions developed in [13, 14]. This point will be explained in a companion paper.

Organization of the paper. In Section 2 we extend the theory behind the GFA developed in [2] for convex obstacles to the case where the set of obstacles is augmented with a set H of pairwise disjoint free bitangents[1] : we start with the definition of the visibility complex whose set of vertices is the set B_H of free bitangents crossing none of the bitangents in H; we introduce the operator $b \mapsto \phi(b; H)$ which associates with the bitangent $b \in B_H$ the bitangent $b' \in B_H$ crossing b from right to left and with minimal angle with respect to the direction of b (see Figure 3), we show that this operator satisfies — with respect to the partial order $\prec_H$ defined as the tran-

[1] This setting is already used in [1]

sitive closure of the χ_1 relations on (some covering space of) B_H — the so-called minimum element and flip properties stated in [2] in the case where $H = \emptyset$ (Theorems 1 and 2), and 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**.) Section 3 is devoted to the statement and proof of our "sum of squares" theorem : we introduce in the picture the dual order $\prec_*$ of $\prec$ ($:=\prec_\emptyset$), the dual versions $\phi_*, R_*, L_*, \ldots$, of the operators $\phi, R, L, \ldots$ (Theorems 3,4) and the operator ϕ_R which associates with a directed bitangent b 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 (see Figure 5 and Theorem 5); then we state (and prove) our "sum of squares" theorem under the cryptic form: the number of pairs of free bitangents (a, b) such that $\phi(a, b) = \phi(b, a)$ and $\phi_*(a, b) = \phi_*(b, a)$ — here $\phi(u, v) := \phi(u; H(v))$ where $H(v)$ is the set of bitangents $\phi_R(v)$ and $\phi_L(v) := \pi^{-1}\phi_R\pi(v)$ (π is the operator that changes the direction of a bitangent) — is a big-O of the number of free bitangents (see Figure 6 and Theorem 6), and we 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 7). Finally we explain how $\phi_R(b)$ can be computed efficiently (in the dynamic setting of the GFA) using only the χ_1 predicate (Theorem 10). In the fourth and last section we describe the optimal flip method based on the predicate χ_1 and analyze its complexity.

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. 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 $\mathrm{seg}_H(r)$ the (directed) line segment in $\mathbb{F}$ joigning 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) — we will say that " ... leave the signed disk O_{+i} (O_{-i})" —; 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-dimensionnal cell complex on $\mathbb{V}_H$:

(1-) its 0-cells (vertices) are the intersection points of the curves γ_i. The application 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.
(2-) its 1-cells (or edges) are the connected components of the curves γ_i, γ_p minus the vertices. The application 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 application 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.
(3-) its 2-cells (faces) are the connected components of $\mathbb{V}_H$ minus the vertices and edges; we denote by $\sigma_{\pm\mathrm{hull}}$ the 2-cells of rays whose supporting lines let the disks on their left (right)[2]. The boundary of a 2-cell $\sigma \neq \sigma_{\pm\mathrm{hull}}$ is composed of two chains — a chain

[2] The cell $\sigma_{\epsilon\mathrm{hull}}$ is homeomorphic to a 2-disk minus a point.

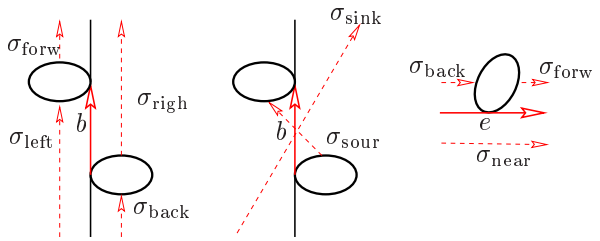

Figure 2: Topology around a vertex and an edge of the visibility complex.

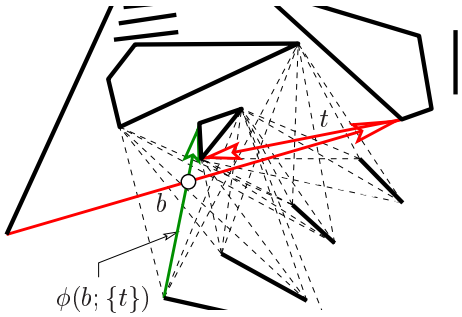

Figure 3: The bitangents crossing b are drawn dashed.

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 $\text{sour}(\sigma)$ the supporting line of a ray $r \in rc(\sigma)$ lies on the right side of the supporting line of the corresponding parallel ray $r' \in lc(\sigma)$.

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 2 for an illustration. The sink of the cell $\sigma_{\text{xxx}}(b;H)$ is denoted $\phi_{\text{xxx}}(b;H)$. For $b \in H$ a similar construction give 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 2. For $b \notin H$ we set $rc(b) = rc(\sigma_{\text{sour}}(b))$.

The minimum element and flip properties. For technical reason it is 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 (signed) 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 two following theorems whose proofs appear in [2] in the case $H = \emptyset$; up to some minor modifications the proofs are also valid in the case $H \neq \emptyset$.

Theo. & Def. 1 (Minimum Element Property). *The application 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 (1) $b \prec_H t$, and (2) t crosses b is well defined and one-to-one. Furthermore (1-) $\phi(b;H) = \phi_{sink}(b;H)$ if $b \notin H'$ and $\phi(b) = \pi(b)$ otherwise, and (2-) if $H \subset H'$ then $\phi(b;H) \preceq \phi(b;H')$.*□

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 (interior) 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 next theorem relates the greedy pseudotriangulation $G(I)$ to the ϕ operator and identifies

the minimal bitangents of the filter I:

Theorem 2. *Let I be a filter of $(X_0(H), \prec_H)$. Then (1-)* **(Flip Property).** *$G(I) = I \setminus \phi(I; H) \subset sink(\hat{I})$. (2-)* **(Minimal elements Property).** *$b \in G(I)$ 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 (recall that a pseudotriangle is a simply connected region of the plane whose boundary consists of a simple closed curve with three cusp points whose convex hull contains the pseudotriangle). 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 $H = \emptyset$.

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 [2], that b' leaves either the chain $R^2(b)$ or the first atom of $R^3(b)$ which in that case is an arc. This motivates the introduction of the chain $R^{\text{lea}}(b)$ as the chain $R^2(b)$ augmented with the first atom, denoted $aR^3(b)$, 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 (and simple) 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 $\phi(b)$. This leads to the following piece of code

Function FLIP2(b : BITANGENT) : PAIR OF ARCS
Comments/* 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** (Top(C), Top(C'))

Function χ_2-WALK(C, C': LIST OF ARCS)
0 **while** $C \neq \emptyset$ and $C' \neq \emptyset$
1 $r \leftarrow$ Top(C); $l \leftarrow$ Top(C'); $\phi \leftarrow$ bit(r, l);
Comments/* bit(r, l) is the bitangent leaving (entering) the signed disk supporting the arc r (l);*/ r^+ (l^+) is the sink of $r(l)$ */
2 **if** (ϕ leaves r and enters l) **exit**;
3 **if** $\chi_2(r^+, l^+)$ $C \leftarrow$ Pop(C) **else** $C' \leftarrow$ Pop(C');
4 **return** 0;

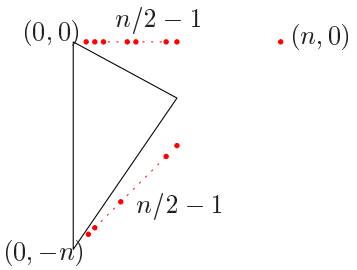

Figure 4: In the configuration $\{(i, 0), (i, i-n)\}_{1 \leq i \leq n/2-1}$ augmented with the point $(n, 0)$ and the triangle joining the three points $(0, 0), (0, -n)$ and $(n/2 - 1, -n/4 - 1/2)$ we have $\sum_b |R'(b)| = \Omega(n^2)$ but $|B| = O(n)$.

Let $R'(b)$ be the set of bitangents t of the chain $R^2(b)$ not in $rc(b)$ such that $b_0 \prec t \prec \phi(b)$ where b_0 is the minimal bitangent (if any) of the chain $R^2(b)$. The complexity of FLIP2 is clearly proportional to the sum of the sizes of $R'(b)$ and $L'(b)(:= \pi^{-1} R' \pi(b))$ plus a constant amortized quantity. It turns out that in general these quantities are not constant amortized. However we will show in the next section that this quantity is constant amortized on the set of bitangents b such that $\phi(b)$ leaves a disk on their right (we say that $\phi(b)$ is a right-xx bitangent):

$$\sum_{b \in B:\, \phi(b) \text{right-xx}} |R'(b)| = O(|B|). \tag{3}$$

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).

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 state our 'sum of squares' theorem.

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 $i = 2, 3$. For example if e is a positive 0-cusp edge then $R(a') = R(\phi(a), I \setminus a)$ and $R^1(a) = aeR^1(a')$, $R^2(a') = R^{2_1}(a)$, $R^3(a') = R^{3_1}(a)\phi(a)L^{2_2}(a)$, $R^4(a') = L^{3_2}(a)\pi(e)$. Similar relations hold in the dual setting. We use these relations to prove by induction the two following theorems which will be very useful in our subsequent analysis.

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

Theorem 4 (Duality Theorem). *We claim that $a \in R^{2_1}(b)$ (resp. $R'(b)$) iff $b \in L_*^{2_1}(a)$ (resp $L'_*(a)$).* □

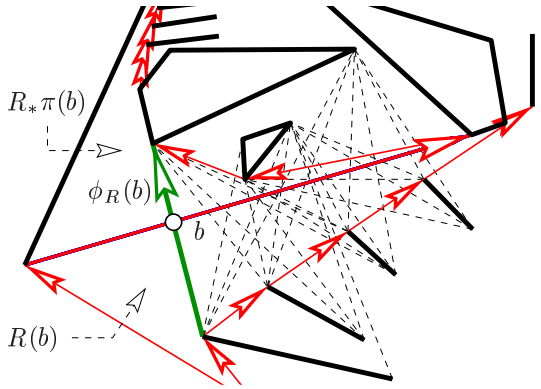

Figure 5: $\phi_R(b)$ leaves $R(b)$ and enters $R_*\pi(b)$.

We now turn to the definition of the operator ϕ_R, under three equivalent forms (see Figures 6 and 5).

Theo. & Def. 5. *We define $\phi_R(b)$ for $b \in X_0$ not on the convex hull to be the bitangent t in the range $[b, \pi(b)]$ that satisfies the three equivalent conditions:*
(1-) *t leaves $R(b)$ and enters $R_*\pi(b)$,*
(2-) *t is the bitangent crossing b whose crossing point with b is the closest to the tail of b,*
(3-) *t is the largest bitangent that leaves $R^{lea}(b)$.*
The left version ϕ_L of ϕ_R is defined by the relation $\pi\phi_R = \phi_L\pi$. □

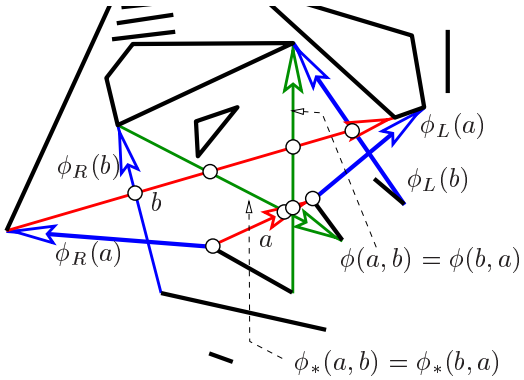

Figure 6: A pair $(a, b) \in \mathcal{G}$.

We are now ready to state the 'sum of squares' theorem.

Theorem 6 ("Sum of Squares" Theorem). *Let $\mathcal{G}$ be the set of pairs $(a, b) \in X_0 \times X_0$ such that $\phi(a, b) = \phi(b, a)$ and $\phi_*(a, b) = \phi_*(b, a)$ where $\phi(u, v) := \phi(u; H(v))$ and $H(v) = \{\phi_R(v), \phi_L(v)\}$. We claim that $|\mathcal{G}/\pi^2| = O(|B|)$. [A sketch of the proof is given in the appendix]* □

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. Before, some notations : $r(b)$ is the set of bitangents a such that $\phi_{\text{left}}(a) = \phi_{\text{righ}}(b)$, — this set is a subset of the set of bitangents that appear on the right boundary of the face $\sigma_{\text{righ}}(b)$ —, and b_R is the successor of b in the greedy pseudotriangle $R(b)$. A right-xx (xx-right) bitangent is a bitangent that leaves (enters) a disk on its right.

Theorem 7. *Let $Y_R(b)$ be the set of a such that one of three following conditions is satisfied*
(1-) $a \in R'(b)$ *and* $\phi_R(b_R)$ *is right-xx*
(2-) $a \in r(b)$, $a_L \in L^1(a)$, $b_R \in R^1(b)$ *and* $\phi_R(b_R)$ *is right-xx,* $\phi_L(a_L)$ *is xx-left and* $\phi_{righ}(b)$ *is right-left*
(3-) $a \in R'_*(b)$ *and* $\phi_L(a_L)$ *is xx-left.*
Let a and b be such that $a \in Y_R(b)$ (iff. $b \in Y_L(a)$). Assume that $a \neq \min Y_R(b)$ and $b \neq \min Y_L(a)$. Then $(a,b) \in \mathcal{G}$. □

Now we relate for a positive 0-cusp edge with source b and sink b_R the chain $rc(b)$ to the bitangent $\phi_R(b_R)$ (Bridge Theorem) and show that $\phi_R(b_R)$ can be computed efficiently with the χ_1 predicate in the dynamic setting of the GFA.

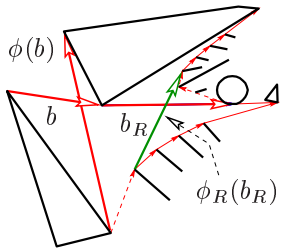

Theorem 8 (Bridge Theorem). *Let e be a positive 0-cusp edge with source b and sink b_R, let $b_R S$ be the prefix subchain of $R_*\pi(b_R)$ (for $\prec$) with sink $\phi_R(b_R)$, and let T be the suffix subchain of $R^{lea}(b_R)$ with source $\phi_R(b_R)$ (its sink is $\phi(b)$). We claim that $R^{lea}(b_R)$ is a prefix subchain of $R^{lea}(b)$ and that $rc(b) = beS\phi_R(b_R)T\phi(b)$.* □

We use the following procedure.

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'))

Function 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;
5 **return** 0;

Clearly the complexity is the sum of the sizes of the chains $R'(b_R)$ and S where S is the chain introduced in Theorem 8. This shows that the complexity of χ_1-WALK is at most $|rc(b)| + |R'(b)|$. Using a tandem walk (as explained in the previous section) we can reduce this complexity to $|rc(b)|$ in case where $\phi_R(b_R)$ is left-xx. 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 $\chi_1(\phi, r_g^+)$
that we deduce from properties of the χ_1 predicate on a collection of 2 disks and a point (= circle with an infinitesimal radius) sum up in the following theorem which deals with "sequences" 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 $\chi_1(t_1, t_p)$).

Theorem 9. *Consider three disjoint disks O_0, O_1, O_2 where O_0 is a point; we denote by b_{ij} the bitangent that leaves the signed disk O_i and enters the signed disk O_j, $i, j \in \{\pm 0, \pm 1, \pm 2\}$. Assume that b_{0i}, b_{0j} are free and that $\chi_1(b_{0i}, b_{0j})$. Then we have the following 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* $i = \pm 1, j = +2$. □

Let T_R and T_L denote respectively the pseudotriangles locally to the right and to the left of ϕ in the pseudo-quadrangle defined by $R(b)$ and $R_*\pi(b_R)$. Let O_i (resp. O_j) be the signed disk containing r_g (resp. l). It should be clear that $t = b_{ij}$. We will use the fact that the supporting lines of two bitangents a, b belonging to the boundary of a pseudotriangle T intersect inside T. We denote $p(a,b)$ this point in the sequel.

We now prove claim (1-). We first prove $\chi_1(l^+, t)$. Since l^+ and ϕ belong to the same side of T_L, $O_0 := p(l^+, \phi)$ is well defined and $l^+ = \lambda^2 b_{0j}, \phi = \mu^2 b_{i0}$ with $\lambda, \mu \in \mathbb{R}$. Moreover, since $l^+ \prec \phi$ we have $\chi_1(l^+, \phi)$ which can be rewritten $\chi_1(b_{0j}, b_{i0})$, i.e. $\chi_1(b_{0-i}, b_{0j})$.

Claims (1-) and (3-) of Theorem 9 give one (or both) of the following sequences: $b_{0-i}b_{0j}b_{ij}$ and $b_{0j}b_{ij}b_{i-j}$. Both sequences yield $\chi_1(b_{0j}, b_{ij})$ which is what we wanted, namely $\chi_1(l^+, t)$.
To prove $\chi_1(t, r_g^+)$, we redefine O_0 as $p(l^+, r_g^+)$. The point O_0 is well defined because r_g^+ and l^+ belong to T_L, and we can still write $l^+ = \lambda^2 b_{0j}$. If O_0 lies before r_g^+ on the directed supporting line of r_g^+ – case (i) – then $r_g^+ = \mu^2 b_{0i}$; otherwise – case (ii) – $r_g^+ = -\mu^2 b_{0,-i}$. Since $l^+ \prec r_g^+$ we have $\chi_2(l^+, r_g^+)$ and thus $\chi_1(b_{0j}, \epsilon b_{0,\epsilon i})$ where ϵ is respectively equal to + and − in cases (i) and (ii). In case (i), r_g is necessarily positive, i.e. $i > 0$. Claim (3-) of Theorem 9 (where i and j are exchanged) gives the sequence $b_{0i}b_{-ji}b_{-j-i}$. Consequently $\chi_1(b_{0i}, b_{-j-i})$ or equivalently $\chi_1(b_{ij}, b_{0i})$. This last conclusion proves $\chi_1(t, r_g^+)$ in case (i). In case (ii), $\chi_1(b_{0j}, \epsilon b_{0,\epsilon i})$ becomes $\chi_1(b_{0,-i}, b_{0j})$. Depending on the signs of i and j, we use claims (1-) or (2-) of Theorem 9 to obtain one of the two sequences $b_{0,-i}b_{0j}b_{ij}$ or $b_{0,-i}b_{-ij}b_{ij}$. In both cases we obtain $\chi_1(b_{0,-i}, b_{ij})$ which can be rewritten as $\chi_1(b_{ij}, -b_{0,-i})$. This last conclusion is exactly $\chi_1(t, r_g^+)$.

We now prove claim (2-). Since $r \neq r_g$, r^+ and ϕ belong to T_R and we can define $O_0 := p(\phi, r^+)$. We can write $r^+ = \lambda^2 b_{0i}$ and since $r^+ \prec \phi$ we have $\chi_1(b_{0i}, b_{0j})$. Since $r \neq r_g$ we have $i < 0$ because r_g is the only eventually positive arc of $R^{\text{lea}}(a)$ (in that case r_g belongs to $R^3(a)$). The sequence of the second claim of Theorem 9 (recall that $i < 0$) yields $\chi_1(b_{0i}, b_{ij})$ which is exactly $\chi_1(r^+, t)$.

We now prove claim (3-). Since $r^+, \phi \in T_R$ and $l^+, \phi \in T_L$ we can define $x := p(r^+, \phi)$ and $y := p(l^+, \phi)$. We set $O_0 = y$; it should be clear that $l^+ = \lambda^2 b_{0j}$. The supporting line of ϕ partitions the convex hull of x and O_i in two (eventually empty) convex disks O_R and O_L lying respectively on the right and on the left. See Figure below:

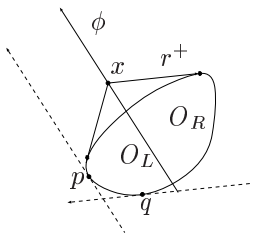

We have $\phi = \mu^2 b_{R0}$ and therefore $l^+ \prec \phi$, i.e. $\chi_2(l^+, \phi)$, can be rewritten as $\chi_1(b_{0j}, b_{R0})$ which is in turn equal to $\chi_1(b_{0,-R}, b_{0j})$. Observe that R is negative because $r \neq r_g$. Theorem 9 claim (1-) therefore gives us the sequence $b_{0-R}b_{0j}b_{Rj}$ and consequently $\chi_1(b_{0j}, b_{Rj})$. To prove $\chi_1(l^+, t)$ it therefore remains to prove that $t = b_{Rj}$ (recall that $l^+ = \lambda^2 b_{0j}$). We show successively that $t \in \{b_{Rj}, b_{Lj}\}$ and $t \neq b_{Lj}$. By hypothesis, we have $\chi_1(t, r^+)$, therefore t is tangent to the convex hull of O_i and x; this proves $t \in \{b_{Rj}, b_{Lj}\}$. If O_L is empty then the proof is finished, otherwise we have $\phi = -\mu'^2 b_{0L}$. We can rewrite $\chi_2(l^+, \phi)$ as $\chi_1(b_{0L}, b_{0j})$. Since L is negative (recall that r is negative because $r \neq r_g$), claim (2-) of Theorem 9 gives $\chi_1(b_{0L}, b_{Lj})$. In other words we have $\chi_1(b_{Lj}, \phi)$. Let p (resp. q) be the tangent point of the line tangent to O_L and parallel to ϕ (resp. $\pi(r^+)$). The conclusion $\chi_1(b_{Lj}, \phi)$ implies that b_{Lj} leaves the arc of O_L with source x and sink p. Since $\chi_1(r^+, \phi)$, the point q does not belong to this arc and therefore b_{Lj} does not leave the arc of O_i with source q and sink the target point of r^+. In other words we have $\chi_1(r^+, b_{Lj})$ and consequently $t \neq b_{Lj}$ because $\chi_1(t, r^+)$.

We now prove claim (4-). Since ϕ and r_g^+ belong to T_R with $\phi \prec r_g^+$ we have $\chi_1(\phi, r_g^+)$. Since ϕ and l^+ belong to T_L with $\phi \prec l^+$, we have $\chi_1(\phi, l_g^+)$.

If b is minimal in a filter I, we know that $R^{\text{lea}}(b) \subset G(I)$. To apply the procedure PHI-R we also need the chain $R_*^{\text{ent}}\pi(b_R)$. Unfortunately, this chain is not known completely in $G_*(I_*)$. Nevertheless it is not hard to show, using $\phi_*(b_R) \in G_*(I_*)$ and Theorem 3 claim (2-), that the subchain of $R_*^{\text{ent}}\pi(b_R)$ which is needed by the procedure PHI-R (i.e. the chain S of the Bridge Theorem) is a subchain of $G_*(I_*)$ and more precisely a subchain of $R_*(\phi_*(b_R), I_*)$. We sum up the discussion in the following theorem.

Theorem 10. *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 $|rc(b)|+|R'(b)|$ if $\phi_R(b_R)$ is right-xx, and time $|rc(b)|$ otherwise.* □

Finally we examine the case where the edge e is a positive 1- or 2-cusp edge, and study the equation $\phi_L(b) = \phi(b)$.

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 $\phi_R(b_R) = \pi(b)$.* □

Theorem 12 (Equation $\phi_L(b) = \phi(b)$). *The three following assertions are equivalent:*
(1-) $\phi_L(b) \in Cross(\phi_{righ}(b))$.
(2-) $\phi(b) \neq \phi_L(b)$.

(3-) $\phi_L(b)$ *does not leave the first atom of* $R^2(b)$ *nor 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* $\phi_R(b) = \pi\phi_*(b)$. □

4 Optimal flip using χ_1

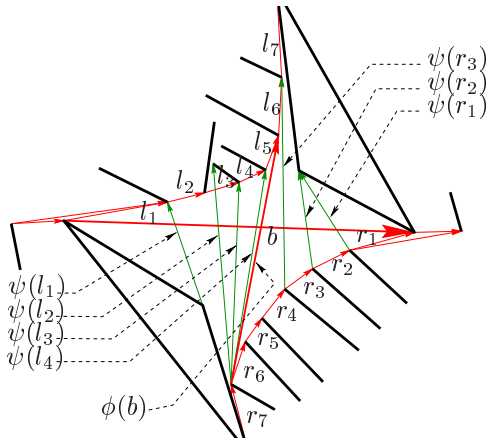

Figure 7: Flipping b using only χ_1. Here $r_u = r_3$, $l_5 = l_q$, $r_6 = r_p$ and $\psi(r_6) = \psi(l_5) = \phi(b)$.

We are now ready to describe the method based only on χ_1. 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.

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 12 we can detect in constant time if $\phi(b)$ coincides with $\phi_R(b)$ or $\phi_L(b)$; in that case we are done[3]. 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 10). A similar computation is done with b_L.

[3]The condition $\phi_L(b)$ leaves the first atom of $\pi R^2_*(b)$ is true (under the weaker assumption that $b, \phi_{\text{righ}}(b) \in G(I)$) iff. the arc of $G_*(I_*)$ that $\pi_*\phi_L(b)$ leaves coincides with the arc of $G_*(I_*)$ at the tail of $\phi_{\text{righ}_*}(b)$.

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 left-xx then $\phi(b)$ is also left-xx and therefore leaves the first atom (which is an arc) of $R^3(b)$. Similarly if l_1 is xx-right then $\phi(b)$ is xx-right 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, 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, omitted from this version, is similar to the proof of Theorem 10.

Theorem 13. *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-) *If* $1 \leq i < q'$ *then the bitangent* $\psi(l_i)$ *enters* $\pi R^{lea}_*(\phi_*(l_{i+1}), I_*)$.

(2-) *If* $q' \leq i < q$ *then the bitangent* $\pi\phi_*(l_{i+1})$ *leaves* $aL^3(b)$ *and* $\psi(l_i)$ *enters* $aL^3(b)$.

(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 intialisation 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** $(\text{Top}(C), \text{Top}(C'))$;

where c is a suffix subarc of $aL^3(b)$ 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 $aL^3(b)$ (added to treat

case (2-) of Theorem 13). 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 10, 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 left-xx the procedure PHI-R-NEW(b, l_{i+1}) can be too costly. Running tandem walks as in Theorem 10 allows us to compute $\psi(l_i)$ in time $|rc(l_i)|$ when $\psi(l_i)$ is left-xx. Assume now that $\psi(l_i)$ is right-xx. 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 $\phi_R(l_{i+1})$ is right-xx. 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 right-xx, 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 3 claim (2-) 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 3 claim (2-)). Using the fact that the bitangents of $R^2(l_i)$ are in the orbit of r_p under the operator ϕ_{back} (Theorem 3 claim (2-) again), we obtain $r_j \in R'(l_i)$ for $p+1 \leq j \leq m$. As $\phi_R(l_{i+1})$ is right-xx, 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 $\phi_R(b_R)$ and ideterminе whether it is right-xx; this is the cost of the procedure PHI-R(b, b_R).
(2) $|Y_L(b)| + |lc(b)|$ to compute $\phi_L(b_L)$ and obtain its type; 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 xx-right; 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 7 connects $Y_R(b)$ to $\mathcal{G}$ so with the help of Theorem 6 we can claim that $|Y_R(b)|$ is amortized constant. It follows that the cost of FLIP1 is amortized constant.

References

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

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

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

[4] F.S. 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.

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

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

[7] 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.

[8] 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.

[9] P. K. Agarwal, J. Basch, L. J. Guibas, J. Hershberger, and L. Zhang. Deformable free space tilings for kinetic collision detection, 2000.

[10] I. Streinu. A combinatorial approach to planar non-colliding robot arm motion planning. In *FOCS'00*, pages 443–453, 2000.

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

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

[13] S. Riviére. *Visibility computations in 2D polygonal scenes.* PhD thesis, Université Joseph Fourier, Grenoble, France, 1997.

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

A Appendix

Proof of the Duality Theorem. 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 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 Flip 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 $R^{\mathrm{lea}}(b)$. 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 $R^{\mathrm{lea}}(b)$ and that the edge e is a positive source cusp or a negative sink cusp in which case a belongs to or leaves $R^{\mathrm{lea}}(b)$, $a' \notin R(b)$ (and thus $b \preceq \phi_*(a')$ according to Claim 3); the induction hypothesis gives b belongs to or leaves $L^{\mathrm{lea}}_*(a)$; the flip property gives that $L^2_*(a) = R^{3_1}_*(a')\phi_*(a')L^{2_2}_*(a')$, $L^3_*(a) = L^{3_2}_*(a') \oplus e$ (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 $L^{\mathrm{lea}}_*(a')$. The other cases are similar. □

We give now a sketch of the proof of the 'sum of squares' Theorem. We use the notation $\Phi(a, b)$ instead of $\Phi(a; H(b))$ for $\Phi = \phi, \phi_{\mathrm{back}}, \phi_{\mathrm{forw}}, \ldots, R, L, \ldots$, etc.

For $b \in X_0$ we denote by x_1, x_2 the intersection points of b with $\phi_R(b)$ and $\phi_L(b)$ and we introduce $\sigma(b)$ the 2-face of the visibility complex $\mathbb{V}(H(b))$ defined by the line segment (x_1, x_2), provided that $x_1 \neq x_2$. (If $x_1 = x_2$ then $\phi_R(b) = \phi_L(b)$ and we set $\sigma(b) = \emptyset$.)

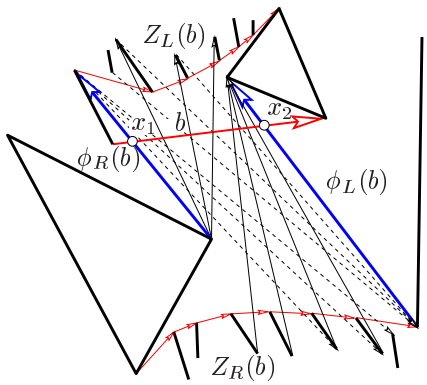

Figure 8: The face $\sigma(b)$. In plain lines are shown the $\phi(a, b)$ such that $\phi(a, b) = \phi(b, a)$ and in dotted lines the $\phi_*(a, b)$ such that $\phi_*(a, b) = \phi_*(b, a)$. We have $|Z_R(b)| = 7, |Z_L(b)| = 6$ and $(a, b) \in \mathcal{G}$ only for the first three bitangents a of $Z_R(b)$ and $Z_L(b)$.

The first step is to relate $\mathcal{G}$ to the set of pairs (a, b) such that $a \in \sigma(b)$. Let $Z_R(b)$ be the set of bitangents not in $\sigma_{\mathrm{sour}}(b)$ nor in $\sigma_{\mathrm{sink}}(b)$ such that $\phi_{\mathrm{left}}(a, b) = \mathrm{sink}(\sigma(b))$. The bitangents of $Z_R(b)$ are consecutive bitangents of the right boundary of $\sigma(b)$. As usual we set $Z_L(b) := \pi^{-1} Z_R \pi(b)$.

Theorem 14. *(1-) $a \in Z_R(b)$ iff. $b \in Z_L(a)$ iff. $a \in R'(b), r(b)$ or $R'_*(b)$.*
(2-) Let $(a, b) \in \mathcal{G}$. Then $a \in Z_R(b)$ or $a \in Z_L(b)$.

Proof. Omitted. □

It is convenient to 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. It should be clear that $(a, b) \in \mathcal{G}_R$ iff. $(\pi(a), \pi(b)) \in \mathcal{G}_L$. Let now I be a filter of $(X_0, \prec)$ chosen once for all, and let $\mathcal{G}(I), \mathcal{G}_R(I), \mathcal{G}_L(I)$ to 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 $|\mathcal{G}_R(I)| = O(|B|)$.

Theorem 15. *Let $a \in Z_R(b)$. Then:*
(1-) $R(a, b) = R(a; \{\phi_L(b)\}) = R(a)$.
(2-) $\phi(a, b) = \phi(a; \{\phi_L(b)\})$.
(3-) Assume that $\phi(a, b) = \phi(b, a)$ and that a enters a disk on its right. Then $\phi_{forw}(a, b)$ enters $L(b)$ or $\phi_L(b) \prec \phi_{forw}(a, b)$.

Proof. Claim (1-). Omitted.

Claim (2-). Since $\phi_L(b) \in H(b)$ one has $\phi(a;\{\phi_L(b)\}) \preceq \phi(a,b)$ (Theorem 1). Let t be a bitangent that leaves $R(a)$ and pierces a. If $t \prec \phi(a,b)$ than t pierces the forward view of the rays in the face $\sigma(b)$, i.e., pierces $\phi_L(b)$; therefore $\phi(a,b) \preceq \phi(a;\{\phi_L(b)\})$.

Claim (3-). Since $\phi(b,a)$ enters $L(b,a)$, the left version of claim (1-) and the hypothesis $\phi(a,b) = \phi(b,a)$ imply that $\phi(a,b)$ enters $L(b)$. We show that a certain subchain of $L^{\rm ent}(a,b)$ is a subchain of $L^{\rm ent}(b)$. Let $C(a,b)$ be the subchain of $L^{\rm ent}(a,b)$ whose source is $\phi(a,b)$ and whose sink is the minimal bitangent of $L^3(a,b)$, and let $C(b)$ be the subchain of $L^{\rm ent}(b)$ whose source is $\phi(a,b)$ and whose sink is whose sink is $\phi_L(b)$. According to the Orbit Theorem we can write that

$$C(a,b) = \epsilon'_0 b'_1 \epsilon'_1 \dots b'_p \epsilon'_p \text{ and } C(b) = \epsilon_0 b_1 \epsilon_1 \dots b_q \epsilon_q$$

where $b'_1 = \phi_{\rm forw}(\phi(a,b),b)$, $b_1 = \phi_{\rm forw}(\phi(a,b))$, and where the sequences $(b_i), (b'_j)$ are defined inductively by $b_{i+1} = \phi_{\rm forw}(b_i)$ and $b'_{j+1} = \phi_{\rm forw}(b'_j, b)$. If $\phi_{\rm forw}(a,b) \prec \phi_L(b)$ then the sink of the arc ϵ'_p – i.e. the minimal bitangent of $L^3(a,b)$ – is $\phi_{\rm forw}(a,b)$ otherwise the sink of ϵ'_p is $\phi_L(b)$. In both cases the elements of $C(a,b)$ are smaller than $\phi_L(b)$. Moreover in the second case we have $\phi_L(b) \prec \phi_{\rm forw}(a,b)$ because $\phi_L(b) \in L^{\rm ent}(a,b)$ and $\phi_{\rm forw}(a,b)$ is the sink of $L^{\rm ent}(a,b)$ (see claim (3-) of the Orbit Theorem).

We claim that $b_i \in X_0(H(b))$ for all $i \leq q$. Indeed, $\phi_L(b)$ enters $L^{\rm ent}(b)$ and therefore intersects no bitangent in $L(b)$. To prove that b_i does not intersect $\phi_R(b)$ we use Theorem 12 which states that $\phi_R(b)$ intersects no bitangent in $L^2(b)$ (case where $\phi_R(b) = \phi(b)$) or one bitangent in $L^2(b)$ namely $\min_\prec L^2(b)$ (case where $\phi_R(b) \neq \phi(b)$). In the former case $\phi(b)$ does not enter the first atom of $L^2(b)$ therefore $b_1 \neq \min_\prec L^2(b)$ and this proves $b_i \in X_0(H(b))$. Consequently $\phi_{\rm forw}(\phi(b)) = \phi_{\rm forw}(\phi(b),b)$ and $\phi_{\rm forw}(b_i) = \phi_{\rm forw}(b_i,b)$. This proves $p \leq q$ and $\epsilon_i = \epsilon'_i$ and $b_i = b'_i$ for $i \leq p$. Thus $C(a,b)$ is a subchain of $C(b)$. Claim (3-) follows immediately. □

Theorem 16. *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. The proof is decomposed into four parts. In part 1, we prove that $\phi_L(b) \neq \phi_L(b')$ and in part 2, we use the Orbit Theorem and the common bitangent $\phi(a,b)$ $(= \phi(a',b'))$ to rebuild the chains $L^{\rm ent}(b), L^{\rm ent}(b')$ and prove that they coincide; in particular we will obtain $b \prec b'$ and the fact that a bitangent t entering $L^{\rm ent}(b')$ enters $L^{\rm ent}(b)$ or $\pi(b) \prec t$. Our candidate for t is $\phi_{\rm forw}(a',b')$; we use Theorem 15 claim (3-) to relate this last bitangent to $L^{\rm ent}(b')$. To prove that $\phi_{\rm forw}(a',b')$ does not enter $L^{\rm ent}(b)$ we show that $\phi_L(b) \prec \phi_{\rm forw}(a',b')$ and use the fact that $\phi_L(b)$ is the maximal bitangent entering $L^{\rm ent}(b)$ (this is part 4).

Part 1. We show that the assumption $b \neq b'$ implies $\phi_L(b) \neq \phi_L(b')$.
Assume that $\phi_L(b) = \phi_L(b')$. We have $\phi(a,b) = \phi(a',b')$, therefore claim (2-) of Theorem 15 implies $\phi(a;\{\phi_L(b)\}) = \phi(a';\{\phi_L(b)\})$ and consequently $a = a'$ by the injectivity of the operator $\phi(\cdot;\{\phi_L(b)\})$ in $X_0(\phi_L(b))$. This in turn implies $\phi_R(a) = \phi_R(a')$ and using the equalities $\phi(a,b) = \phi(b,a)$ and $\phi(a',b') = \phi(b',a')$ we obtain $b = b'$.

Part 2. Let $t_0 = \phi(a,b)$. As $L(b,a) = L(b)$ and t_0 enters $L(b,a)$ we can introduce the the suffix subchain $C(b)$ of $L(b)$ with source t_0. Similarly we define the suffix subchain $C(b')$ of $L(b')$ with source $t'_0 = \phi(a',b')$. We prove successively that:
(1-) b and b' are comparable and enter the same disk.
(2-) $C(b') = C(b) \oplus \pi(e)$. Where e is the arc with source b and sink b'.
(3-) $\pi(b) \prec \phi_L(b')$.
Let $t_1, \dots$ be the sequence of bitangents defined by $t_i = \phi_{\rm forw}(t_{i-1})$ for $i \geq 1$. According to Theorem 3 we can write that $C(b) = t_0 \epsilon_0 t_1 \epsilon_1 \dots t_p \epsilon$, 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; t_p and b enter the same object).
Similarly we can write $C(b') = t'_0 \epsilon'_0 t'_1 \epsilon'_1 \dots t'_{p'} \epsilon'$ where $t'_i = \phi_{\rm forw}(t'_{i-1})$. Since by hypothesis $t_0 = t'_0$ it follows that $t_i = t'_i$ for all i and consequently that b and b' enter (on their left) the same object. Therefore b and b' are comparable. As $b' \not\preceq b$ we have $b \prec b'$ and therefore $C(b') = C(b) \oplus \pi(e)$.
As $\phi_L(b)$ enters $C(b)$ $(\subset L^{\rm ent}(b))$, $\phi_L(b)$ enters $C(b')$; therefore $\phi_L(b) \preceq \phi_L(b')$ because $\phi_L(b')$ is the maximal bitangent entering $C(b')$. The bitangent $\phi_L(b')$ enters $C(b)$ or enters $\pi(e)$; in the first case we obtain that $\phi_L(b') = \phi_L(b)$ which contradicts the assumption $b \neq b'$ (see Part 1). So we conclude that $\phi_L(b')$ enters $\pi(e)$ and therefore $\pi(b) \prec \phi_L(b')$.

Part 3. We claim that $\phi_*(a) \prec b$ and $\phi_*(b) \prec a$.
Theorem 14 tells us successively that $b \in Z_L(a)$ and that b belongs to one of the three chains $L'(a), l(a)$ or $L'_*(a)$. We know (claim (2-) of Theorem 14), that $\phi(a,b)$ leaves $R(a)$. Moreover $\phi(a,b) = \phi(b,a)$ implies that $\phi(a,b)$ intersects b and therefore $\phi(a,b) \neq \phi(a)$. This in turn implies $\phi(a) \neq \phi_R(a)$ otherwise $\phi(a)$ would be the only bitangent leaving $R(a)$ and in-

tersecting a. Theorem 12 tells us that $\phi_*(a)$ does not leave the first atom of $L^2_*(a)$ and consequently $\min_{\prec_*} L^2_*(a) \prec_* \phi_*(a)$. But Theorem 3 claim (2) gives $\min_{\prec_*} L^2_*(a) = \phi_{\text{left}_*}(a)$ thus $\phi_{\text{left}_*}(a) \prec \phi_*(a)$. This proves that if b belongs to $l(a)$ (whose source is $\phi_{\text{left}_*}(a)$) or $L'_*(a)$ we have $b \prec_* \phi_*(a)$. When $b \in L'(a)$ we simply write the two inequalities $\phi_*(a) \prec a$ and $a \prec b$. The proof of $\phi_*(b) \prec a$ is similar.

Part 4. We prove that $\pi(b) \prec \phi_{\text{forw}}(a', b')$.
According to the claim (3-) of Theorem 15, $\phi_{\text{forw}}(a', b')$ enters $L(b')$ or $\phi_L(b') \prec \phi_{\text{forw}}(a', b')$. In the latter case we are done since claim (3-) of Part 2 gives $\pi(b) \prec \phi_L(b')$. Assume now that $\phi_{\text{forw}}(a', b')$ enters $L(b')$. It follows that $\phi_{\text{forw}}(a', b')$ enters $C(b')$. To prove that $\pi(b) \prec \phi_{\text{forw}}(a', b')$ we argue that $\phi_{\text{forw}}(a', b')$ enters $\pi(e)$, i.e. $\phi_{\text{forw}}(a', b')$ does not enter $C(b)$. This is a consequence of the following claims:
(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 we deduce that $\phi_L(b) \prec \phi_{\text{forw}}(a', b')$ and consequently that $\phi_{\text{forw}}(a', b')$ does not enter $C(b)$ because $\phi_L(b)$ is the maximal bitangent entering $C(b)$.

We prove claim (1-) with the help of the minimum element property in $X_0(\phi_L(b))$ by showing that:
$\phi_*(a) \prec \pi_*\phi_L(b') \prec \phi(a; \phi_L(b))$
from which we deduce that a and $\phi_L(b')$ are non-crossing because $\phi_L(b') \in X_0(\phi_L(b))$.

The first inequality follows from $b \prec \pi_*\phi_L(b')$ (claim (3-) of Part 2) and $\phi_*(a) \prec b$ (Part 3). 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 $\phi(a, b) \prec \phi_L(b')$.

To prove claim (2-) assume that $\phi_L(b)$ and a' are not crossing. The bitangents $\phi_L(b)$ and $\phi_L(b')$ both enter $L^{\text{ent}}(b')$ therefore they do not cross. Claim (1-) tells us that $\phi_L(b')$ and a are not crossing. Therefore we can introduce $G = \{\phi_L(b), \phi_L(b')\}$ and state that the bitangents a, a' belong to $X_0(G)$. Moreover $\phi_L(b')$ and $\phi(a', b')$ enter $L^{\text{ent}}(b')$ so they do not intersect. With the help of claim (2-) of Theorem 15 we can therefore write $\phi(a, b) = \phi(a; \{\phi_L(b)\}) = \phi(a; G)$. Similarly $\phi(a', b') = \phi(a'; G)$. The injectivity of the operator $\phi(\cdot; G)$ in $X_0(G)$ gives us $a = a'$ which contradicts the assumption $b \neq b'$ (see Part 1).
The inequality of claim (3-) is a consequence of $\phi_L(b) \preceq \phi_L(b')$, $\phi_L(b') \preceq \pi\phi_*(b')$ and $\pi\phi_*(b') \prec \pi(a')$ (see Part 3 for this last inequality).

To prove claim (4-) observe that $\phi_L(a', b')$ enters $L^{\text{ent}}(a', b')$ and use claim (3) of Theorem 3 which states that $\phi_{\text{forw}}(a', b')$ is the sink of $L^{\text{ent}}(a', b')$ (recall that a' is xx-right). Claim (5-) is trivial. Claim (6-) is a consequence of claims (2-) and (3-), $\phi_L(b) \in X_0(T)$ and the minimum element property in $X_0(T)$. □

We are now ready to prove the "sum of squares" Theorem.

Let $\mathcal{G}^+_R(I)$ $(\mathcal{G}^-_R(I))$ be the set of $(a, b) \in \mathcal{G}_R(I)$ such that $\phi(a, b) \notin \pi(I)$ $(\phi_*(a, b) \in I)$. 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.

(1-) follows $\phi(a, b) \preceq \pi\phi_*(a, b)$ (Minimum Property).
(2-) $(\mathcal{G}^+_R)_*(J) = \{(a, b) \in (\mathcal{G}_R)_*(J) \mid \phi_*(a, b) \in I\}$. The result follows from $(\mathcal{G}_R)_*(J) = \mathcal{G}_R(I)$.
(3-) The number of pairs (a, b) in $\mathcal{G}_R(I)$ such that a is xx-left is a $O(|B|)$ because $Z_R(b)$ contains at most one xx-left bitangent. Similar arguments allow us to restrict our attention to the case where b,b' are xx-left and a,a' are xx-right. We prove that $(a, b) \mapsto \phi(a, b)$ is an injective application from $\bigcup_{b \in I \cap J} \mathcal{G}^+_R(b, I) \setminus \min_{\prec} \mathcal{G}^+_R(b, I) \times \{b\}$ into $I \cap J$. As a' and a'' belong to $Z_R(b)$ with $a'' \prec a'$, claim (1-) of Theorem 3 gives $a' \in L^1(a'', b')$. It is not hard to see, using the flip property, that $\min_{\prec} L^3(a', b') \preceq \phi(a'', b')$ (with equality iff. a' is the successor of a'' in $\sigma(b)$). With the help of Theorem 3 claim (3-) we get $\min_{\prec} L^3(a', b') = \phi_{\text{forw}}(a', b')$ (recall that a' is xx-right) from which we deduce $\phi_{\text{forw}}(a', b') \preceq \phi(a'', b')$ and therefore $\pi(b) \prec \phi(a'', b')$. Let $\mathcal{G}^+_R(b', I)$ denote the set of bitangents c such that $(c, b') \in \mathcal{G}^+_R(I)$. We have $a'' \notin \mathcal{G}^+_R(b', I)$ therefore $a' = \min_{\prec} \mathcal{G}^+_R(b', I)$.