Convex blocking and partial orders on the plane

José Miguel Díaz-Báñez* Marco A. Heredia† Canek Peláez†

J. Antoni Sellarès‡ Jorge Urrutia§ Inmaculada Ventura*

April 30, 2013

Abstract

Let $C = \{c_1, \ldots, c_n\}$ be a collection of disjoint closed bounded 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. We study the problem of finding a direction α_0 such that the number of elements that have to be removed along α_0 before we can remove c_1 is minimized. We prove that if we have the sorted set $\mathcal{D}$ of directions defined by the tangents between pairs of elements of C, we can find α_0 in $O(n^2)$ time. We also discuss the problem of sorting $\mathcal{D}$ in $o(n^2 \log n)$ time.

1 Introduction

Consider a set $C = \{c_1, \ldots, c_n\}$ of pairwise disjoint closed bounded convex sets. It is well known that the elements of C can be removed one at a time by translating them upwards while avoiding collisions with other elements of C; see [11, 16]. For example, the elements of the set $C = \{c_1, \ldots, c_9\}$ shown in Figure 1(a) can be removed in the order $c_2, c_3, c_1, c_9, c_6, c_4, c_5, c_7, c_8$. Clearly this result is also valid if we remove the elements of C by translating them along any direction $\alpha \in [0, 2\pi)$.

Suppose that $c_1 \in C$ is a special object that we want to uncover, and that we are allowed to choose a direction $\alpha \in [0, 2\pi)$ along which we can remove the elements of C one at a time while avoiding collisions. We want to find the

*Departamento Matemática Aplicada II, Universidad de Sevilla, Spain; partially supported by the Spanish Government under Project MEC MTM2009-08652.

†Posgrado en Ciencia e Ingeniería de la Computación, Universidad Nacional Autónoma de México, Mexico; partially supported by CONACYT of Mexico.

‡Institut d'Informàtica i Matemàtica Aplicada, Universitat de Girona, Spain; partially supported by the Spanish MCI grant TIN2010-20590-C02-02.

§Instituto de Matemáticas, Universidad Nacional Autónoma de México, Mexico; partially supported by SEP-CONACYT of Mexico, Proyecto 80268, and by the Spanish Government under Project MEC MTM2009-08652.

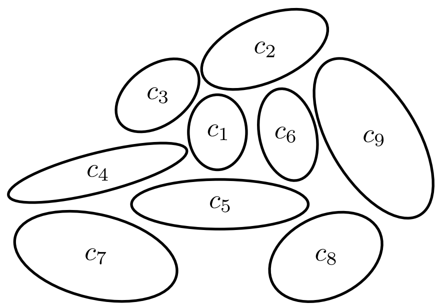

(a) The elements of C can be removed in the upwards direction in the order $c_2, c_3, c_1, c_9, c_6, c_4, c_5, c_7, c_8$.

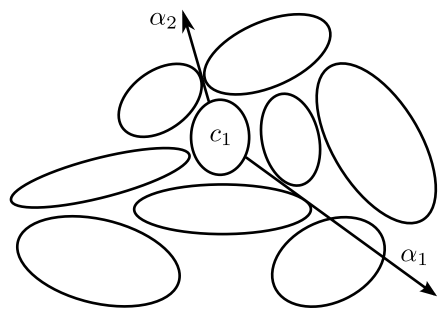

(b) Two different directions in which we can remove elements from C.

Figure 1: Disassembly in different directions.

direction α_0 that minimizes the number of elements we need to remove before we reach c_1. For example, in Figure 1(b), it is easy to see that if we remove the elements of C in the direction α_1, four elements of C have to be removed before c_1 is uncovered, while for α_2 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* [5, 4, 9, 17]. Similar problems are studied in [1, 7], and it is also related to *spherical orders* determined by light obstructions [10].

In this paper, we present an $O(n^2)$ time algorithm to solve this problem, assuming that we have the sorted set $\mathcal{D}$ of directions defined by the tangents between pairs of elements of C. To ease our presentation, in the remainder of the paper we will assume that the interior of the convex sets is not empty. It is not hard to see that the result holds for families of closed sets.

In Section 2 we give basic definitions and state the problem in these terms. In Section 3 we explain how we can reduce the search space of our problem to the set $\mathcal{D}$ of critical directions. In Section 4 we present the data structure that we use to solve our problem. In Section 5, we present an algorithm to solve the main problem and we prove its time complexity. In Section 6, we discuss the difficulty of sorting $\mathcal{D}$ in less than $O(n^2 \log n)$ time. Lastly, in Section 7 we present our conclusions.

2 Partial orders and blocking

Let X be a finite set, and $<$ a relation on the elements of X that satisfies the following conditions:

1. If $x < y$ and $y < z$ then $x < z$ (transitivity), and
2. $x \nless x$ (anti-reflexivity).

The set X together with $<$ is called a partial order, and 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 which has 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 following conditions are satisfied:

a) the elements of X are represented by points on the plane,

b) if y is a cover of x, the edge joining them is a monotonically increasing curve (with respect to the y-axis) starting at x and ending in y,

c) no edges of $P(<, X)$ intersect except perhaps at a common endpoint.

Given two elements $x, y \in X$, a *supremum* of x, y 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 partial order $P(<, X)$ is called a *truncated planar lattice* if the order that results when both a least and a greatest element are added to $P(<, X)$ is a planar lattice.

Let $C = \{c_1, \ldots, c_n\}$ be a set of disjoint closed bounded convex sets on the plane and $\alpha \in [0, 2\pi)$. 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.

Clearly, if c_j is an α-cover of c_i, then to uncover c_i along the α direction we need first to remove c_j. 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 $c_{\sigma(r)}$, with $r = 1, \ldots, k-1$ (Figure 2). The following observation will be useful:

Observation 2.1. *If $c_i \prec_\alpha c_j$, then $c_j \prec_{\alpha+\pi} c_i$.*

Clearly if $c_i \prec_\alpha c_j$ and $c_j \prec_\alpha c_k$, then $c_i \prec_\alpha c_k$. Since $c_i \not\prec_\alpha c_i$, C together with the blocking relation $\prec_\alpha$ is a partial order $P(\prec_\alpha, C)$. It is known that $P(\prec_\alpha, C)$ is a truncated planar lattice [16].

Observe that the diagram of $P(\prec_\alpha, C)$ 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). Since $\prec_\alpha$ is defined using α-coverings, the elements of C that we need to remove

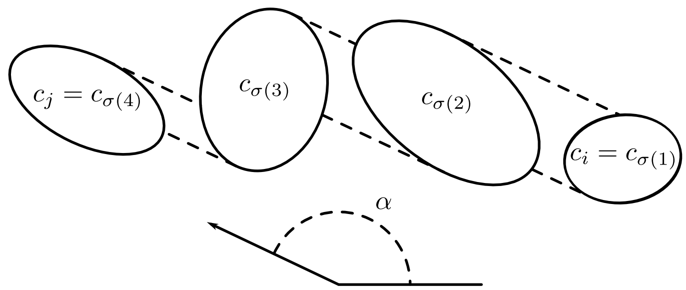

Figure 2: $c_{\sigma(r+1)}$ is an α-cover of $c_{\sigma(r)}$, $r = 1, 2, 3$, and $c_{\sigma(r)} \prec_\alpha c_{\sigma(s)}$, $r < s$. In particular, $c_i \prec_\alpha c_j$.

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 upper set of c_i in the α direction, or for short, the *α-up-set* of c_i. Thus our problem reduces to that of finding the direction α_0 such that the cardinality of the α_0-up-set of c_1 is minimized.

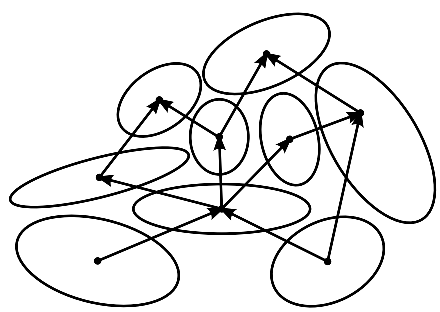

Figure 3: Diagram of $P(\prec_\alpha, C)$ for $\alpha = \pi/2$.

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 $\Theta(n^2)$ times.

In the next section we prove some properties of $P(\prec_\alpha, C)$ which will simplify the search space for α_0.

3 The critical directions

A line ℓ is called a *supporting line* of a closed convex set c if it intersects c, and c is contained in one of the closed half-planes determined by ℓ. In what follows, we will assume that no line is a supporting line of three or more elements of C, and that there are no two different parallel lines that each support two elements of C.

Given two closed convex sets c_i and c_j, a line ℓ is called an *internal tangent* of c_i and c_j if ℓ is a supporting line to both convex sets, and c_i is contained in one of the closed half-planes determined by ℓ while c_j is contained in the other. Similarly, a line ℓ is called an *external tangent* of c_i and c_j if ℓ supports them and c_i and c_j are contained in the same closed half-plane determined by ℓ (Figure 4).

Given $c_i, c_j \in C$, if we orient their common supporting lines from c_i to c_j, we can classify them as *left internal*, *right internal*, *left external*, and *right external* as in Figure 4. By definition, it is not hard to see that the internal tangents define critical directions where two convex sets can change their blocking relation.

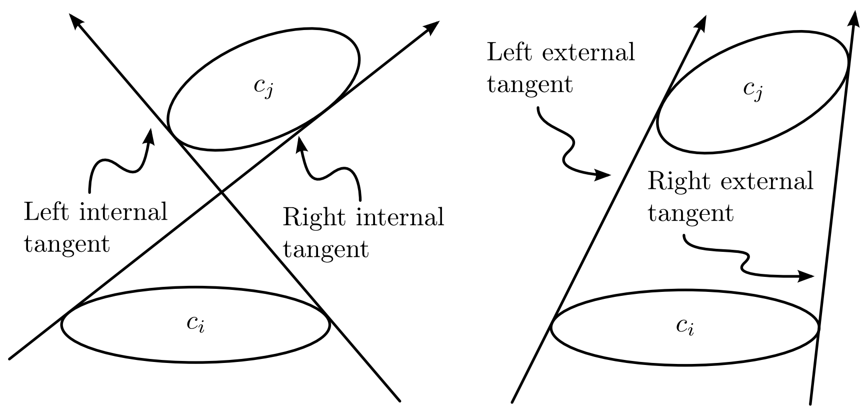

(a) Internal tangents of c_i and c_j. **(b)** External tangents of c_i and c_j.

Figure 4: Internal and external tangents of c_i and c_j.

Note that if α is the direction defined by a tangent of c_i and c_j, from c_i to c_j, then $\alpha + \pi$ is the direction of the same tangent of c_i and c_j, but directed from c_j to c_i.

Observation 3.1. *There are at most $4\binom{n}{2}$ distinct values of α where $P(\prec_\alpha, C)$ may change; these changes occur in the slopes defined by the internal tangents between pairs of elements of C (in both directions).*

Given $\alpha \in [0, 2\pi)$, and $\beta = \alpha + \theta$, $\theta \in [0, 2\pi)$, the interval $I = [\alpha, \beta]$ will denote the set of directions γ such that $\gamma = \alpha + \delta$, $0 \le \delta \le \theta$, addition taken mod 2π. Note that we consider that said directions grow counter-clockwise.

Although the changes in $P(\prec_\alpha, C)$ may only happen at directions defined by internal tangents, we also consider directions defined by external tangents as they will be used in Section 4.

Let $\mathcal{D} = \{\gamma_1, \ldots, \gamma_{8\binom{n}{2}}\}$ be the set of directions determined by the internal and external tangents of pairs of elements of C, and suppose that they are labeled in such a way that for $r < s$, $\gamma_r < \gamma_s$. Observe that if we change the value of α continuously from 0 to 2π, $P(\prec_\alpha, C)$ may change only when α

crosses an element of $\mathcal{D}$. Thus for any α and β in the open interval (γ_k, γ_{k+1}), $P(\prec_\alpha, C) = P(\prec_\beta, C)$, $\gamma_k, \gamma_{k+1} \in \mathcal{D}$. Thus for any $\alpha \in (\gamma_k, \gamma_{k+1})$, $P(\prec_\alpha, C)$ will be denoted by $P(\prec_{\gamma_k}, C)$.

Lemma 3.2. *If $c_i \prec_\alpha c_j$, then there is a direction $\beta \in (\alpha, \alpha + \pi)$ such that c_i and c_j are not comparable in $P(\prec_\beta, C)$.*

Proof. By Observation 2.1, $c_j \prec_{\alpha+\pi} c_i$. Suppose that the lemma is false; then there is a critical direction $\gamma_i \in \mathcal{D}$, with $\alpha < \gamma_i < \alpha + \pi$, such that $c_i \prec_{\gamma_i} c_j$ and $c_j \prec_{\gamma_{i+1}} c_i$. Observe that at each critical direction in $\mathcal{D}$ in which the partial order changes, either two elements from C stop being comparable or they become comparable. Thus if $c_i \prec_{\gamma_i} c_j$, then in γ_{i+1} we have that $c_i \prec_{\gamma_{i+1}} c_j$, or c_i and c_j are incomparable. In either case it cannot happen that $c_j \prec_{\gamma_{i+1}} c_i$. $\square$

We now prove:

Lemma 3.3. *Let c_i and c_j be two convex sets in C. The set of directions in which c_j blocks c_i forms a unique non-empty interval $\mathcal{I}_{i,j}$.*

Proof. If $c_i \prec_\alpha c_j$ for some α, by definition there is a sequence $S = c_{\sigma(1)}, c_{\sigma(2)}, \ldots, c_{\sigma(k)}$ of elements of C such that $c_{\sigma(r+1)}$ is an α-cover of $c_{\sigma(r)}$, with $c_i = c_{\sigma(1)}$, $c_j = c_{\sigma(k)}$ and $r = 1, \ldots, k-1$. Denote by $\mathcal{I}_{\sigma(r),\sigma(r+1)}$ the interval of directions determined by the right and left interior tangents of $c_{\sigma(r)}$ and $c_{\sigma(r+1)}$, with $r = 1, \ldots, k-1$. Observe that $c_{\sigma(r+1)}$ is a γ-cover of $c_{\sigma(r)}$ for all $\gamma \in \mathcal{I}_{\sigma(r),\sigma(r+1)}$. The sequence S determines then a set of directions $\mathcal{I}(S)$ along which c_i is blocked by c_j, where $\mathcal{I}(S) = \bigcap_{r=1}^{k-1} \mathcal{I}_{\sigma(r),\sigma(r+1)}$. Since at least $\alpha \in \mathcal{I}(S)$ and the intersection of intervals is an interval, $\mathcal{I}(S)$ is a non-empty interval.

When c_j blocks c_i in the direction α, there could be more than one sequence determining such blocking. Any two such blocking sequences, S and S', differ at least in an element and, in general, $\mathcal{I}(S_{i,j}) \neq \mathcal{I}(S'_{i,j})$.

If we now consider all the directions $\delta \in [0, 2\pi)$ where $c_i \prec_\delta c_j$, then there is a finite number, $m \geq 1$, of distinct blocking sequences $S_1, \ldots, S_m$ given by those directions; and each S_k determines a non-empty interval $\mathcal{I}(S_k)$ of directions. Let $\mathbf{S} = \{S_1, S_2, \ldots, S_m\}$. The set $\mathcal{I}_{i,j}$ of directions in which c_j blocks c_i, determined by all the sequences in $\mathbf{S}$, is then the non-empty set $\mathcal{I}_{i,j} = \bigcup_{k=1}^{m} \mathcal{I}(S_k)$. If $\mathcal{I}_{i,j}$ is in fact an interval, then our result holds.

By Lemma 3.2, there is a direction β where $c_i \not\prec_\beta c_j$. Without loss of generality, suppose that $\beta = 0$; thus $0 \notin \mathcal{I}(S_k)$ for each $S_k \in \mathbf{S}$. Let $\mathcal{I}(S_k) = [l_k, r_k]$ for each $S_k \in \mathbf{S}$, and let $I = [\theta_1, \theta_2]$ where $\theta_1 = \min\{l_1, \ldots, l_m\}$ and $\theta_2 = \max\{r_1, \ldots, r_m\}$. We will show that $\mathcal{I}_{i,j} = I$.

Clearly $\mathcal{I}_{i,j} \subseteq I$, hence it remains to be proved that $I \subseteq \mathcal{I}_{i,j}$. Let $\gamma \in [\theta_1, \theta_2] = I$, we will prove that $c_i \prec_\gamma c_j$, and therefore $I \subseteq \mathcal{I}_{i,j}$. Let $\mathcal{B}_i$ be the band enclosed between the two supporting lines of c_i in the γ direction. We have three cases:

1. The convex set c_j intersects $\mathcal{B}_i$. Thus, clearly $c_i \prec_\gamma c_j$.

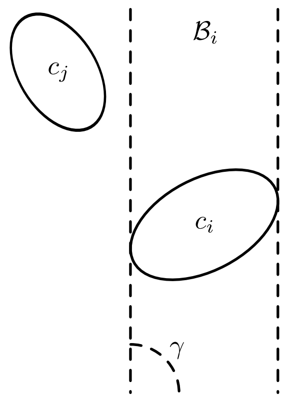

Figure 5: c_j to the left of $\mathcal{B}_i$.

2. The convex set c_j is to the left of $\mathcal{B}_i$ (Figure 5).

   Since $c_i \prec_{\theta_1} c_j$, we know that there is a sequence $c_i = c_{\sigma(1)}, c_{\sigma(2)}, \ldots, c_{\sigma(k_1)} = c_j$ such that $c_{\sigma(r+1)}$ is a θ_1-cover of $c_{\sigma(r)}$ for $r = 1, \ldots, k_1 - 1$.

   Since $c_{\sigma(2)}$ is a θ_1-cover of $c_i = c_{\sigma(1)}$, there is a line segment parallel to the direction θ_1 with endpoints in $c_i = c_{\sigma(1)}$ and $c_{\sigma(2)}$ such that it does not intersect any other convex set of C. Similarly, for $c_{\sigma(r)}$ and $c_{\sigma(r+1)}$ there is a line segment parallel to the direction θ_1 with endpoints in $c_{\sigma(r)}$ and $c_{\sigma(r+1)}$, $2 \leq r \leq k_1 - 1$ such that it does not intersect any other convex set in C. Each $c_{\sigma(r)}$, $2 \leq r \leq k_1 - 1$ contains two endpoints from two of these line segments, and these 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 consists of two types of line segments: Those parallel to the θ_1 direction, and those contained in $c_{\sigma(r)}$, $2 \leq r \leq k_1 - 1$. But $\theta_1 < \gamma$, so the first type always goes upwards and to the right of the γ direction. The second type may go to the right or to the left; see (Figure 6).

   By construction, such curve intersects $\mathcal{B}_i$, and then at least one element of $\{c_{\sigma(2)}, \ldots, c_{\sigma(k_1-1)}\}$ also intersects $\mathcal{B}_i$, say $c_{\sigma(s)}$, and thus $c_i \prec_\gamma c_{\sigma(s)}$.

   Denote by $\mathcal{B}_{\sigma(s)}$ the band bounded by the supporting lines of $c_{\sigma(s)}$ in the γ direction. If c_j intersects $\mathcal{B}_{\sigma(s)}$ then $c_{\sigma(s)} \prec_\gamma c_j$, and by transitivity, $c_i \prec_\gamma c_j$.

   Suppose then that c_j does not intersect $\mathcal{B}_{\sigma(s)}$. It is easy to see that c_j should lie to the left of $\mathcal{B}_{\sigma(s)}$. By substituting $c_{\sigma(s)}$ for c_i, and applying our previous argument repeatedly, we obtain a subsequence $\{c_i = c_{\sigma(i_1)}, c_{\sigma(i_2)}, \ldots, c_{\sigma(i_t)} = c_j\}$ of $\{c_{\sigma(1)}, c_{\sigma(2)}, \ldots, c_{\sigma(k_1)}\}$, with $i_1 < \cdots < i_t$, such that $c_{\sigma(i_1)} \prec_\gamma c_{\sigma(i_2)}, \ldots, c_{\sigma(i_{t-1})} \prec_\gamma c_{\sigma(i_t)}$, and thus $c_i \prec_\gamma c_j$.

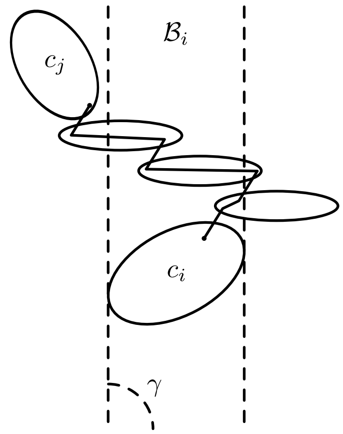

Figure 6: A sequence of θ_1-covers from c_j to c_i, and the curve that passes through the elements of the sequence.

3. The convex set c_j is to the right of the band $\mathcal{B}_i$. The proof is analogous to that of the previous case, but now using the direction θ_2 instead of θ_1.

Since in all the three cases we have that $c_i \prec_\gamma c_j$, then $I \subseteq \mathcal{I}_{i,j}$, and therefore $I = \mathcal{I}_{i,j}$, which is a unique non-empty interval where c_j blocks c_i. □

In what follows, we will show how to maintain $P(\prec_\alpha, C)$ as α changes from 0 to 2π in such a way that we can obtain $P(\prec_{\gamma_{k+1}}, C)$ from $P(\prec_{\gamma_k}, C)$, or more precisely, a *triangulation* $\mathcal{T}_{k+1}$ from $\mathcal{T}_k$ in constant time, $\mathcal{T}_k$ to be defined in the following section. This will enable us to obtain α_0 in $O(n^2)$.

4 α-triangulations

We observe first that our problem can be solved by calculating the truncated lattices $P(\prec_{\gamma_k}, C)$ for every direction $\gamma_k \in \mathcal{D}$, obtaining the γ_k-up-set of c_1 for each lattice, and then selecting a $\gamma_i \in \mathcal{D}$ which yields the smallest γ_i-up-set. Since calculating each $P(\prec_{\gamma_k}, C)$ can be done in $O(n \log n)$ [16], and $\mathcal{D}$ has $8\binom{n}{2}$ elements, this yields an $O(n^3 \log n)$ time algorithm to solve our problem.

To improve this complexity, we will show that $P(\prec_{\gamma_{k+1}}, C)$ can be obtained from $P(\prec_{\gamma_k}, C)$ (more precisely, the γ_k-triangulation to be described shortly) in constant time.

For $\alpha \in [0, 2\pi)$, 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 such that for each $c_i \in C$, $s \prec_\alpha c_i \prec_\alpha t$. We can picture t (respectively s) as a very large convex set blocking all of the elements of C (respectively blocked by all of the elements of C) in the direction α (Figure 7).

For each α, we now extend $P'(\prec_\alpha, C)$ to a triangulation $\mathcal{T}_\alpha$; that is, a planar multigraph where every internal face—except the external one—is a triangle.

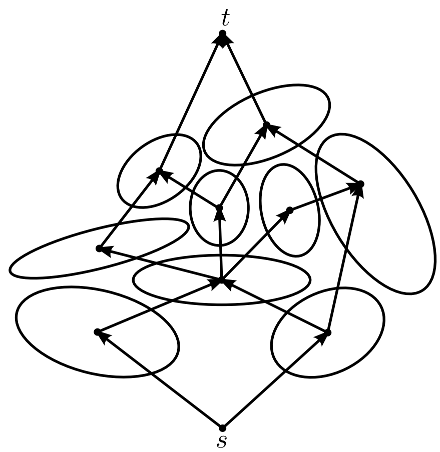

Figure 7: The lattice $P'(\prec_\alpha, C)$ for $\alpha = \pi/2$.

We will call $\mathcal{T}_\alpha$ the *α-triangulation* of C. To construct it, we use what we call α-visibility.

Given two convex sets $c_i, c_j \in C$, we will say that c_j is *α-visible* from c_i if there is an oriented line segment with direction α starting at a point on the boundary of c_i and ending at a point on the boundary of c_j such that it does not intersect any other convex set in C (Figure 8). Such line segments will be called (c_i, c_j) α-visibility line segments.

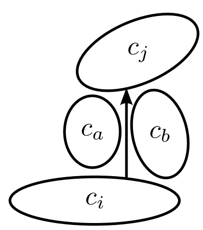

Figure 8: c_j is $\frac{\pi}{2}$-visible from c_i. Note that c_a is not $\frac{\pi}{2}$-visible from c_b, and vice versa.

If c_j is α-visible from c_i, the *α-visibility zone* of c_j and c_i is the union of all the (c_i, c_j) α-visibility line segments (Figure 9). Note that the visibility zone of c_i and c_j is not necessarily a connected region; see Figure 10(a).

It is important to remark that if c_j is α-visible from c_i, then $c_i \prec_\alpha c_j$, however it is not necessarily true that c_j is an α-cover of c_i. On the other hand, if c_j is an α-cover of c_i, then c_j is α-visible from c_i and their α-visibility zone is not empty and connected.

To obtain $\mathcal{T}_\alpha$ we proceed as follows: If c_j is α-visible from c_i, and c_j is not an α-cover of c_i, then we add to $P'(\prec_\alpha, C)$ an oriented arc from c_i to c_j for each connected component of the α-visibility zone of c_i and c_j. Each of these

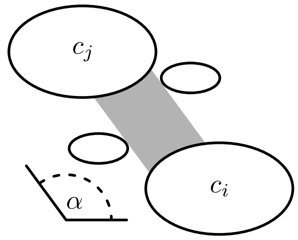

Figure 9: The shaded region is the α-visibility zone of c_j from c_i.

arcs can be drawn passing through their corresponding connected component of the α-visibility zone of c_j and c_i. Clearly $\mathcal{T}_\alpha$ is planar, and the above procedure yields an embedding of $\mathcal{T}_\alpha$ on the plane such that all of its faces, except for the external face, are triangular faces; see Figures 10(a), 10(b).

The extra arcs added to $P'(\prec_\alpha, C)$, will be called *α-visibility arcs*, to distinguish them from the regular arcs of $P'(\prec_\alpha, C)$. In particular, for each direction α, the source s and the sink t will always be joined by two visibility arcs bounding the external face of $\mathcal{T}_\alpha$. In all our figures, α-visibility arcs will be drawn with dashed curves, and the arcs from $P'(\prec_\alpha, C)$ with solid line segments. The triangulation $\mathcal{T}_\alpha$ arising from the lattice shown in Figure 7 is given in Figure 11.

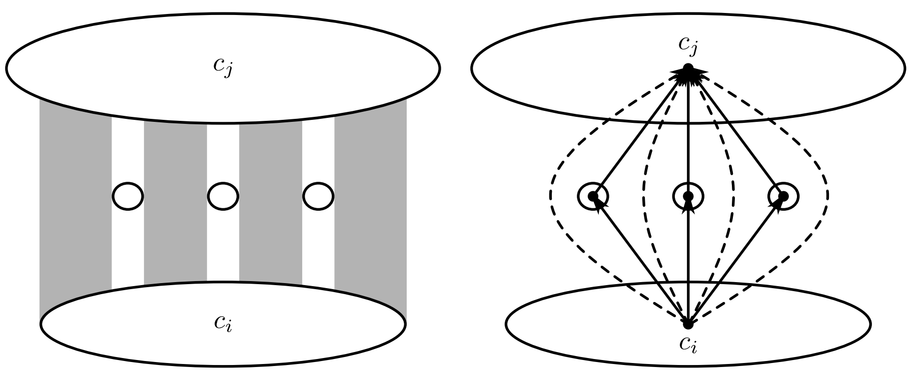

(a) Disconnected $\frac{\pi}{2}$-visibility zone of c_i and c_j.

(b) Oriented multi-edge connecting c_i and c_j, in dashed.

Figure 10: Visibility zone and its corresponding oriented multi-edge.

Observe that all the arcs in $\mathcal{T}_\alpha$ belong to two triangular *faces* of $\mathcal{T}_\alpha$, except for two arcs connecting s and t. Let e be an arc of $\mathcal{T}_\alpha$ that belongs to two triangular faces f and f' of $\mathcal{T}_\alpha$. Each of these faces contains a vertex (an element of C) that is not an endpoint of e. These elements will be called *opposite* elements with respect to e.

Since $\mathcal{D}$ has $8\binom{n}{2}$ elements, there are at most $8\binom{n}{2}$ triangulations $\mathcal{T}_\alpha$, $\alpha \in \mathcal{D}$. We now study how to obtain $\mathcal{T}_{\gamma_{k+1}}$ from $\mathcal{T}_{\gamma_k}$. We remark first that there are

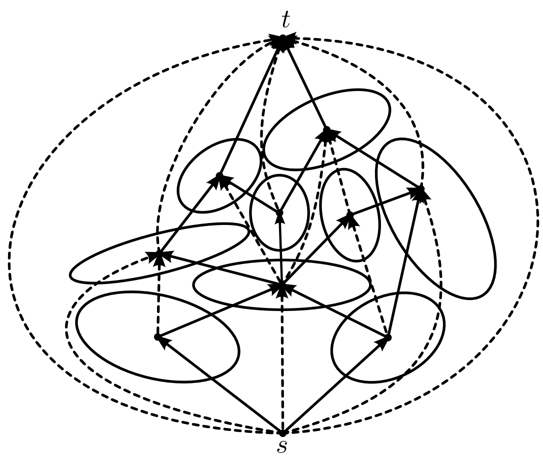

Figure 11: The triangulation $\mathcal{T}_\alpha$ for $\alpha = \frac{\pi}{2}$.

many cases in which $\mathcal{T}_{\gamma_k}$ and $\mathcal{T}_{\gamma_{k+1}}$ are the same; see Figure 12. It is also possible that $\mathcal{T}_{\gamma_k} \neq \mathcal{T}_{\gamma_{k+1}}$, but $P'(\prec_{\gamma_k}, C) = P'(\prec_{\gamma_{k+1}}, C)$. This could happen if $\mathcal{T}_{\gamma_{k+1}}$ differs from $\mathcal{T}_{\gamma_k}$ only in visibility arcs.

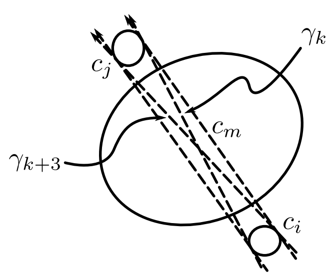

Figure 12: The triangulations $\mathcal{T}_{\gamma_k}$ through $\mathcal{T}_{\gamma_{k+3}}$ are the same, since the partial order does not change, and the vertices involved preserve their γ-visibility.

Let c_i and c_j be the elements of C that define γ_{k+1}. By definition, γ_{k+1} is parallel to one of the four tangents defined by c_i and c_j. If $\mathcal{T}_{\gamma_k}$ differs from $\mathcal{T}_{\gamma_{k+1}}$, then either γ_{k+1} is defined by an external tangent, and this caused a visibility change, or γ_{k+1} is defined by an internal tangent, and this caused a change in the partial order.

We will now prove that the difference between the triangulations $\mathcal{T}_{\gamma_k}$ and $\mathcal{T}_{\gamma_{k+1}}$ (if any) will be an arc flip, as defined in [14]; that is, we will remove an arc e from $\mathcal{T}_{\gamma_k}$, and replace it by another arc connecting two elements of C which are opposite with respect to e.

Lemma 4.1. *Let $\gamma_k, \gamma_{k+1} \in \mathcal{D}$, and let c_i and c_j be the convex sets defining γ_{k+1}. If $\mathcal{T}_{\gamma_k} \neq \mathcal{T}_{\gamma_{k+1}}$, then $\mathcal{T}_{\gamma_{k+1}}$ can be obtained from $\mathcal{T}_{\gamma_k}$ by flipping an arc in $\mathcal{T}_{\gamma_k}$. Such an arc flip involves arcs incident to c_i, c_j, or both of them.*

Proof. Without loss of generality suppose that $\gamma_{k+1} = \frac{\pi}{2}$ and that the tangent defining γ_{k+1} is oriented from c_i to c_j. Such tangent can be left external, right external, left internal, or right internal. However, the analysis for the case when γ_{k+1} is defined by a right external tangent is equivalent to the one for the left external case, and the same applies to the right internal and left internal cases.

- γ_{k+1} is defined by a left external tangent ℓ. Since $\mathcal{T}_{\gamma_k} \neq \mathcal{T}_{\gamma_{k+1}}$, no other element of C intersects ℓ between c_i and c_j (Figure 13(a)). Therefore there is a γ_{k+1}-visibility arc from c_i to c_j, and there is also a γ_k-visibility arc from c_i to c_j.

  Let c_a be the first element below c_i that intersects ℓ, and c_b the first element above c_j that intersects ℓ. If no element of C intersects ℓ below c_i, then $c_a = s$; similarly, if no element of C intersects ℓ above c_j, then $c_b = t$.

  It is not hard to see that c_j is γ_k-visible from c_a but not γ_{k+1}-visible from it (to the left of c_i) because c_i blocks any line segment parallel to γ_{k+1} between them. Also, c_b is γ_{k+1}-visible from c_i and not γ_k-visible from c_i (to the left of c_j) because c_j gets in the way of visibility. Finally, c_b is γ_k- and γ_{k+1}-visible from c_j, c_i is γ_k- and γ_{k+1}-visible from c_a, and c_a and c_b are γ_k- and γ_{k+1}-visible (Figure 13(b)).

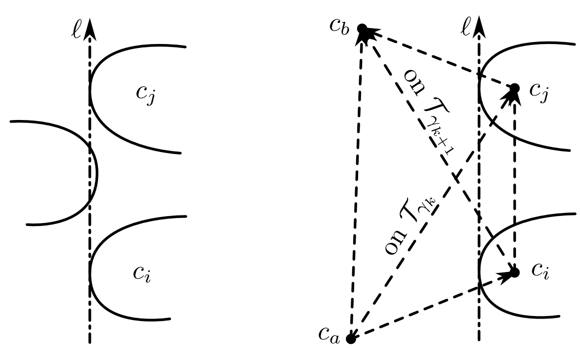

(a) There is no change in visibility for the direction γ_{k+1} when there is an element between c_i and c_j.

(b) The quadrangle defined by c_i, c_j, c_a, and c_b and the flip when going from γ_k to γ_{k+1}.

Figure 13: The case when γ_{k+1} is a left external tangent.

  In other words, the elements c_i, c_j, c_a, and c_b form a quadrangle in both $\mathcal{T}_{\gamma_k}$ and $\mathcal{T}_{\gamma_{k+1}}$, with the γ_k-visibility arc from c_a to c_j being a diagonal in γ_k of such a quadrangle, and this diagonal flips to the γ_{k+1}-visibility arc from c_i to c_b in $\mathcal{T}_{\gamma_{k+1}}$.

- γ_{k+1} is defined by a left internal tangent ℓ. Since $\mathcal{T}_{\gamma_k} \neq \mathcal{T}_{\gamma_{k+1}}$, no other element of C intersects ℓ between c_i and c_j (Figure 14(a)). Therefore there is a regular arc, defined in $P'(\prec_{\gamma_k}, C)$, from c_i to c_j.

Let c_a be the first element below c_i that intersects ℓ, and c_b the first element above c_j that intersects ℓ. If no element of C intersects ℓ below c_i, then $c_a = s$; similarly, if no element of C intersects ℓ above c_j, then $c_b = t$.

It is not hard to see that c_j is a γ_k-cover of c_i, but c_j is not a γ_{k+1}-cover of c_i, and the arc from c_i to c_j in $\mathcal{T}_{\gamma_k}$ is replaced by the γ_{k+1}-visibility arc from c_a to c_b in $\mathcal{T}_{\gamma_{k+1}}$. Finally, c_b is γ_k- and γ_{k+1}-visible from c_i and c_j, and c_i and c_j are γ_k- and γ_{k+1}-visible from c_a (Figure 14(b)).

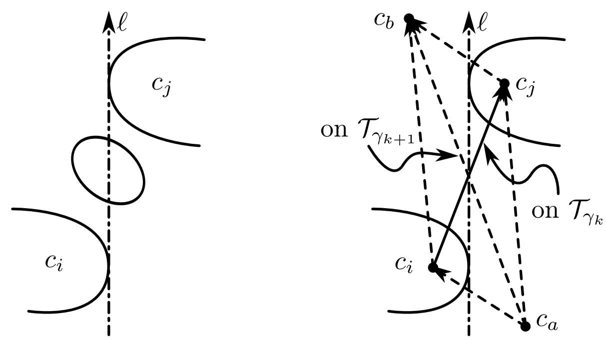

(a) There is no change in visibility for the direction γ_{k+1} when there is an element between c_i and c_j.

(b) The quadrangle defined by c_i, c_j, c_a, and c_b and the flip when going from γ_k to γ_{k+1}.

Figure 14: The case when γ_{k+1} is a left internal tangent.

In other words, the elements c_i, c_j, c_a, and c_b form a quadrangle in both $\mathcal{T}_{\gamma_k}$ and $\mathcal{T}_{\gamma_{k+1}}$, with the arc from c_i to c_j in $\mathcal{T}_{\gamma_k}$ being a diagonal of such quadrangle, and this diagonal flips to the γ_{k+1}-visibility arc from c_a to c_b in $\mathcal{T}_{\gamma_{k+1}}$.

Therefore $\mathcal{T}_{\gamma_{k+1}}$ can be obtained from $\mathcal{T}_{\gamma_k}$ by performing an arc flip, and our result holds. Even more, in each case we know if the arc to flip is an α-visibility arc or a regular arc, and if it flips to a regular arc or to an α-visibility arc. $\square$

Figure 15 shows an example of the arc flip performed to transform $\mathcal{T}_{\gamma_k}$ into $\mathcal{T}_{\gamma_{k+1}}$.

We will assume that for each direction γ in $\mathcal{D}$, we also have associated to it the two convex sets in C that define it. The next result follows:

Corollary 4.2. *Given $\mathcal{T}_{\gamma_k}$, we can obtain $\mathcal{T}_{\gamma_{k+1}}$ in $O(1)$ time.*

5 An algorithm to find α_0

In this section we prove that if we have the elements of $\mathcal{D}$ sorted then we can find the direction α_0 for which the up-set of c_1 is minimized in $O(n^2)$ time.

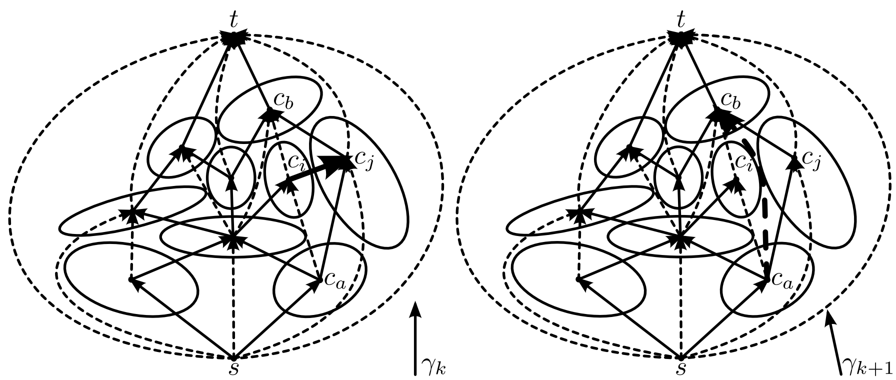

(a) The arc $c_i \to c_j$ before flipping. **(b)** The arc $c_a \to c_b$ after flipping.

Figure 15: Flipping of an arc when going from $\mathcal{T}_{\gamma_k}$ to $\mathcal{T}_{\gamma_{k+1}}$.

We will need two lemmas to prove our result.

Lemma 5.1. *For any element $c_i \in C$, as we go from γ_1 to $\gamma_{8\binom{n}{2}}$, the up-set of c_i changes $O(n)$ times.*

Proof. Let $c_i, c_j \in C$, $c_i \neq c_j$. By Lemma 3.3, the set of directions for which c_j blocks c_i is an interval $\mathcal{I}_{i,j}$. This means that as we go from γ_1 to $\gamma_{8\binom{n}{2}}$, c_j enters and leaves the up-set of c_i once. Therefore the up-set of c_i changes a linear number of times. □

Suppose next that for a direction $\gamma_k \in \mathcal{D}$ we have $\mathcal{T}_{\gamma_k}$, that c_1 and all the elements of C that belong to the up-set of c_1 are colored red, and the remaining elements of C are colored blue. We now show how we can detect in constant time whether the up-set of c_1 changes.

Lemma 5.2. *Given $\mathcal{T}_{\gamma_k}$ such that it vertices are colored as above, we can detect whether the up-set of c_1 changes in $\mathcal{T}_{\gamma_{k+1}}$ in constant time.*

Proof. Observe that if γ_{k+1} is defined by an external tangent of two elements $c_i, c_j \in C$, then $P'(\prec_{\gamma_k}, C) = P'(\prec_{\gamma_{k+1}}, C)$, and therefore the up-set of c_1 remains unchanged, and the coloring of C for γ_{k+1} is the same as that for γ_k.

Suppose then that γ_{k+1} is defined by an internal tangent of two elements c_i and c_j of C. Two cases arise depending on whether γ_{k+1} is a left or a right tangent of c_i and c_j.

Suppose first that γ_{k+1} is a right internal tangent. In this case, c_i and c_j are comparable in $\mathcal{T}_{\gamma_{k+1}}$. Assume without loss of generality that $c_i \prec_{\gamma_{k+1}} c_j$. If c_i and c_j were comparable in $\mathcal{T}_{\gamma_k}$, 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_i and c_j are not comparable in $\mathcal{T}_{\gamma_k}$. Observe first that if c_i is red, and c_j is blue, then c_j becomes comparable to c_1, and the up-set of c_1

changes. In all the other cases when γ_{k+1} is a right internal tangent, the up-set of c_1 remains unchanged.

Suppose next that γ_{k+1} is a left internal tangent of c_i and c_j. In this case, it must happen that c_i and c_j are comparable in $\mathcal{T}_{\gamma_k}$. Assume that $c_i \prec_{\gamma_k} c_j$. If both c_i and c_j are red, then c_j could leave the up-set of c_1, but only if it was a γ_k-cover of c_i. The case when c_i is red and c_j is blue in $\mathcal{T}_{\gamma_k}$ cannot happen, since $c_i \prec_{\gamma_k} c_j$. In all the other cases when γ_{k+1} is a left internal tangent, the up-set of c_1 remains unchanged.

Therefore, the up-set of c_1 can change only when γ_{k+1} is a right internal tangent and c_i is red and c_j is blue; or when γ_{k+1} is a left internal tangent and both c_i and c_j are red, and c_j is a γ_k-cover of c_i. We can test either case in constant time. For the second case, we can check if c_j is a γ_k-cover of c_i, and by Lemma 4.1 we know before each flip if that is the case. $\square$

Observe that if the up-set of c_1 does not change, then the red and blue coloring of the elements of $\mathcal{T}_{\gamma_k}$ is maintained in $\mathcal{T}_{\gamma_{k+1}}$ in the sense that the red elements in $\mathcal{T}_{\gamma_k}$ are the elements in the up-set of c_1 in $\mathcal{T}_{\gamma_{k+1}}$.

Theorem 5.3. *Suppose that we have the sorted set of directions $\mathcal{D} = \{\gamma_1, \ldots, \gamma_{8\binom{n}{2}}\}$, and that for each γ_k, we are also given the pair of elements c_i and c_j that generated it. Then we can find α_0 in $O(n^2)$.*

Proof. Construct $P'(\prec_{\gamma_1}, C)$ and $\mathcal{T}_{\gamma_1}$ in $O(n \log n)$ time. Next we calculate the γ_1-up-set of c_1 in $O(n)$ time by using BFS on $P'(\prec_{\gamma_1}, C)$.

By Corollary 4.2, we can obtain, one by one, the $\mathcal{T}_{\gamma_1}, \ldots, \mathcal{T}_{\gamma_{8\binom{n}{2}}}$ in overall quadratic time. By Lemma 5.2, we can find, also in overall quadratic time, the set of directions in which the up-set of c_1 changes. By Lemma 5.1, the up-set of c_1 changes a linear number of times. Each time this happens, we recolor the elements of our current partial order in linear time. Thus we can also maintain the coloring of the vertices of $\mathcal{T}_{\gamma_1}, \ldots, \mathcal{T}_{\gamma_{8\binom{n}{2}}}$ in quadratic time.

Therefore we can find α_0 in $O(n^2)$ time. $\square$

6 Some remarks about sorting $\mathcal{D}$

If we assume that for each pair of elements of C, we can calculate their tangent lines in constant time, then we can sort the elements of $\mathcal{D}$ in $O(n^2 \log n)$ time. A similar problem to that of sorting the elements of $\mathcal{D}$ arises from the problem of sorting the intersections generated by arrangements of curves on the plane.

A family of x-monotone Jordan curves is called *well behaved* if each time two curves intersect they cross each other, and any two curves intersect at most s times, where s is constant ([12], pages 399 and 404). In this context, it is also assumed that the intersections of any two curves can be calculated in constant time; this is usually referred to as being *under a proper model of computation* ([3, 15]).

It is known that for arrangements of well behaved curves with n elements in which any two of them intersect at most two times, the arrangement generated by them, including the set of all of their intersections, can be constructed in $O(n^2 \cdot 2^{\alpha(n)})$, where $\alpha(n)$ is the inverse of the Ackermann function [8]. However it is not known how to sort these intersections according to the x-axis in $o(n^2 \log n)$ time, even when we consider arrangements of lines.

The well known sorting $X+Y$ open problem (Problem 41 in [6]) says: Given two sets X and Y of numbers, each of size n, how quickly can the set $X+Y$ of all pairwise sums be sorted? In [13] it is proved that the sorting $X+Y$ problem is a particular case of the problem of sorting the intersections of an arrangement of lines according to the x-axis. The first reference to the sorting $X+Y$ problem was made in 1976, and it remains open: By the result proved in [13], sorting the intersections of an arrangement of lines according to the x-axis is an even stronger result.

In what follows, and to ease our presentation, we assume that the boundary of the elements of C is smooth. This avoids an unenlightening case analysis that leaves our results unchanged. To see that our problem can be reduced to that of sorting the intersection points of arrangements of well behaved curves in which any two of them intersect at most twice, we proceed as follows.

Let $c_i \in C$, and let U_i and L_i be the upper and lower chains of the boundary of C. Under the dual transformation which maps a non-vertical line ℓ defined by the equation $y = mx - n$ to the dual point $\ell^* = (m, n)$ and a point $p = (a, b)$ to the line $p^*: y = ax - b$, the points in U_i will be mapped to lines whose lower envelope will be a concave x-monotone curve that we will call U_i^*, and the points in L_i will be mapped to lines whose upper envelope will be a convex x-monotone curve that we will call L_i^*.

In the dual space, every line that intersects c_i is mapped to a point bounded from above by L_i^* and from below by a U_i^*, and every point inside c_i is mapped to a line enclosed between U_i^* and L_i^* respectively ([2, Section 7.4], (Figure 16).

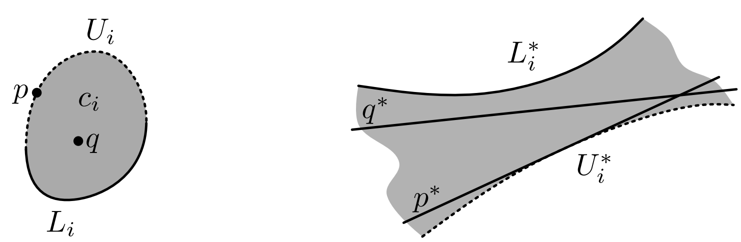

Figure 16: An element c_i of C and its mapping in the dual space.

Let $c_i, c_j \in C$. If a line ℓ is a tangent of c_i and c_j, then it intersects both convex sets. Without loss of generality, suppose that it does so in U_i and L_j. In the dual space this results in ℓ^* being the intersection point of U_i^* and L_j^*. For simplicity, we will assume that there are no vertical tangents between pairs of elements in C; we can always slightly rotate the whole set if necessary.

Let $\Gamma = \{U_i^*, L_i^* | i = 1, \ldots, n\}$ be an arrangement of curves. The next lemma follows:

Lemma 6.1. *Any two curves in Γ intersect at most twice.*

Proof. Let $c_i \in C$, and let ℓ be a non-vertical line tangent to c_i. Observe that if ℓ intersects c_i at a point in U_i, then c_i lies below ℓ; if ℓ intersects the boundary of c_i at a point in L_i, c_i is above ℓ.

Let $c_i, c_j \in C$, and let ℓ be a line tangent to both of them. If ℓ is an external tangent, then c_i and c_j are both contained in the same closed half-plane determined by ℓ. Therefore ℓ intersects c_i in L_i (respectively U_i), and ℓ intersects c_j in L_j (respectively U_j) (Figure 17).

If ℓ is an internal tangent to c_i and c_j, then one of them lies above ℓ and the other below it. Thus if ℓ intersects c_i in U_i (resp. L_i), it intersects c_j in L_i (resp. U_i).

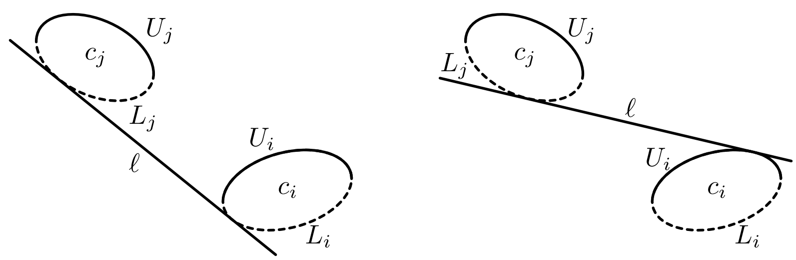

Figure 17: Intersections of a tangent ℓ with the upper and lower chains of two convex sets in C.

We now prove that any two $\tau_i^*, \tau_j^* \in \Gamma$ intersect at most twice. Suppose on the contrary that they intersect at least three times. Assume that τ_i^*, τ_j^* were generated by upper or lower chains of two elements $c_{i'}, c_{j'}$ of C. Since the intersection points of τ_i^*, τ_j^* correspond to common tangents of $c_{i'}, c_{j'}$, one of these tangents is an internal tangent, and the other an external tangent of $c_{i'}, c_{j'}$. But an internal tangent touches a lower and an upper chain of $c_{i'}, c_{j'}$, and an external chain intersects either two lower or two upper chains of $c_{i'}, c_{j'}$. Thus τ_i^*, τ_j^* intersect at most twice. $\square$

There are two ways in which the curves U_i^*, L_i^*, U_j^*, and L_j^* can intersect: If c_i is not contained in the vertical strip defined by the vertical tangents to c_j (or vice versa as in Figure 18(a)), then each pair of curves will intersect at most once (Figure 18(b)). If c_i is contained in the vertical strip defined by the vertical tangents to c_j (or vice versa; Figure 18(c)), then each pair of curves will intersect at most twice (Figure 18(d)).

Therefore the problem of calculating and sorting $\mathcal{D}$ is equivalent to calculating and sorting the intersections of an arrangement of curves that intersect each other at most twice.

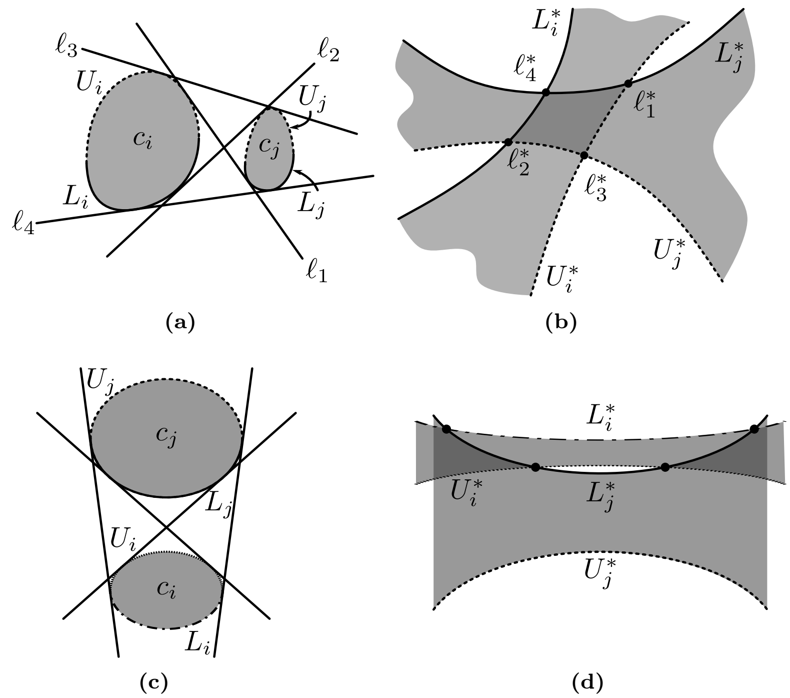

Figure 18: Top: Each pair of curves intersects at most once in the dual space. Bottom: Each pair of curves intersects at most twice in the dual space.

7 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 in the plane, find a direction α_0 minimizing the number of elements of C that have to be removed, along the direction α_0, in order to reach a particular element $c_1 \in C$. We present an $O(n^2)$ time algorithm to solve this problem, under the assumption we have the sorted set $\mathcal{D}$ of slopes of tangents to pairs of elements of C.

References

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

[2] K. P. Agarwal and M. Sharir. Davenport–Schinzel sequences and their geometric applications. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, chapter 1, pages 1–47. Nord Holland, 2000.

[3] J. Bentley and T. Ottmann. Algorithms for reporting and counting geometric intersections. *Computers, IEEE Transactions on*, C-28(9):643 –647, sept. 1979.

[4] B. Chazelle, T. Ottmann, E. Soisalon-Soininen, and D. Wood. The complexity and decidability of separation. In J. Paredaens, editor, *Automata, Languages and Programming*, volume 172 of *Lecture Notes in Computer Science*, pages 119–127. Springer Berlin / Heidelberg, 1984.

[5] N. G. de Bruijn. Aufgaben 17 and 18. *Nieuw Archief voor Wikskunde*, 2(954):67, 1954.

[6] E. D. Demaine, J. S. B. Mitchell, and J. O'Rourke. The open problems project. `http://cs.smith.edu/∼orourke/TOPP/Welcome.html`.

[7] A. Dumitrescu and M. Jiang. On reconfiguration of disks in the plane and related problems. *Computational Geometry: Theory and Applications*, 46(3):192–202, 2013.

[8] H. Edelsbrunner, L. Guibas, J. Pach, R. Pollack, R. Seidel, and M. Sharir. Arrangements of curves in the plane-topology, combinatorics, and algorithms. *Theoretical Computer Science*, 92(2):319 – 336, 1992.

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

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

[11] 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*, STOC '80, pages 154–160, New York, NY, USA, 1980. ACM.

[12] D. Halperin. Arrangements. In J. E. Goodman and J. O'Rourke, editors, *Handbook of discrete and computational geometry*, chapter 21, pages 389–412. CRC Press, Inc., Boca Raton, FL, USA, 1997.

[13] A. Hernández Barrera. Finding an $o(n^2 \log n)$ algorithm is sometimes hard. In *Proceedings of the 8th Canadian Conference on Computational Geometry*, pages 289–294. Carleton University Press, 1996.

[14] F. Hurtado, M. Noy, and J. Urrutia. Flipping edges on triangulations. *Discrete & Computational Geometry*, 22(3):333–346, 1999.

[15] F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag New York, Inc., New York, NY, USA, 1985.

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

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