On Piercing Sets of Objects *

Matthew J. Katz † Franck Nielsen ‡

Abstract

A set of objects is k-pierceable if there exists a set of k points such that each object is pierced by (contains) at least one of these points. Finding the smallest integer k such that a set is k-pierceable is NP-complete. In this paper, we present efficient algorithms for finding a piercing set (i.e., a set of k points as above) for several classes of convex objects and small values of k. In some of the cases, our algorithms imply known as well as new Helly-type theorems, thus adding to previous results of Danzer and Grünbaum who studied the case of axis-parallel boxes. The problems studied here are related to the collection of optimization problems in which one seeks the smallest scaling factor of a centrally symmetric convex object K, so that a set of points can be covered by k congruent homothets of K.

1 Introduction

Let $\mathcal{S}$ be a set of n d-dimensional convex objects. $\mathcal{S}$ is *k-pierceable* if there exists a set $\mathcal{P}$ of k points in $\mathcal{E}^d$ such that each object in $\mathcal{S}$ is pierced by (contains) at least one point in $\mathcal{P}$. Perhaps one of the most famous theorems in convex geometry is Helly's theorem [DGK63, GW93] that states that $\mathcal{S}$ is 1-pierceable (i.e., has a non-empty intersection) if and only if every subset of $\mathcal{S}$ of cardinality $d+1$ is 1-pierceable. The *Helly-number* $h = h(\mathcal{C}, P)$ associated with a class of objects $\mathcal{C}$ and a property P is the smallest integer (if such exists) so that for any set $\mathcal{S} \subseteq \mathcal{C}$ we have: if every subset of $\mathcal{S}$ of cardinality h has property P, then $\mathcal{S}$ also has property P. If such an integer does not exist, we set h by convention to ∞ (infinity). Let $\prod^k$ denote the k-pierceability property and $\mathcal{C}^d$ be the class of d-dimensional convex objects. Then Helly's theorem states that $h(\mathcal{C}^d, \prod^1) = d+1$. There are many Helly-type theorems in convex geometry (see [GW93] for an up-to-date survey).

Danzer and Grünbaum [DG82] studied the case of d-dimensional (axis-parallel) boxes $\mathcal{B}^d$ and obtained Helly-type theorems whenever they exist. They prove in particular that $h(\mathcal{B}^d, \prod^k) = \infty$, for $d, k \geq 3$. They conclude their paper with the following conjecture: $h(\mathcal{T}^d(K), \prod^2) < \infty$ if and only if K is a convex polytope, where $\mathcal{T}^d(K)$ is the class of all translates of K. This conjecture has been refuted very recently by Katchalski and

*Part of this work was done while the first author was visiting INRIA Sophia-Antipolis.

†Department of Computer Science, Utrecht University, P.O. Box 80.089, 3508 TB Utrecht, The Netherlands, E-mail: `matya@cs.ruu.nl`. Supported by the Dutch Organization for Scientific Research (N.W.O.).

‡INRIA, BP 93, 06902 Sophia-Antipolis cedex (France), E-mail: `Franck.Nielsen@sophia.inria.fr`.

Nashtir [KN96], who proved that $h(\mathcal{T}^2(K), \prod^2) = \infty$, for a centrally symmetric convex hexagon K.

The problems studied in this paper are related to a collection of optimization problems. Let $\mathcal{P}$ be a set of n points in $\mathcal{E}^d$ and K a centrally symmetric convex object. Let o be the center point of K and let K_λ be a homothet of K with scaling factor λ. K defines a semi-distance function as follows: For two points a and b, $d_K(a,b) = \lambda$, if K_λ is the smallest homothet of K that contains b when it is centered at a.

In these optimization problems we seek the smallest real λ such that $\mathcal{P}$ can be covered by k translates of K_λ. Note that these problems are always solvable by setting λ to an arbitrary large real. The decision problem associated with these optimization problems is the following. Given a value λ determine whether $\mathcal{P}$ can be covered by k translates of K_λ, i.e., try to find k appropriate loci for the (center points of the) k copies of K_λ. A point $p \in \mathcal{P}$ lies in a translate of K_λ centered at o iff $d_K(o,p) \leq \lambda$, and, since $d_K(\cdot,\cdot)$ is a symmetric function, iff $d_K(p,o) \leq \lambda$. Let $\mathcal{P}_{K_\lambda}$ be the set that is obtained from $\mathcal{P}$ by replacing each point $p \in \mathcal{P}$ with the translate of K_λ that is centered at p. Then, $\mathcal{P}$ can be covered by k translates of K_λ iff $\mathcal{P}_{K_\lambda}$ is k-pierceable. Generally, once the decision problem is solved efficiently, we link it with the parametric searching technique of Megiddo [Meg83], in order to obtain an efficient solution for the original optimization problem. For example, consider the 2-center problem. In this problem we wish to cover $\mathcal{P}$ by two congruent disks of minimum radius. The corresponding decision problem is: Given a fixed value for the radius, say, 1, determine whether $\mathcal{P}$ can be covered by two disks of radius 1. This is equivalent to the following problem. Let $\mathcal{P}_K$ be the set of unit disks centered at the points of $\mathcal{P}$, determine whether $\mathcal{P}_K$ is 2-pierceable. Recently, Sharir [Sha96] has presented a nearly linear algorithm for the 2-center decision problem, which immediately implies a nearly linear solution to the equivalent piercing problem. It is well known that if K is a centrally symmetric strictly convex object in the plane (e.g., a disk), then $h(\mathcal{T}^2(K), \prod^2) = \infty$ (see [HD60] for an elegant proof). This immediately implies that the 2-center decision problem is not an LP-type problem, since every LP-type problem implies a Helly-type theorem [Ame94] (see discussion below).

Amenta [Ame94] studied the relationships between LP-type problems and Helly-type theorems. (See [MSW92] for the definition of LP-type problems and an efficient linear-time randomized algorithm.) She showed that every LP-type problem has a corresponding Helly-type theorem but that the converse is not necessarily true. Nevertheless, she gave two paradigms for obtaining a LP-type problem from a Helly-type theorem. These paradigms fit the Helly-type theorems that are derived in this paper. We thus can obtain a randomized linear-time algorithm for the corresponding piercing problems (which can be derandomized, in turn, using the algorithm of Chazelle and Matoušek [CM93]). However, these algorithms are much more complicated than ours. Moreover, our algorithms consist of the proofs for the Helly-type theorems that we obtain.

Table 1 summarizes our results. We also show that $h(\mathcal{C}, \prod^2) = \infty$, where $\mathcal{C}$ is the class of 4-sided convex polygons. All these results are a significant improvement over the naïve solution that consists of computing the arrangement which is of size $O(n^d)$, and checking, for each subset of k cells, whether the underlying set of objects is pierced by this subset.

Objects	k	Time
homothetic triangles	2	$O(n)$ (Helly-type)
$4,5$-oriented polygons	2	$O(n \log n)$
d-dim. c-oriented polytopes	2	$O(n^{\min\{\lfloor \frac{c}{2} \rfloor, d\}} \log n)$
$(d+1)$-oriented simplices	2	$O(n^{\lceil \frac{d}{2} \rceil} \log n)$
d-dim. boxes	2	$O(n)$ (Helly-type)
homothetic triangles	3	$O(n \log n)$

Table 1: Summary of the results presented in this paper.

Thus, the naïve method requires $O(n^{dk+1})$ time (assuming k is a constant). The naïve solution can be slightly improved. Each cell of the arrangement is labelled with the subset of objects containing it. A cell is *maximal* if its corresponding subset is not contained in the subset of another cell. Clearly, if there exists a solution, then there also exists a solution that consists of at most k points drawn from maximal cells. Note, however, that the number of maximal cells can be as large as $\Omega(n^d)$. (Consider the following set of n convex objects: For each direction x_i, $i = 1, \ldots, d$, we take the n/d objects $[0, n/d]^d \cap (j < x_i < j+1)$, $j = 0, \ldots, n/d - 1$. Clearly, each d-cell of the arrangement of this set is maximal.) We can refine the naïve algorithm as follows. For each subset of $k-1$ cells, check whether the remaining set of objects that are not pierced by these cells has a non-empty intersection (this is an LP-type problem). Thus, we obtain an $O(n^{d(k-1)+1})$ expected time algorithm.

Since the problem of finding the minimum integer k so that a set of objects $\mathcal{S}$ is k-pierceable was shown to be NP-complete [Kar72, FPT81], many authors have focused on the problem of approximating k. There exist some polynomial-time algorithms for the latter problem with bounded error ratio [Chv79, Hoc82]. Bellare et al. [BGLR93] show that no polynomial-time algorithm can approximate the optimal solution within a factor of $(\frac{1}{8} - \epsilon) \log |\mathcal{S}|$, unless $NP \subseteq DTIME[n^{\log \log n}]$, where $\epsilon > 0$. This problem (in its non-geometric formulation) is called in the literature the *set cover* problem.

In this paper, we essentially study a larger class of objects (in comparison with Danzer and Grünbaum's paper), namely, the class of c-oriented convex polytopes, which contains the special case of all homothets of some c-facet polytope. An object O is *c-oriented* if it can be defined as the intersection of at most c translates of members of a given set of c halfspaces. For example, d-dimensional (axis-parallel) boxes are $2d$-oriented.

This paper is organized as follows. We first demonstrate our general method (Section 2) on the case of boxes and for $k = 2$, and derive a Helly-type theorem of Danzer and Grünbaum. We then describe this method (Section 3) for the case of c-oriented convex polytopes. In Section 4, several special cases are considered in which we are able to obtain more efficient solutions than those that are implied by the general method. In this section we obtain a new Helly-type theorem (this theorem with the exact Helly-number was also discovered independently by Katchalski and Nashtir [KN96]). Finally, we conclude in Section 5.

2 Piercing Sets of Boxes with 2 Points

Let $h(d,k)$ be the Helly-number for d-dimensional (axis-parallel) boxes and k points, that is, $h(d,k)$ is the smallest integer (if such exists) such that, if every subset of cardinality $h(d,k)$ of a set $\mathcal{B}$ of d-boxes is k-pierceable, then $\mathcal{B}$ is k-pierceable. If such an integer does not exist, we set $h(d,k)=\infty$. Danzer and Grünbaum [DG82] have proven the following theorem, which we present as a table:

Theorem 1 (Danzer and Grünbaum)

$h(d,k)$	$k=1$	$k=2$	$k=3$	$k\geq 4$
$d=1$	2	3	4	$k+1$
$d=2$	2	5	16	∞
$d\geq 3$	2	$\begin{cases} 3d & \textit{if } d \textit{ is odd} \\ 3d-1 & \textit{if } d \textit{ is even} \end{cases}$	∞	∞

Their proof for the case of d-dimensional boxes, $d\geq 2$, and two points is constructive and induces a simple algorithm. We briefly describe this algorithm in order to compare it subsequently with ours. If there exists a hyperplane $x_i=c$ that meets all the boxes of $\mathcal{B}$, then $\mathcal{B}$ is 2-pierceable if and only if the $(d-1)$-dimensional set of boxes $\mathcal{B}'=\{B\cap(x_i=c)|B\in\mathcal{B}\}$ is 2-pierceable. Hence, we may assume w.l.o.g. that the projection $\mathcal{P}_i$ of $\mathcal{B}$ onto the x_i axis yields at least one pair of disjoint segments, for $i=1,\ldots,d$. Denote by a_i (resp. b_i) the lower bound (resp. upper bound) of the right (resp. left) endpoints of the segments of $\mathcal{P}_i$. W.l.o.g., assume that $a_i=0$ and $b_i=1$, for $i=1,\ldots d$, and put $v=(1,1,...,1)$. Then, $\mathcal{B}$ is 2-pierceable if and only if there are two opposite points $p\in\{0,1\}^d$ and $v-p$ that pierce $\mathcal{B}$. This gives us an $O(2^d n)$ time algorithm. Note that if $\mathcal{B}$ is not 2-pierceable, then we can obtain a counterexample of size $O(2^d)$ from which we can select a subset of size $3d$ (for odd d) or $3d-1$ (for even d) in $O(1)$ time, for fixed dimension d.

We describe in this section a more general and algorithmic approach to the k-piercing problem, and demonstrate it for the case of d-boxes with 2 points. (The algorithm above of Danzer and Grünbaum is very specialized and it cannot be generalized to handle other classes of objects.) Our approach does not yield the exact Helly-numbers but only upper bounds. In the subsequent sections, we will apply our approach to sets of more complex objects to obtain new results. We begin with the case of axis-parallel rectangles (2-dimensional boxes) in order to focus on the method.

2.1 2-dimensional boxes (rectangles)

Let $\mathcal{B} = \{B_1, ..., B_n\}$ be a set of n 2-boxes. We wish to find a piercing set for $\mathcal{B}$ which consists of 2 points, if such a set exists. In other words, we wish to solve the 2-piercing problem for $\mathcal{B}$. We first check whether $\mathcal{B}$ is 1-pierceable by computing the region $\cap_{k=1}^{n} B_k$ (which is of course also a box) incrementally. If at some stage, j, the current region becomes empty, then obviously $\mathcal{B}$ is not 1-pierceable. Denote the box $\cap_{k=1}^{j-1} B_k$ by R. Since $R \cap B_j = \emptyset$, there exists either a vertical line or a horizontal line that separates between R and B_j. Assume for example that the separating line is vertical and that B_j lies to the right of this line, then the box B_i of $\mathcal{B}$ whose right edge contains the right edge of R and B_j are disjoint. We have thus found in linear time two disjoint boxes B_i and B_j in $\mathcal{B}$ (and have proven the well known result: $h(2, 1) = 2$). Clearly, if $\mathcal{B}$ is 2-pierceable, then one of the two piercing points must lie in B_i and the other in B_j.

Consider the following partition of $\mathcal{B}$ into four subsets:

- $\mathcal{B}_1 = \{B \in \mathcal{B} | B \cap B_i \neq \emptyset \text{ and } B \cap B_j = \emptyset\}$
- $\mathcal{B}_2 = \{B \in \mathcal{B} | B \cap B_i = \emptyset \text{ and } B \cap B_j \neq \emptyset\}$
- $\mathcal{B}_3 = \{B \in \mathcal{B} | B \cap B_i = \emptyset \text{ and } B \cap B_j = \emptyset\}$
- $\mathcal{B}_4 = \{B \in \mathcal{B} | B \cap B_i \neq \emptyset \text{ and } B \cap B_j \neq \emptyset\}$

Clearly, if $\mathcal{B}_3 \neq \emptyset$ then $\mathcal{B}$ is not 2-pierceable, and we have a counterexample of size 3. Assume, therefore, that $\mathcal{B}_3 = \emptyset$. Also, all boxes in $\mathcal{B}_1$ must be pierced in B_i, and all boxes in $\mathcal{B}_2$ must be pierced in B_j. Put $B'_1 = \cap \mathcal{B}_1$ and $B'_2 = \cap \mathcal{B}_2$. If one of these boxes is empty, then $\mathcal{B}$ is not 2-pierceable (and we have a counterexample of size 3), and if both are non-empty and $\mathcal{B}_4 = \emptyset$ then we are done.

Now, choose the box $B \in \mathcal{B}_4$ with the highest bottom edge b. B must be pierced either in B'_1 or in B'_2. If B is pierced in B'_k, $1 \leq k \leq 2$, then clearly the best place to put a piercing point inside B'_k is at the intersection point formed by b and the vertical edge of B'_k that is closer to the separating vertical line. We first try to pierce B in B'_1 (if possible). We then check whether the set of all boxes in $\mathcal{B}$ that are still not pierced is 1-pierceable. If this does not yield a solution (i.e., this set is not 1-pierceable and we obtain two disjoint boxes as above), we try the analogue case. If this too does not yield a solution then we may conclude that $\mathcal{B}$ is not 2-pierceable, and we have obtained a counterexample of size at most 13. (The four boxes defining B'_1, plus the four boxes defining B'_2, plus B, plus the two disjoint boxes that are obtained when piercing B in B'_1, and the two disjoint boxes that are obtained when piercing B in B'_2.)

We have described a linear-time algorithm for the 2-piercing problem for a set of planar boxes. Our algorithm immediately implies a Helly-type theorem, since it always provides a counterexample of size at most 13 whenever the input set is not 2-pierceable. (A more careful inspection shows that the maximum size of a counterexample is only 5, see Theorem 1.)

2.2 d-dimensional boxes

In this section we describe our algorithm for the 2-piercing problem for a set $\mathcal{B}$ of n d-dimensional boxes. This algorithm will imply that $h(d, 2)$ is finite. As mentioned, the exact value for $h(d, 2)$ was given by Danzer and Grünbaum. However, our emphasis here is on the general method which will lead to new results in the subsequent sections.

As before, we first check in linear time whether $\mathcal{B}$ is 1-pierceable, and obtain two disjoint boxes B_i and B_j, if it is not. A box B is a special type of polytope; it is defined by $2d$ halfspaces of the form $x_i \geq a_i$ or $x_i \leq b_i$, for $1 \leq i \leq d$ and real values a_i and b_i. We will distinguish between the 'orientations' of opposite facets of B; thus B has $2d$ facets of $2d$ different orientations.

Our algorithm consists of two stages. In the first stage we obtain a collection of at most 2^{2d} pairs of disjoint regions (boxes) $\mathcal{C} = \{(E_1, F_1), \ldots, (E_m, F_m)\}$ with the following two properties: (i) If there exists a solution, then there exists a pair $(E_k, F_k) \in \mathcal{C}$ and points $p \in E_k$ and $q \in F_k$ such that $\{p, q\}$ is a piercing set for $\mathcal{B}$. (ii) Let $\mathcal{A}(B)$ denote the arrangement of $\mathcal{B}$ restricted to the (open) box B. Then the arrangements $\mathcal{A}(E_k)$ and $\mathcal{A}(F_k)$ do not have a common facet orientation. That is, if $\mathcal{A}(E_k)$ has a facet of some orientation o, then $\mathcal{A}(F_k)$ does not have a facet of this orientation (and vice versa). In the second stage, for each of the pairs (E, F) in $\mathcal{C}$, we search for a solution in which one of the points lies in E and the other lies in F.

We begin the first stage with the disjoint regions B_i and B_j. If there exists a third box in $\mathcal{B}$ which is disjoint from both B_i and B_j, then clearly $\mathcal{B}$ is not 2-pierceable, and we are done. Let o denote a facet orientation that appears in both the clipped arrangements $\mathcal{A}(B_i)$ and $\mathcal{A}(B_j)$. We show below how to eliminate o from one of them. More precisely, we will replace the pair (B_i, B_j) by two pairs of regions such that the first region in both these pairs is contained in B_i, the second region is contained in B_j, and for each of these pairs the orientation o appears in only one of the corresponding clipped arrangements, and property (i) above holds for the collection consisting of these two pairs. Among all the facets of orientation o of the clipped arrangements $\mathcal{A}(B_i)$ and $\mathcal{A}(B_j)$, choose the 'topmost' facet f (i.e., if the halfspaces corresponding to these facets are of the form $x_i \geq a_i$, choose the facet with largest a_i, and if they are of the form $x_i \leq b_i$, choose the facet with smallest b_i). Let B be the box that is partially defined by f. In any solution, B is either pierced in B_i or in B_j. If it is pierced in B_k, $1 \leq k \leq 2$, then we can replace the region B_k by the smaller region $B_k \cap B$ for which the orientation o does not appear in its corresponding clipped arrangement. We thus replace the pair (B_i, B_j) by the (at most) two pairs $(B_i \cap B, B_j)$ and $(B_i, B_j \cap B)$. Note that it is necessary to choose a facet orientation o that appears in both clipped arrangements. We apply this procedure now to the (at most) two pairs that were obtained (thus obtaining (at most) four new pairs) and so on, until none of the current pairs can be further processed. At the end of this *tree process* we are left with a collection $\mathcal{C}$ of at most 2^{2d} pairs of regions so that properties (i) and (ii) hold.

In the second stage, for each of the pairs $(E, F) \in \mathcal{C}$, we search for a solution in which one of the points lies in E and the second in F. Consider a pair $(E, F) \in \mathcal{C}$. Our goal is to shrink the two regions E and F without loosing all the solutions (if such exist), until eventually one of them becomes a cell of the arrangement of $\mathcal{B}$. We then can place a point anywhere inside this cell, and check whether the set of boxes in $\mathcal{B}$ that are not pierced by this point is 1-pierceable.

Assume w.l.o.g. that E and F are vertically separable. That is, there exists a vertical hyperplane h such that E lies on one of its sides and F lies on the other. As in the planar case, we partition the set $\mathcal{B}$ into the four subsets $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3, \mathcal{B}_4$ (with respect to E and F), and compute the regions B'_1 and B'_2. The problematic boxes are those boxes in $\mathcal{B}_4$ that still intersect both B'_1 and B'_2, since such boxes can be pierced both in B'_1 and in B'_2. The hyperplane h intersects all but the two vertical facets of these boxes. Consider a pair of opposite facet orientations o and $\overline{o}$. By property (ii) above, o ($\overline{o}$) appears in at most one of the two corresponding clipped arrangements $\mathcal{A}(B'_1)$ and $\mathcal{A}(B'_2)$. If o and $\overline{o}$ do not appear in the same clipped arrangement, then we can eliminate both of them easily (see Figure 1 – Case 1). Assume, for example, that $\overline{o}$ appears in $\mathcal{A}(B'_1)$. We can replace the region B'_1 (resp. B'_2) by the region $B'_1 \cap \overline{s}$ (resp. $B'_2 \cap s$), where $\overline{s}$ (resp. s) is the topmost halfspace (in the sense explained above) corresponding to a facet of orientation $\overline{o}$ (resp. o) in $\mathcal{A}(B'_1)$ (resp. $\mathcal{A}(B'_2)$), without loosing all solutions, since all the closing facets of orientation o (resp. $\overline{o}$) of boxes in $\mathcal{B}_4$ that intersect B'_1 (resp. B'_2) are necessarily 'below' the closing facet of B'_1 (resp. B'_2) of this orientation. So now assume that o and $\overline{o}$ appear in the same clipped arrangement, say, $\mathcal{A}(B'_1)$. Let s (resp. $\overline{s}$) be the topmost halfspace corresponding to a facet of orientation o (resp. $\overline{o}$) in $\mathcal{A}(B'_1)$. If $s \cap \overline{s} \neq \emptyset$, then again we can replace the region B'_1 by the region $B'_1 \cap (s \cap \overline{s})$ without loosing any solution (see Figure 1 – Case 2). If, however, $s \cap \overline{s} = \emptyset$, then, if B'_2 is disjoint from both these halfspaces, then clearly there is no solution (see Figure 1 – Case 3), and, otherwise (B'_2 is contained in one of these halfspaces), one of the two boxes that are partially defined by these halfspaces, say the one that is defined by $\overline{s}$, must be pierced in B'_1 since it does not intersect B'_2, so we replace B'_1 by the region $B'_1 \cap \overline{s}$ (see Figure 1 – Case 4) .

We now consider another pair of opposite orientations, and so on, until one of our two regions reduces to a cell of the arrangement of $\mathcal{B}$, and we can conclude as described above. To summarize, we have presented a linear-time algorithm for the 2-piercing problem for the set $\mathcal{B}$. Moreover, we have shown that $h(d,2)$ is finite, since our algorithm will always provide a counterexample of bounded constant size.

Remark 1: The above tree process can be applied to any number of disjoint regions to remove an orientation that appears in all the corresponding clipped arrangements.

Remark 2: It is easy to transform the linear algorithm that is induced by the proof of Danzer and Grünbaum (described at the beginning of this section) into a linear algorithm for the 3-piercing problem for rectangles in the plane. Combining the approach of this algorithm with some dynamic data structures, Sharir and Welzl [SW96] have recently obtained $O(n \text{ polylog}(n))$ solutions to the 4- and 5-piercing problems for rectangles in the plane.

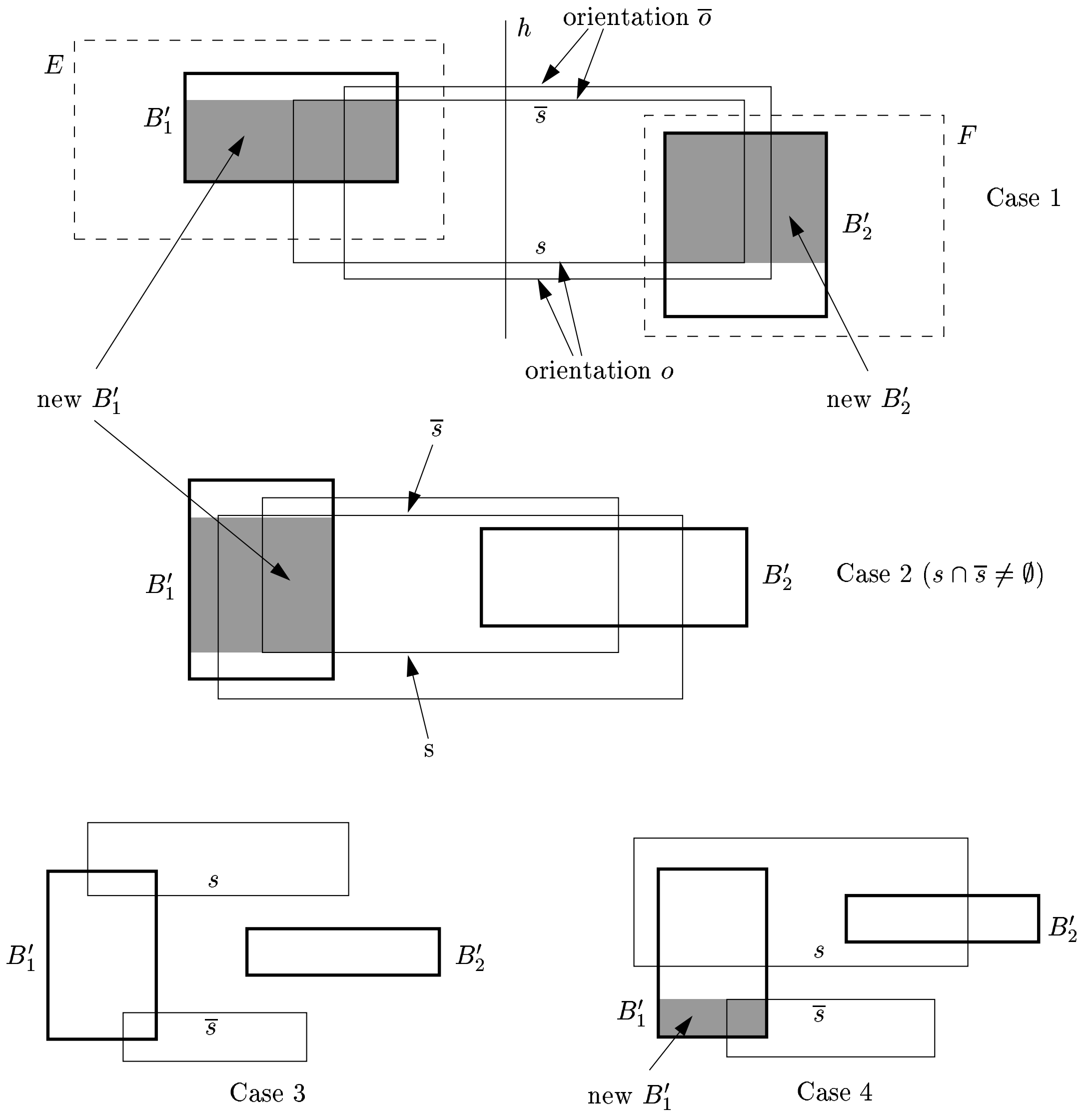

Figure 1: Considering a pair of opposite orientations o and $\overline{o}$ in (B_1', B_2').

3 Piercing Sets of c-Oriented Polytopes with 2 Points

Let $\mathcal{H}$ be a set of c halfspaces in $\mathcal{E}^d$, for some constant c. In this section, we consider the class $\mathcal{C}_\mathcal{H}$ of c-oriented convex polytopes above $\mathcal{H}$ (that is, every polytope in $\mathcal{C}_\mathcal{H}$ can be defined as the intersection of at most c translates of halfspaces in $\mathcal{H}$). Let $\mathcal{S}$ be a subset of $\mathcal{C}_\mathcal{H}$ of size n. The main result of this section is a general algorithm for the 2-piercing problem for $\mathcal{S}$, whose time complexity is roughly $O(n^{\min\{\lfloor \frac{c}{2} \rfloor, d\}})$. Thus, whenever $\lfloor \frac{c}{2} \rfloor < d$ (e.g., when $c = d+1$ and $\mathcal{S}$ is a set of simplices), this algorithm is much faster than the naïve algorithm described in the introduction.

Piercing $\mathcal{S}$ with a single point. As the intersection of any subset of $\mathcal{S}$ is either empty or a member of $\mathcal{C}_\mathcal{H}$ (and thus defined by at most c objects of $\mathcal{S}$), we can compute the intersection $R = \cap\mathcal{S}$ in linear time. More precisely, we process the objects of $\mathcal{S}$ one by one, spending $O(c)$ time per object. If $R \neq \emptyset$, then $\mathcal{S}$ is 1-pierceable, and we may pick any point in R. If $R = \emptyset$, then we obtain a subset $\mathcal{S}'$ of cardinality at most $c+1$ such that $\cap\mathcal{S}' = \emptyset$. This implies that $h(\mathcal{C}_\mathcal{H}, \prod^1) \leq \min\{c+1, d+1\}$. However, the case $c < d$ is not interesting, since, in this case, either $\cap\mathcal{S} \neq \emptyset$ or it is possible to decrease the dimension. We can obtain in constant time a minimal subset $\mathcal{S}'' \subseteq \mathcal{S}'$ of size at most $d+1$ such that $\cap\mathcal{S}'' = \emptyset$.

Piercing $\mathcal{S}$ with two points. We distinguish between two cases: $\lfloor \frac{c}{2} \rfloor < d$ and $\lfloor \frac{c}{2} \rfloor \geq d$. For the former case we present below a roughly $O(n^{\lfloor \frac{c}{2} \rfloor})$ time algorithm, while for the latter case we basically apply the naïve algorithm described in the introduction (exploiting the fact that the underlying objects are c-oriented), and obtain a roughly $O(n^d)$ time algorithm. We begin with the case $\lfloor \frac{c}{2} \rfloor < d$. We first check whether $R = \cap\mathcal{S}$ is empty. If $R \neq \emptyset$ then $\mathcal{S}$ is 1-pierceable. Otherwise, we obtain a subset $\mathcal{S}''$, as above, of size at most $d+1$ so that $\cap\mathcal{S}'' = \emptyset$. We now consider all the (constant number of) ways to partition $\mathcal{S}''$ into two subsets $\mathcal{S}_1$ and $\mathcal{S}_2$ such that $X = \cap\mathcal{S}_1 \neq \emptyset$ and $Y = \cap\mathcal{S}_2 \neq \emptyset$. Clearly, the regions X and Y are disjoint. For each such partition (producing the pair of disjoint regions X and Y) we proceed as follows. Apply the 'tree process' of stage 1 of the algorithm of the previous section for the case of d-dimensional boxes. At the end of this stage, we are left with a collection of at most 2^c pairs of regions for which properties (i) and (ii) of the previous section hold. Consider each of these pairs (E, F) separately. As before let $\mathcal{A}(R)$ denote the arrangement of $\mathcal{S}$ within the region R, then property (ii) assures that the arrangements $\mathcal{A}(E)$ and $\mathcal{A}(F)$ do not have facets of the same orientation, and therefore at least one of these arrangements, say, $\mathcal{A}(E)$, does not have more than $\lfloor \frac{c}{2} \rfloor$ facet orientations. Thus the arrangement $\mathcal{A}(E)$ has only faces of dimension k for $d - \lfloor \frac{c}{2} \rfloor \leq k \leq d$, and its combinatorial complexity is therefore only $O(n^{\lfloor \frac{c}{2} \rfloor})$. We compute the arrangement $\mathcal{A}(E)$, and traverse its cells, moving from a cell to an adjacent cell, and maintaining the intersection of the remaining objects that are not pierced by the current cell. (This is done by maintaining c sorted lists, one per orientation; an insert/delete operation costs $O(\log n)$.) This gives us an algorithm of time complexity $O(n^{\lfloor \frac{c}{2} \rfloor} \log n)$. Notice that if $\mathcal{S}$ consists of simplices of $d+1$ orientations, then this algorithm yields an $O(n^{\lceil d/2 \rceil} \log n)$ time algorithm. The following theorem summarizes the main result of this section.

Theorem 2 *Let $\mathcal{S}$ be a set of n c-oriented convex polytopes in $\mathcal{E}^d$, where $\lfloor \frac{c}{2} \rfloor < d$. It is possible to solve the 2-piercing problem for $\mathcal{S}$ in $O(n^{\lfloor \frac{c}{2} \rfloor} \log n)$ time. In particular, if $\mathcal{S}$*

consists of simplices of $d+1$ orientations, then the running time is $O(n^{\lceil d/2\rceil}\log n)$.

For the case $\lfloor \frac{c}{2} \rfloor \geq d$, we apply the naïve solution that was described in the introduction, exploiting the additional property that the intersection of any subset of $\mathcal{S}$ is either empty or is a c-oriented polytope. That is, we compute the arrangement of $\mathcal{S}$ and traverse its cells, moving from one cell to an adjacent cell, and maintaining as above the intersection of the remaining set of objects that are not pierced by the current cell. The running time of this algorithm is $O(n^d \log n)$.

In the next section we mention a few cases for which we are able to obtain algorithms that are more efficient that those that are implied by the general algorithm described above.

4 Applications

This section is divided into three parts. In the first part we obtain a new Helly-type theorem as a corollary of the general method described in the previous section. In the second part we show that for the class $\mathcal{C}$ of 4-sided convex polygons $h(\mathcal{C},\prod^2)=\infty$. Finally, in the third part, we mention several cases for which better bounds than those that are implied by the general method are attainable, due to some 'boundary property' that holds in these cases.

4.1 A new Helly-type theorem for homothetic triangles

Consider the case of 3-oriented polygons in the plane, i.e, homothetic triangles, and the appropriate version of the 2-piercing problem. We apply the method of the previous section to obtain a collection of at most $3\times 6=18$ pairs (A_i,B_i) of disjoint regions, such that, in each of them at least one of the regions, say, A_i, is crossed only by edges of a single orientation. Now, considering a pair (A,B), it is clear where to pick the piercing point in A, and we check whether the remaining set of triangles that are not pierced by this point is 1-pierceable. Thus, we have a linear-time algorithm for the 2-piercing problem for this case. Moreover, the algorithm implies a Helly-type theorem, since, if there is no solution, then it is always possible to find a counterexample of size at most $18\times(3+4)$, so that the Helly-number h is at most 126.

Theorem 3 *Let $\mathcal{T}$ be a class consisting of all homothets of a fixed triangle in the plane. Then $h(\mathcal{T},\prod^2)\leq 126$.*

Remark 1: As mentioned above, this theorem was also discovered independently by Katchalski and Nashtir [KN96], who proved that $h(\mathcal{T},\prod^2)=9$ in a direct way.

Remark 2: In order to prove a similar result for the case of homothetic quadrilaterals, i.e., all homothets (or even translates) of a fixed convex quadrilateral, we must specify a way to pick a 'best' point as above in one of the regions of a pair of regions that is obtained by the method of the previous section. However, this does not seem possible. In section 4.3 we consider a slightly more general case, the case of 5-oriented convex polygons, and obtain an $O(n\log n)$-time algorithm.

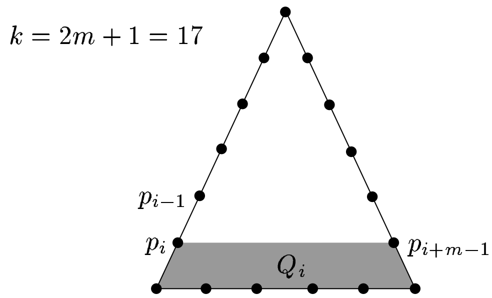

Figure 2: The construction for $k = 17$. Every proper subset is 2-pierceable, while the whole set requires 3 points.

4.2 A non Helly-type theorem

Let $\mathcal{C}$ be the class of 4-sided (but not 4-oriented) convex polygons in the plane. We prove below that $h(\mathcal{C}, \prod^2) = \infty$. Consider any triangle Δ, and place $k = 2m+1$ points $p_1, \ldots, p_k$ on the boundary of Δ; 3 points at the vertices of Δ and $\frac{k-3}{3}$ at each of its edges. Define the 4-sided convex polygon Q_i as the convex hull of $p_i, \ldots, p_{i+m-1}$, for $i = 1, \ldots, k$ (addition of subscripts is done modulo k). Put $\mathcal{S} = \{Q_1, \ldots, Q_k\}$ (see Figure 2). Clearly, $Q_i \cap Q_j \neq \emptyset$ iff $\{p_i, \ldots, p_{i+m-1}\} \cap \{p_j, ..., p_{j+m-1}\} \neq \emptyset$. Therefore p_l belongs to exactly m members of $\mathcal{S}$, $l = 1, \ldots, k$, and thus $\mathcal{S}$ is not 2-pierceable. On the other hand, $\mathcal{S} \backslash \{Q_j\}$ is 2-pierceable; take the points p_{j+m} and p_{j+2m}. Since $\lim_{m \to +\infty} k(m) = \infty$, the assertion $h(\mathcal{S}, \prod^2) = \infty$ follows.

Theorem 4 *Let $\mathcal{C}$ be the class of 4-sided convex polygons, then $h(\mathcal{C}, \prod^2) = \infty$.*

4.3 Looking at the boundary

In this section, we show that in some cases it is possible to obtain better bounds than those implied by the method of the previous section. In these cases we also apply the general method, but, in the last step, when considering the pairs of regions that were obtained, we are able to decrease the size of the search space. In other words, instead of having to consider all the cells in the clipped arrangement $\mathcal{A}(E)$, we observe that it is sufficient to consider only the arrangement $\mathcal{A}(\partial E)$, where ∂E is the boundary of E, since all the maximal cells of $\mathcal{A}(E)$ are adjacent to ∂E.

4.3.1 ≤ 5-oriented polygons with 2 points

Let us consider any class $\mathcal{C}_{\mathcal{H}}$, $|\mathcal{H}| \leq 5$, in the plane. That is, a class of at most 5-oriented convex polygons. In this case, let A be the region of a pair (A, B) for which at most $\lfloor \frac{5}{2} \rfloor = 2$ orientations appear in its clipped arrangement $\mathcal{A}(A)$. It is easy to see that all the maximal cells of $\mathcal{A}(A)$ are adjacent to the boundary ∂A of A. We therefore compute in $O(n \log n)$ time the intersection of ∂A with the n objects. Then, we traverse the $O(n)$ intervals along ∂A (each object may give rise to at most $O(c)$ intervals). The last step is done by maintaining dynamically the intersection property (1-pierceability) of the complementary set of objects. We thus obtain an $O(n \log n)$ time algorithm.

Theorem 5 *Let $\mathcal{S}$ be a set of n 5-oriented convex polygons. Then it is possible to find a piercing pair for $\mathcal{S}$ (if such exists) in $O(n \log n)$ time.*

4.3.2 Homothetic triangles with 3 points

We can use the linear-time algorithm of section 4.1 in order to obtain an $O(n \log n)$ time algorithm for piercing a set $\mathcal{S}$ of n homothetic triangles in the plane with 3 points. Indeed, we first check in linear time whether $\mathcal{S}$ is $\leq$ 2-pierceable. If $\mathcal{S}$ is $\leq$ 2-pierceable then we are done, otherwise we consider the counterexample that was obtained which is of constant size. We make all the possible assumptions for piercing this set with three points. Each assumption gives rise to a set of three disjoint convex regions (homothetic triangles). We apply the 'tree process' method to each of these triplets, and end up with $O(1)$ triplets of convex disjoint regions. Consider one such triplet (A, B, C). One of its regions, say A, has at most 2 orientations appearing in its clipped arrangement $\mathcal{A}(A)$ (since, if all regions have 3 orientations, then we can remove at least one orientation by applying the 'tree process'). It is easy to see that A has the boundary property, i.e., all the maximal cells of $\mathcal{A}(A)$ are adjacent to the boundary ∂A of A. We may therefore restrict our attention to the $O(n)$ intervals along ∂A (each object intersecting A defines at most 2 intervals along ∂A). We traverse ∂A moving from one endpoint of an interval to the next. At the current endpoint p, we check whether the set of objects that are not pierced by p is 2-pierceable. We can maintain the 2-pierceability property of this complementary set under insertions and deletions of objects in logarithmic time, using a linear-size dynamic data structure that is reminiscent of the data structures constructed in [SW96] for the 4- and 5-piercing problems for rectangles in the plane. We describe this data structure in the proof of Lemma 7 below. We thus obtain an $O(n \log n)$ algorithm.

Theorem 6 *Let $\mathcal{S}$ be a set of n homothetic triangles. Then it is possible to find a piercing triplet for $\mathcal{S}$ (if such exists) in $O(n \log n)$ time.*

The previous theorem is based on the following lemma.

Lemma 7 *The 2-pierceability property of a set of n homothetic triangles can be maintained dynamically in $O(\log n)$-time per update.*

Proof. Let R be a set of n homothetic triangles in the plane. We need to show how to maintain the 2-pierceability property (i.e., whether R is 2-pierceable or not) when a homothetic triangle is either inserted to or deleted from R. For each orientation o, draw a line through the 'topmost' edge of orientation o. Let L_R denote the set of these three lines.

According to Sharir and Welzl [SW96], if R is 2-pierceable, then it can be pierced by two points lying on the lines of L_R and we may assume that one of the piercing points lies on a vertex of the arrangement of L_R.

Our data structure consists of three balanced binary trees, one per each orientation. Assume w.l.o.g. that the first orientation is the vertical orientation and that the triangles lie to the left of their vertical edge. Then the first tree is built as follows. Let $x_1 < x_2 < \cdots < x_n$ be the projections of the vertical edges onto the x-axis. Store these values in sorted order in the leaves of a balanced binary tree; x_1 in the leftmost leaf, and so on. A node u of the tree represents the corresponding subset R_u of R (i.e. R_u consists of the triangles in R for which the projection of their vertical edge is stored in a leaf of the subtree rooted at u). We associate with u the region $T_u = \cap R_u$ (which is either empty or a homothetic triangle). The tree can be constructed in a bottom-up manner in $O(n)$ time, after sorting the leaf values in $O(n \log n)$ time. The second and third trees are constructed in a similar way.

We need to determine in logarithmic time whether R is 2 pierceable or not. First notice that we can obtain L_R, the three lines forming the 'location domain' of R, in constant time, since they are the lines corresponding to the minimal (or maximal) values stored in the three trees. Now let v be a vertex of the arrangement of L_R. We need to determine whether the subset R_v of R consisting of the triangles in R that are not pierced by v has a non-empty intersection. We can obtain R_v (implicitly) as the union of a logarithmic number of represented subsets (see below), and since we have the intersection of each of these subsets, we can compute $\cap R_v$ in logarithmic time.

How do we obtain R_v as a collection of represented subsets ? Notice that v lies in a triangle Δ iff it lies in the three half-planes defined by the edges of Δ. Let R_v^i, $i = 1, 2, 3$, denote the subset of R consisting of the triangles whose half-plane of the i'th orientation doesn't contain v. Clearly $R_v^1 \cup R_v^2 \cup R_v^3 = R_v$. We will perform a query in the i'th tree to obtain the subset R_v^i as the union of a logarithmic number of represented subsets. Consider for example the first tree, and let v_x be the projection of v onto the x-axis. We search with v_x in the first tree (beginning at the root) for the smallest value x that is stored in a leaf of the tree for which $v_x \leq x$. This search defines a root-leaf path in the tree. Consider the nodes of the tree that do not belong to this path and are the left children of nodes on this path. The number of these nodes is logarithmic and the union of their corresponding represented sets is exactly R_v^1.

Using standard techniques we can update this data-structure when a triangle is either inserted to or deleted from R in logarithmic time. Thus the 2-pierceability property can be maintained in logarithmic time. □

Remark: We cannot obtain a similar theorem for the case of at most 5-oriented convex polygons considered in the previous subsection, since we do not have a Helly-type theorem for the 2-piercing problem in this case.

5 Conclusion

The main result of this paper is a general method for solving the 2-piercing problem for a set $\mathcal{S}$ of d-dimensional c-oriented polytopes. This method is much more efficient than the naïve method whenever $\lfloor \frac{c}{2} \rfloor < d$ (e.g., when $c = d + 1$ and $\mathcal{S}$ is a set of simplices). It yields a linear-time algorithm and a new Helly-type theorem for the 2-piercing problem for homothetic triangles in the plane. This Helly-type theorem allows us to apply the method also to the 3-piercing problem for homothetic triangles and obtain an $O(n \log n)$ algorithm, since, in general, the method is based on the existence of a Helly-type theorem for the corresponding $(k-1)$-piercing problem.

Our method applied to the 2-piercing problem for d-dimensional axis-parallel boxes yields an alternative linear-time algorithm to the one implied in a paper of Danzer and Grünbaum [DG82] and also implies the corresponding Helly-type theorem of [DG82].

References

[Ame94] N. Amenta. Helly-type theorems and generalized linear programming. *Discrete Comput. Geom.*, 12:241–261, 1994.

[BGLR93] M. Bellare, S. Goldwasser, C. Lund, and A. Russel. Efficient probabilistically checkable proofs and applications to approximation. In *Proc. 25th ACM Symp. on Theory of Computing*, pages 294–304, 1993.

[Chv79] V. Chvátal. A greedy heuristic for the set-covering problem. *Math. Oper. Res.*, 4:233–235, 1979.

[CM93] B. Chazelle and J. Matoušek. On linear-time deterministic algorithms for optimization problems in fixed dimension. In *Proc. 4th ACM-SIAM Symp. Discrete Algorithms*, pages 281–290, 1993.

[DG82] L. Danzer and B. Grünbaum. Intersection properties of boxes in R^d. *Combinatorica*, 2(3):237–246, 1982.

[DGK63] L. Danzer, B. Grünbaum, and V. Klee. Helly's theorem and its relatives. In *Proc. Symp. Pure Math.*, pages 101–180, 1963.

[FPT81] R. J. Fowler, M. S. Paterson, and S. L. Tanimoto. Optimal packing and covering in the plane are NP-complete. *Inform. Process. Lett.*, 12(3):133–137, 1981.

[GW93] P. M. Gruber and J. M. Wills, editors. *Handbook of Convex Geometry*, volume A. North-Holland, Amsterdam, the Netherlands, 1993.

[HD60] H. Hadwiger and H. Debrunner. *Kombinatorische Geometrie in der Ebene*. Monographies de l'Enseignement Mathématique, N2, Geneva, 1960.

[Hoc82] D. S. Hochbaum. Approximation algorithms of the set covering and vertex cover problems. *SIAM J. Computing*, 11(3):555–556, 1982.

[Kar72] R. Karp. *Reducibility among combinatorial problems, complexity of computer computations*. E. Miller and J.W. Thatcher, New York, 1972.

[KN96] M. Katchalski and D. Nashtir. On a conjecture of Danzer and Grünbaum. In *Proc. AMS*, to appear.

[Meg83] N. Megiddo. Applying parallel computation algorithms in the design of serial algorithms. *J. ACM*, 30:852–865, 1983.

[MSW92] J. Matoušek, M. Sharir, and E. Welzl. A subexponential bound for linear programming. In *Proc. 8th ACM Symp. on Computational Geometry*, pages 1–8, 1992.

[Sha96] M. Sharir. A near-linear algorithm for the planar 2-center problem. In *Proc. 12th ACM Symp. on Computational Geometry*, pages 106–112, 1996.

[SW96] M. Sharir and E. Welzl. Rectilinear and polygonal p-piercing and p-center problems. In *Proc. 12th ACM Symp. on Computational Geometry*, pages 122–132, 1996.