Convex blocking and partial orders on the plane[1]

José Miguel Díaz-Báñez* Marco A. Heredia† Canek Peláez† J. Antoni Sellarès‡ Jorge Urrutia§
Inmaculada Ventura*

Abstract

Let $C = \{c_1, \ldots, c_n\}$ be a collection of disjoint closed convex sets in the plane. Suppose that one of them, say c_1, represents a valuable object we want to uncover, and we are allowed to pick a direction $\alpha \in [0, 2\pi)$ along which we can translate (remove) the elements of C one at a time while avoiding collisions. In this paper we find an $O(n^2 \log n)$ time algorithm that finds a direction α that minimizes the number of elements of C that have to be removed before we can reach c_1.

1 Introduction

Consider a set $C = \{c_1, \ldots, c_n\}$ of pairwise disjoint closed bounded convex sets, and a direction $\alpha \in [0, 2\pi)$; e.g., the vertical upwards direction. It is well known that the elements of C can be translated (removed) one at a time by moving them upwards while avoiding collisions with other elements of C [7, 10]. Suppose that c_1 is a special object that we want to uncover, and that we are allowed to choose a direction α along which we can remove the elements of C one at a time while avoiding collisions.

We want to find the direction α that minimizes the number of elements we need to remove before we reach c_1. In Figure 1, it is easy to see that if we remove the elements of C in the direction α_2, four elements of C have to be removed before we reach c_1, while for α_1 we only need to remove two.

This problem can be seen as a variant of the problem known in computational geometry as the *separability problem* [2, 5, 9]. It is also related to *spherical orders* determined by light obstructions [6].

In this paper we present an $O(n^2 \log n)$ time algorithm to solve this problem, assuming that for every pair of elements of C we can compute a tangent line to both of them in constant time.

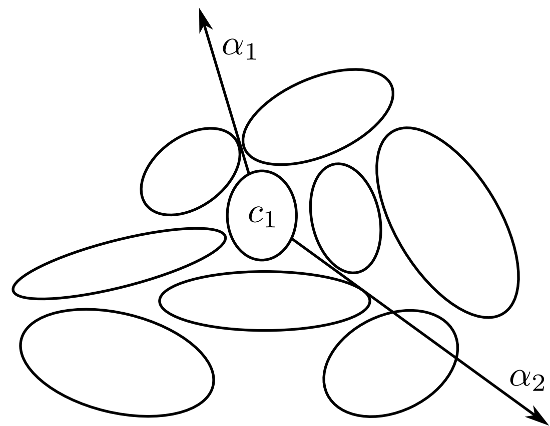

Figure 1: A set C of convex sets.

2 Preliminaries

Let X be a finite set, and $<$ a relation on the elements of X that satisfies the following conditions: (a) If $x < y$ and $y < z$ then $x < z$ (transitivity), and (b) $x \not< x$ (antireflexivity). The set X together with $<$ is called a partial order, and it is usually denoted as $P(X, <)$.

Given $x, y \in X$, we say that y *covers* x if $x < y$ and there is no element $w \in X$ such that $x < w < y$. The *diagram* of $P(X, <)$ is the directed graph whose vertices are the elements of X and there is an oriented edge from x to y if y covers x. We say that the diagram of $P(X, <)$ is planar if it can be drawn on the plane in such a way that the elements of X are represented by points on the plane, no edges of G intersect, except perhaps at a common endpoint, and if y is a cover of x, then they are joined by a monotonically increasing oriented edge from x to y (in the vertical direction).

Given two elements $x, y \in X$, a *supremum* of them is an element $w \in X$ such that $x < w$, $y < w$, and for any other element $z \in X$ such that $x < z$ and $y < z$ we have that $w < z$. An *infimum* is defined in a similar way, except that we require w to be $w < x$ and $w < y$. An ordered set is called a *lattice* if any two elements have a unique supremum and infimum. A lattice is called a *planar lattice* if its diagram is planar. Finally, a finite order $P(<, X)$ is called a *truncated planar lattice* if by

*Departamento Matemática Aplicada II, Universidad de Sevilla, Spain.
†Posgrado en Ciencia e Ingeniería de la Computación, Universidad Nacional Autónoma de México, Mexico.
‡Institut d'Informàtica i Matemàtica Aplicada, Universitat de Girona, Spain.
§Instituto de Matemáticas, Universidad Nacional Autónoma de México, Mexico.
[1]A shorter version of this paper appeared in the XIV Spanish Meeting on Computational Geometry held June 27-30, 2011.
Corresponding author: `canek@ciencias.unam.mx`

adding to $P(<, X)$ both a least and a greatest element the resulting order is a planar lattice.

Let $C = \{c_1, \ldots, c_n\}$ be a set of disjoint closed convex sets on the plane. Given two convex sets c_i and c_j in C, we say that c_j is an *upper cover* of c_i in the direction α (for short, an α-cover) if the following conditions are satisfied:

1. There is at least one directed line segment with direction α starting at a point in c_i and ending at a point in c_j.

2. Any directed line segment with direction α starting at a point in c_i and ending at a point in c_j does not intersect any other element of C.

Observe that if c_j is an α-cover of c_i, then c_i is an $(\alpha + \pi)$-cover of c_j. We say that c_j *blocks* c_i in the direction α, written as $c_i \prec_\alpha c_j$, if there is a sequence $c_i = c_{\sigma(1)}, c_{\sigma(2)}, \ldots, c_{\sigma(k)} = c_j$ of elements of C such that $c_{\sigma(r+1)}$ is an α-cover of $c_{\sigma(r)}$, $r = 1, \ldots, k-1$ (Figure 2).

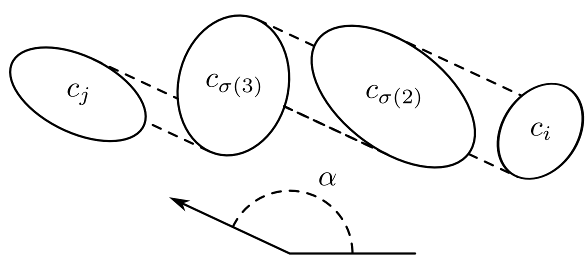

Figure 2: c_j is an α-cover of $c_{\sigma(3)}$ and $c_i \prec_\alpha c_j$.

Clearly if $c_i \prec_\alpha c_j$ and $c_j \prec_\alpha c_k$, then $c_i \prec_\alpha c_k$, and thus C together with the blocking relation $\prec_\alpha$ is a partial order on C, which we will denote as $P(\prec_\alpha, C)$. It is known that $P(\prec_\alpha, C)$ is a truncated planar lattice [10].

The diagram of such truncated lattice has the elements of C as vertices and there is an oriented edge from c_i to c_j if c_j is an α-cover of c_i (Figure 3). The elements of C that we need to remove in the α direction before an element c_i of C is reached are those convex sets c_j such that $c_i \prec_\alpha c_j$, the set containing these elements will be called the α-upper set of c_i, or for short, the *α-up-set* of c_i in α. Thus our problem reduces to that of finding the direction α such that the cardinality of the α-up-set of c_1 is minimized.

Observe that as α changes, so does $P(\prec_\alpha, C)$. In fact, it is easy to find families of convex sets for which $P(\prec_\alpha, C)$ changes a quadratic number of times. We proceed now to prove some properties of $P(\prec_\alpha, C)$, $0 \leq \alpha < 2\pi$ which will allow us to find an α such that the α-up-set of c_1 has minimum cardinality in $O(n^2 \log n)$ time.

Given a convex set c, a line ℓ is called a supporting line of c if it intersects c, and c is contained in one of the closed half planes determined by ℓ. Given two convex sets c_i and c_j, a line ℓ is called an internal tangent of them if ℓ supports them, and c_i is contained in one of the closed half planes determined by ℓ, and c_j in the other. A set of directions I is called an interval, if there are $\alpha, \beta \leq 2\pi$ such that the elements of I are angles of the form $\alpha + \delta$, $0 \leq \delta \leq \beta$, addition taken mod 2π.

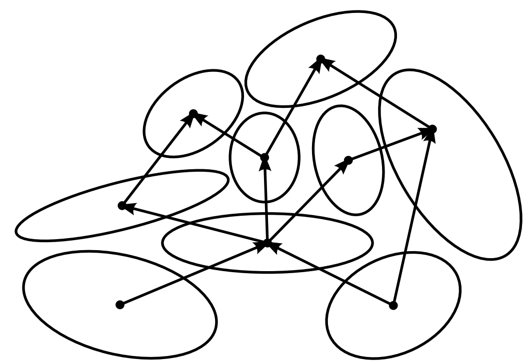

Figure 3: Diagram of $P(\prec_\alpha, C)$ for $\alpha = \pi/2$.

Lemma 1 *Let c_i and c_j be two convex sets in C. The set of directions in which c_j blocks c_i is a non-empty interval $\mathcal{I}_{i,j}$.*

Proof. Clearly a direction in which c_j does not block c_i always exists. Without loss of generality we will assume that such direction is 0.

Let θ_1 be the first direction greater than 0 where $c_i \prec_{\theta_1} c_j$: Such θ_1 exists because c_j always blocks c_i in a set of directions enclosed by the two internal tangents defined by c_i and c_j.

Let θ_2 be the *last* direction greater than θ_1 such that for any $\gamma \in [\theta_1, \theta_2]$ $c_i \prec_\gamma c_j$. If there is no other direction $\gamma \in [\theta_2, 2\pi]$ where $c_i \prec_\gamma c_j$ then our result holds. Suppose then that there are θ_3 and θ_4 such that i) $\theta_2 < \theta_3$, ii) $\theta_3 < \theta_4 < 2\pi$, and for $\gamma \in [\theta_3, \theta_4]$, $c_i \prec_\gamma c_j$, and iii) for any $\gamma \in [\theta_2, \theta_3]$, $c_i \not\prec_\gamma c_j$, (Figure 4).

Clearly $\theta_3 - \theta_2 < \pi$, or $\theta_1 - \theta_4 < \pi$, where the second angle is taken modulo 2π. Assume without loss of generality that $\theta_3 - \theta_2 < \pi$, and that $0 < \theta_2 < \frac{\pi}{2} < \theta_3$, for otherwise we can rotate C appropriately until this condition holds.

Let $\gamma = \frac{\pi}{2}$, then $c_i \not\prec_\gamma c_j$. Since $c_i \prec_{\theta_2} c_j$, we know that there is a sequence $c_i = c_{\sigma(1)}, c_{\sigma(2)}, \ldots, c_{\sigma(k)} = c_j$ such that $c_{\sigma(r+1)}$ is a θ_2-cover of $c_{\sigma(r)}$ for $r = 1, \ldots, k-1$. For the same reason, there is a sequence $c_i = c_{\omega(1)}, c_{\omega(2)}, \ldots, c_{\omega(m)} = c_j$ such that $c_{\omega(r+1)}$ is a θ_3-cover of $c_{\omega(r)}$ for $r = 1, \ldots, m-1$. The two sequences differ in at least one element, otherwise our gap would not exist.

Let ℓ_1 and ℓ_2 be the supporting lines of c_i in the γ direction: Since $c_i \not\prec_\gamma c_j$, c_j cannot intersect the interior

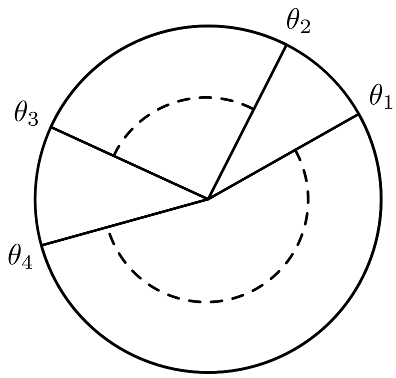

Figure 4: We can assume that $\theta_3 - \theta_2 < \pi$.

of the strip bounded by ℓ_1 and ℓ_2. Suppose first that c_j is to the left of this strip (Figure 5).

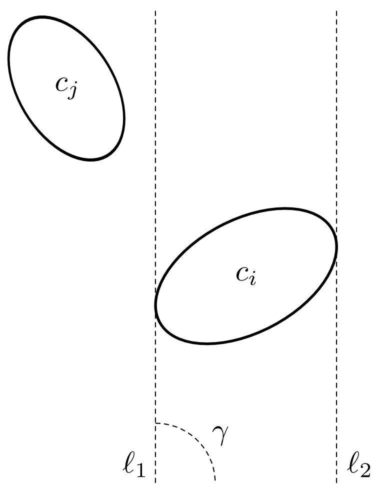

Figure 5: c_j to the left of ℓ_1 and ℓ_2.

Since $c_{\sigma(2)}$ is a θ_2-cover of $c_i = c_{\sigma(1)}$, there is a line segment parallel to the direction θ_2 with endpoints in $c_i = c_{\sigma(1)}$ and $c_{\sigma(2)}$. Similarly for $c_{\sigma(r)}$ and $c_{\sigma(r+1)}$, there is a line segment parallel to the direction θ_2 with endpoints in $c_{\sigma(r)}$ and $c_{\sigma(r+1)}$, $r \in \{2, \ldots, k-1\}$. Each $c_{\sigma(r)}$, $r \in \{2, \ldots, k-1\}$, contains two endpoints from two of this segments, and this endpoints can be joined with a line segment totally contained in $c_{\sigma(r)}$.

This forms a connected curve that starts in c_i and ends in c_j, passing through all the elements of the sequence. This curve consist of two types of line segments: Those parallel to the θ_2 direction, and those completely contained in the elements of the sequences. But $\theta_2 < \gamma$, so the first type always goes to the right. And the second type may go to the left, but contained in an element of the sequence (Figure 6).

The only way such a curve exists, is if at least one element in $\{c_{\sigma(1)}, c_{\sigma(2)}, \ldots,\ c_{\sigma(k)}\}$ intersects the strip

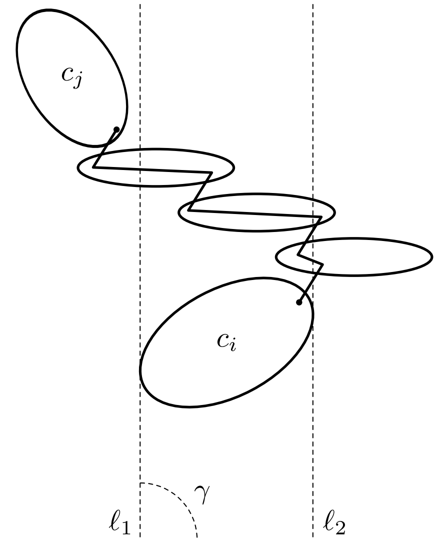

Figure 6: A sequence of θ_2-covers from c_j to c_i, and the curve that passes through the elements of it.

bounded by ℓ_1 and ℓ_2, which implies that $c_i \prec_\gamma c_j$, a contradiction!

If we suppose that c_j is to the right of ℓ_2, a contradiction arises, but using the sequence $c_i = c_{\omega(1)}, c_{\omega(2)}, \ldots, c_{\omega(m)} = c_j$ in the θ_3 direction. Therefore, the directions where c_j blocks c_i form a non-empty interval. □

It follows from the proof of Lemma 1 that the endpoints of the intervals $\mathcal{I}_{i,j}$ are defined by the internal tangents of pairs of elements in C. The next observation follows:

Observation 1 *There are at most $4\binom{n}{2}$ combinatorially distinct values of α where $P(\prec_\alpha, C)$ may change; these changes occur in slopes generated by internal tangents between pairs of elements of C.*

We can then reduce the search space for α_0 to the set $\mathcal{D} = \{\gamma_1, \ldots, \gamma_{4\binom{n}{2}}\}$ containing these directions. For the sake of clarity, we are supposing that no two internal tangents are parallel and that the elements of $\mathcal{D}$ are ordered as $\gamma_i < \gamma_j$ if $i < j$.

We observe that $\mathcal{D}$ can be calculated in $O(n^2 \log n)$ if we suppose that the internal tangents between any two convex sets in C can be determined in constant time. For each γ_k we can store the indexes i, j of the convex sets that define the internal tangent.

3 α-triangulations

Our problem can be solved by calculating the truncated lattices $P(\prec_{\gamma_i}, C)$ for every direction $\gamma_i \in \mathcal{D}$, and then obtaining the γ_i-up-set of c_1 in each one of them. Selecting a $\gamma_i \in \mathcal{D}$ which produces a smallest γ_i-up-set yields an optimal solution. Since calculating the truncated lattice has a cost of $O(n \log n)$ time for each of the $4\binom{n}{2}$ directions in $\mathcal{D}$ [10], this results in an $O(n^3 \log n)$-time algorithm to solve our problem.

To improve this complexity, we will show that we need to calculate from scratch only one truncated lattice. For the remaining directions of $\mathcal{D}$ the corresponding truncated lattice (more precisely, the α-triangulation, to be described shortly) can be updated in constant time.

For each direction $\alpha \in [0, 2\pi)$, we extend $P(\prec_\alpha, C)$ to a planar lattice $P'(\prec_\alpha, C)$ by adding two special vertices, a source s and a sink t, i.e. for each $c_i \in C$, $s \prec_\alpha c_i \prec_\alpha t$. For a fixed direction we can picture t as a very large convex set standing above all of C, and s as a very large convex set standing below all of C (Figure 7).

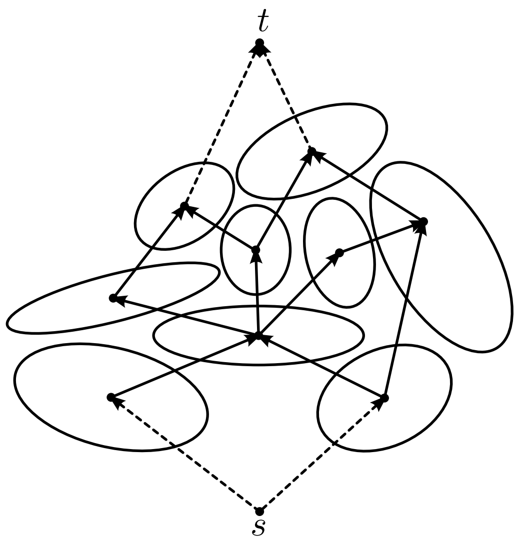

Figure 7: The lattice $P'(\prec_\alpha, C)$ for $\alpha = \pi/2$.

For each α, we now extend $P'(\prec_\alpha, C)$ to a triangulation $\mathcal{T}_\alpha$, that is, a planar graph where every internal face is a triangle, which we will call an *α-triangulation*, by adding oriented edges avoiding creating oriented cycles (Figure 8).

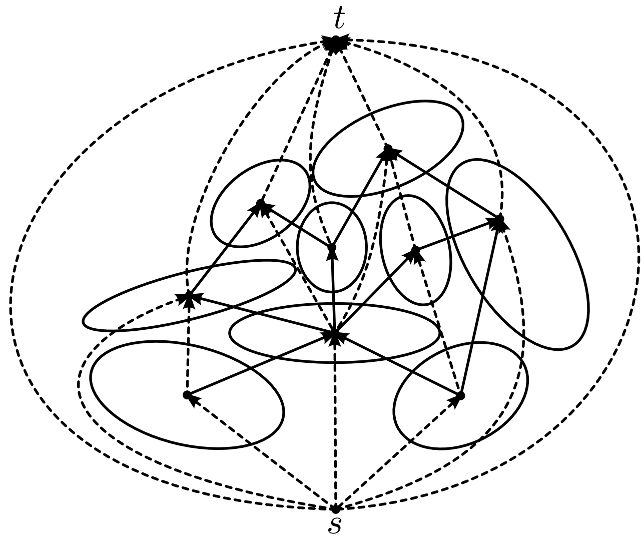

Figure 8: The triangulation $\mathcal{T}_\alpha$ for $\alpha = \frac{\pi}{2}$.

By Observation 1 there are at most $4\binom{n}{2}$ triangulations, and we want to know how $\mathcal{T}_\alpha$ changes as α goes from γ_k to γ_{k+1}. We remark that there are cases when the triangulations $\mathcal{T}_{\gamma_k}$ and $\mathcal{T}_{\gamma_{k+1}}$ are the same (Figure 9).

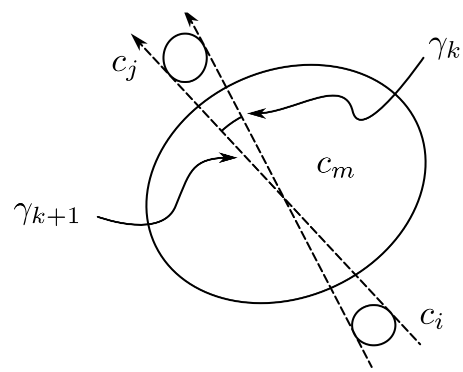

Figure 9: The triangulations $\mathcal{T}_{\gamma_k}$ and $\mathcal{T}_{\gamma_{k+1}}$ are the same, since the partial order does not change.

The next observation will be used:

Observation 2 *Let α, β be such that $P(\prec_\alpha, C) \neq P(\prec_\beta, C)$, then there is at least one pair of elements $c_i, c_j \in C$ such that c_j is an α-cover of c_i in $P(\prec_\alpha, C)$, and it is not a β-cover of c_i in $P(\prec_\beta, C)$, or vice versa; that is, the set of edges of the diagram of $P(\prec_\alpha, C)$ is different from the set of edges of the diagram of $P(\prec_\beta, C)$. Moreover, if $\alpha, \beta \in \mathcal{D}$, then c_i and c_j define α or β.*

It turns out that the difference between the $\mathcal{T}_{\gamma_k}$ and $\mathcal{T}_{\gamma_{k+1}}$ triangulations is an arc flip, as defined in [8]:

Lemma 2 *Given the triangulation $\mathcal{T}_{\gamma_k}$, the triangulation $\mathcal{T}_{\gamma_{k+1}}$ can be obtained from $\mathcal{T}_{\gamma_k}$ (if they are different) by flipping an arc in $\mathcal{T}_{\gamma_k}$. Such an arc flip either adds or removes an arc between the convex sets c_i and c_j that define γ_{k+1}.*

Proof. Suppose that $P(\prec_{\gamma_k}, C)$ and $P(\prec_{\gamma_{k+1}}, C)$ are different. By Observation 2 two cases arise:

1. There are two elements c_i and c_j of C that generate γ_k such that c_j is a γ_k-cover of c_i, but it is not a γ_{k+1}-cover of c_i.
2. The elements c_i and c_j that generate γ_{k+1} become comparable in $P(\prec_{\gamma_{k+1}}, C)$, and one of them, say c_j is a γ_{k+1}-cover of c_i.

In case 1 when we flip the edge connecting c_i to c_j in $\mathcal{T}_{\gamma_k}$ they become not comparable in $\mathcal{T}_{\gamma_{k+1}}$. Furthermore it is easy to see that there is a line parallel to the γ_{k+1} direction that separates c_i from c_j, and that this line intersects two elements of $C \cup \{s,t\} - \{c_i, c_j\}$. This are the two vertices that become adjacent as we flip the edge connecting c_i and c_j. Thus $\mathcal{T}_{\gamma_{k+1}}$ can be obtained from $\mathcal{T}_{\gamma_k}$ in constant time.

In the second case the inverse occurs. □

In Figure 10 and Figure 11 we can see an example of the arc flip performed in the proof of Lemma 2.

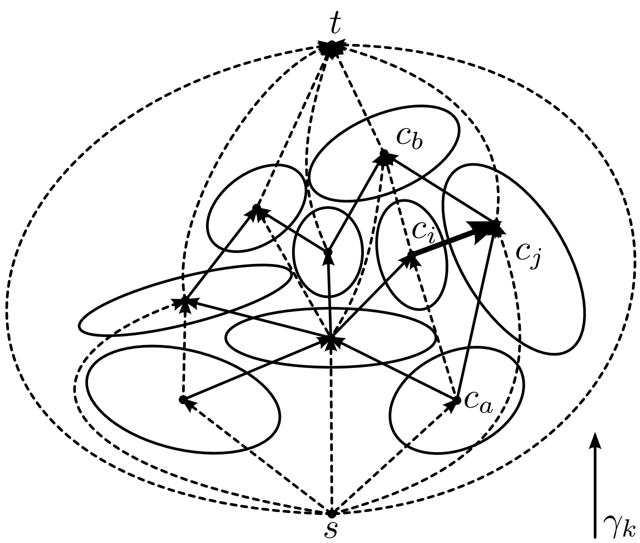

Figure 10: The arc $c_i \to c_j$ before flipping.

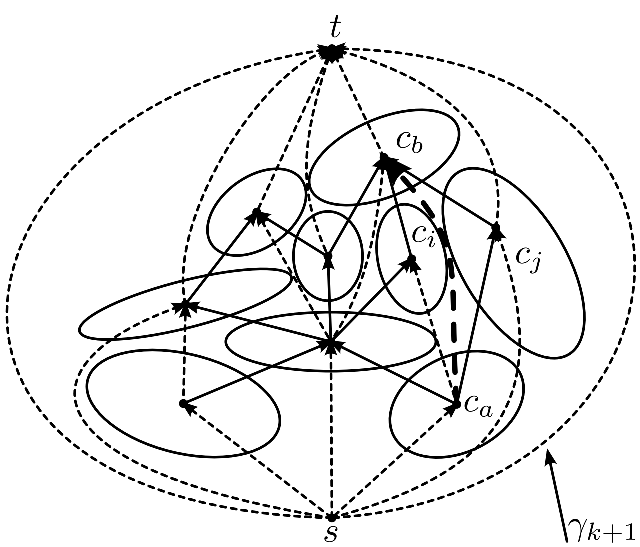

Figure 11: The arc $c_a \to c_b$ after flipping.

4 An $O(n^2 \log n)$ algorithm to find α_0

Theorem 3 *Finding α_0 can be done in $O(n^2 \log n)$.*

To prove Theorem 3 we need the following result:

Lemma 4 *For any element c_i, as we go from γ_1 to $\gamma_{4\binom{n}{2}}$, the up-set of c_i changes $O(n)$ times.*

Proof. By Lemma 1, the set of directions for which c_j blocks c_i is an interval $\mathcal{I}_{i,j}$. This means that for each $c_j \neq c_i$ in C, as we go from γ_1 to $\gamma_{4\binom{n}{2}}$, c_j starts to block c_i once and stops blocking it once. Therefore the up-set of c_i changes a linear number of times, that is any $c_j \in C$ enters and exits it once. □

We proceed now with the proof of Theorem 3.

Proof. We first generate the set $\mathcal{D}$ of critical directions in $O(n^2)$ time. Observe that this can be done in quadratic time since we are assuming that the tangents generated by two elements of C can be calculated in constant time. Next we sort the elements of $\mathcal{D}$ in $O(n^2 \log n)$ time. When we store each $\gamma_i \in \mathcal{D}$ we also store the elements of C that generate it. Next we construct $\mathcal{T}_{\gamma_1}$ in $O(n \log n)$ time, coloring the vertices of the triangulation as follows:

- We color red the elements of C in the γ_1-up-set of c_1, including c_1.
- We color blue the remaining elements of C.

At this stage, we also calculate the number of incoming arcs to each c_i whose initial vertex is blue, or red. Such a coloring can be done in $O(n)$ time. Let c_i and c_j be the elements that generated γ_{k+1}. It is easy to check that if c_j was not a γ_k-cover of c_i or vice versa, then $P(\prec_{\gamma_k}, C) = P(\prec_{\gamma_{k+1}}, C)$ and the up-set of c_1 does not change. Suppose then that c_j was a γ_k-cover of c_i. By Lemma 2, $P(\prec_{\gamma_k}, C) \neq P(\prec_{\gamma_{k+1}}, C)$ and we can obtain $\mathcal{T}_{\gamma_{k+1}}$ from $\mathcal{T}_{\gamma_k}$ in constant time. The crucial part of our procedure is how to update the up-set of c_1.

Suppose first that the elements c_i and c_j that determine γ_{k+1} are different from c_1.

Several cases arise.

1. $c_1 \prec_{\gamma_k} c_i$, $c_1 \prec_{\gamma_k} c_j$, and c_i is not comparable to c_j in $P(\prec_{\gamma_k}, C)$, but c_i is comparable to c_j in $P(\prec_{\gamma_{k+1}}, C)$. In this case, the up-set of c_1 remains unchanged.

2. $c_1 \prec_{\gamma_k} c_i$, $c_1 \prec_{\gamma_k} c_j$, and c_i is comparable to c_j in $P(\prec_{\gamma_k}, C)$, but c_i is not comparable to c_j in $P(\prec_{\gamma_{k+1}}, C)$. In this case the up-set of c_1 may change. Suppose that c_j is a γ_k-cover of c_i. Observe that c_i remains in the up-set of c_1, but c_j may not belong to it anymore. In this case the arc from c_i to c_j is flipped. If at least one arc from a red element c_r to c_j remains then c_j remains in the up-set of c_1, otherwise the up-set of c_1 changes, and is recalculated in linear time.

3. c_i and c_j do not belong to the up-set of c_1. In this case, the up-set of c_1 does not change.

4. $c_i \prec_{\gamma_k} c_j$ and c_i is not in the up-set of c_1 in $P(\prec_{\gamma_k}, C)$. The up-set of c_1 remains unchanged in $P(\prec_{\gamma_{k+1}}, C)$.

5. $c_i \not\prec_{\gamma_k} c_j$, c_j is not in the up-set of c_1, and c_i belongs to the up-set of c_1. In this case, c_j is an γ_{k+1}-cover of c_i and the up-set of c_1 changes. Therefore we must recalculate the up-set of c_1.

Each time we recalculate the up-set of c_1, we also recalculate for each c_i the number of incoming arcs starting at a blue or red point.

A similar case analysis is carried out when $c_i = c_1$, the details are left to the reader. By Lemma 4, we have to update the up-set of c_1 only a linear number of times, and thus the whole process takes $O(n^2 \log n)$ time. This proves Theorem 3. □

5 Conclusions

In this paper we studied a variant of the classic *separability problem.* Given a set $C = \{c_1, \ldots, c_n\}$ of pairwise disjoint closed convex sets, find a direction α minimizing the number of elements of C that have to be removed, in the direction α, before we reach a particular element $c_1 \in C$. We presented an $O(n^2 \log n)$-time algorithm to solve this problem, under the assumption that the internal tangents between any two sets of C can be calculated in constant time: For example, this is the case for circles and ellipses, convex polygons with a constant number of sides, and shapes defined by a constant number of Bézier curves.

We suspect that the complexity of our problem is $\Omega(n^2 \log n)$. In particular any approach that has to sort the elements of $\mathcal{D}$ may in general take $O(n^2 \log n)$ time unless some extra restrictions on the elements of C are imposed. For example for circles, we can sort the slopes generated by them in quadratic time (using the dual space), improving the complexity of our algorithm to $O(n^2)$. The details of this will be given in the full version of this paper.

It is easy to see that if we want to calculate for each $c_i \in C$ the number of elements that have to be removed before we can remove c_i, this can be done in $O(n^3)$.

6 Acknowledgments

Research of José Miguel Díaz-Báñez and Inmaculada Ventura partially supported by Spanish Government under Project MEC MTM2009-08652; research of Marco A. Heredia and Canek Peláez partially supported by CONACYT of Mexico; research of J. Antoni Sellarès partially supported by the Spanish MCI grant TIN2010-20590-C02-02; and research of Jorge Urrutia partially supported by SEP-CONACYT of Mexico, Proyecto 80268, and by Spanish Government under Project MEC MTM2009-08652.

References

[1] M. Abellanas, S. Bereg, F. Hurtado, A. García Olaverri, D. Rappaport, J. Tejel. Moving coins. *Computational Geometry*, 34(1):35-48, 2006.

[2] N. G. de Bruijn. Aufgaben 17 and 18 (in Dutch). *Nieuw Archief voor Wikskunde*, 2(954):67.

[3] B. Chazelle, T. Ottmann, E. Soisalon-Soininen, D. Wood. The Complexity and Decidability of Separation ICALP 119-127, 1984.

[4] A. Dumitrescu, M. Jiang. On Reconfiguration of Disks in the Plane and Related Problems WADS 254-265, 2009.

[5] L. Fejes-Tóth and A. Heppes. Über stabile Körpersysteme (in German). *Compositio Mathematica*, 15(2):119–126, 1963.

[6] S. Foldes, I. Rival and J. Urrutia. Light sources obstructions and spherical orders. *Discrete Mathematics*, 102(1):13–23, 1992.

[7] L. J. Guibas and F. F. Yao. On translating a set of rectangles. In *Proceedings of the Twelfth Annual ACM Symposium on Theory of Computing*, 154–160, 1980.

[8] F. Hurtado, M. Noy and Jorge Urrutia. Flipping Edges in Triangulations. *Discrete & Computational Geometry*, 22(3):333-346, 1999.

[9] G. Toussaint. Movable separability of sets. *Computational Geometry*, North-Holland, Amsterdam, 335–375, 1985.

[10] I. Rival and J. Urrutia. Representing orders on the plane by translating convex figures. *Order*, 4(4):319–339, 1988.