Discrete Helly type theorems

Frederik Brinck Jensen* Aadi Joshi† Saurabh Ray‡

Abstract

Let P be a set of points and $\mathcal{S}$ a family of regions in the plane. We consider the following type of problems. Let k be some fixed positive integer. If for every k points in P, there exists a region $S \in \mathcal{S}$ containing all of the k points, then what is the size of the smallest subset of the regions whose union covers P? We also consider the *dual* problems: if every k of the regions intersect at a common point in P, what is the smallest subset of points in P that together *hit* all regions in $\mathcal{S}$?

The families of regions we consider are halfspaces, convex pseudodisks, and axis-parallel rectangles in the plane. We prove tight results for a small value of k for some of these. For all of our questions, it can be shown that the answer depends only on k, i.e. it is a constant for any fixed k, by using a technique similar to that used in the proof of the Hadwiger-Debrunner (p, q) theorem due to Alon and Kleitman. However, even for small values of k, determining the right answer seems non-trivial.

1 Introduction

Let $\mathcal{C}$ be a finite collection of convex sets in $\mathbb{R}^d$. Helly's theorem [1] (see also Chapter 1 of Matousek's book [2]) states that if every $d+1$ of these sets intersect at a common point in $\mathbb{R}^d$ then all the convex sets in $\mathcal{C}$ intersect at a common point in $\mathbb{R}^d$.

It is natural to ask if a discrete version of Helly's theorem is true. Instead of requiring that every $d+1$ of convex sets intersect at some point in $\mathbb{R}^d$, suppose that we require that they intersect at some point in a discrete set of points P. Then, can we conclude that all the convex sets intersect at some point in P? Unfortunately, this statement is not true even if we require that every k of the convex sets intersect at a point in P for some large constant $k \geq d+1$ and we want to conclude that all convex sets in $\mathcal{C}$ can be hit by some large constant number of points. To see this, consider a set P of n points in convex position in $\mathbb{R}^d$ and let $\mathcal{C}$ be the set of the convex hulls of every subset of P of size greater than $n - \frac{n}{k}$. Then the total size (in terms of the number of points of P contained) of any subset of k sets in $\mathcal{C}$ is more than $(k-1)\,n$ and thus they must have a common point in P. On the other hand, no subset of P of size less than n/k can hit all the convex sets in $\mathcal{C}$.

We show that such a statement is true for some simple regions in the plane. As mentioned earlier, for most of the cases we consider, the technique used by Alon and Kleitman [3] (see [4, 5, 6] and the references therein for more recent work on the problem) can be used to show that the answer is a constant. The basic idea is as follows. Suppose that we are given a set of points P and a set of regions $\mathcal{S}$ such that for each subset of k points in P, there exists a region $S \in \mathcal{S}$ containing all points in that subset. Then, for the families of regions we consider, it can be shown that there exists some region $S \in \mathcal{S}$ that contains a constant fraction (depending on k) of the points in P. Then, using LP-duality, one can assign weights to the regions such that the total weight of regions containing any particular point is at least ϵW where ϵ is a constant (which depends on k) and W is the total weight of all the regions. Finally, using strong ϵ-nets whose size depends only on ϵ, we can obtain a set of points of constant size that hits all the regions. A similar approach works in the dual setting too. However, the constants obtained from such techniques are large mainly due to the constants in the bounds on the size of ϵ-nets. Kleitman et. al. [7] considered the question of hitting convex sets in the plane with the minimum number of points (using arbitrary points in the plane) where three out of every four intersect at a common point. The main goal was to find a better constant using a more direct approach. They proved a lower bound of 3 but the upper bound obtained even for such a (apparently) simple problem is significantly higher: 13. This has not yet been improved.

Questions similar to ours are considered in [8]. In their setting the family of regions either consists of a single region (from some family of regions) or regions obtained by translating and rotating a single region. Halman [9] also studied 'discrete Helly-type theorems' that are different from ours. An excellent survey on Helly type theorems is [10].

2 Halfspaces in $\mathbb{R}^2$

We start with a simple Helly-type result for halfspaces in the plane.

Theorem 1 *Let P be a finite set of n points and $\mathcal{H}$ a set of halfspaces in the plane. If every subset of 3 points*

*New York University Abu Dhabi, fj414@nyu.edu
†New York University Abu Dhabi, aj1566@nyu.edu
‡New York University Abu Dhabi, saurabh.ray@nyu.edu

in P belongs to some halfspace $H \in \mathcal{H}$ then there exists two halfspaces $H_1, H_2 \in \mathcal{H}$ whose union covers P.

Proof. Let $P_{CH} \subseteq P$ denote the subset of points in P that lie on the boundary of the convex hull $\mathrm{CH}(P)$ of P. Consider the halfspace $H_1 \in \mathcal{H}$ containing the largest number of points from P_{CH}. Note that since there is a halfspace in $\mathcal{H}$ covering any triple of points in P_{CH}, H_1 contains at least three points of P_{CH}.

If H_1 contains all points in P_{CH}, then H_1 contains all points in P and the theorem follows. Otherwise, there are two edges on the boundary of $\mathrm{CH}(P)$ such that H_1 contains exactly one endpoint of each of the edges. Let p and q be the endpoints of the edges that are contained in H_1. See figure 1. Note that p and q cannot be the same point since this would mean that H_1 contains only one point of P_{CH} and we argued earlier that H_1 contains at least three points of P_{CH}.

The line through p and q splits $\mathrm{CH}(P)$ into two regions, one of which is covered by H_1. Let A be the region covered by H_1 and let B be the other region. Let $r \in P_{CH}$ be a point not contained in H_1. By assumption, there exists a halfspace, H_2, that contains p, q, and r. Since H_2 contains p and q, H_2 covers either A or B. If H_2 covered A, then it would contain all points of P_{CH} in H_1 and the additional point $r \in P_{CH}$. This would contradict the maximality of H_1. Thus H_2 must cover B. Thus $H_1 \cup H_2$ covers $\mathrm{CH}(P)$ and the theorem follows. □

We now show that the '3' in Theorem 1 is tight i.e., it cannot be replaced by '2'. To this end, we construct a set of points P and a set of halfspaces $\mathcal{H}$ so that every pair of points in P is covered by a halfspace in $\mathcal{H}$ and yet no pair of halfspaces in $\mathcal{H}$ cover all points in P.

Figure 2 shows a disk D and three arcs with a very large radius of curvature so that any tangent to any of the arcs passes through the disc D and does not intersect any of the other arcs. We construct a point set P with n points by distributing $n/3$ points uniformly on each of the three arcs. We define the set of halfspaces $\mathcal{H}$ as follows. For each point p on some arc l_i, $i \in \{0, 1, 2\}$, let H_p be a halfspace not containing p and containing all other points on l_i so that the boundary of H_p is parallel and arbitrarily close to the tangent to l_i at p. Note that H_p contains all points on l_j where $j = (i+1) \mod 3$ and does not contain any of the points in l_k where $k = (i-1) \mod 3$. $\mathcal{H}$ is the set $\{H_p : p \in P\}$.

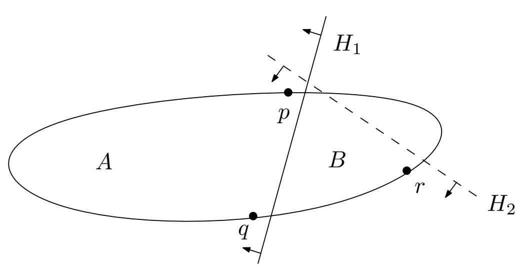

Figure 1: Let H_1 be the halfspace that contains the most points in P_{CH}. As depicted, the only positioning of any halfspace H_2 containing p, and q, and a point $r \notin H_1$ that avoids parts of $P \setminus H_1$ must contain more points on the convex hull of P than H_1, which is a contradiction. The dashed line, H_2, can therefore not exist and we are left with a halfspace that covers the remaining points not contained in H_1.

For any two points $p, q \in P$, we argue that there is a halfspace in $\mathcal{H}$ containing both p and q. There are two cases depending on whether p and q lie on the same arc or on different arcs. If p and q lie on the same arc l_i, then H_r where r is any point on l_k where $k = (i-1) \mod 3$ contains both p and q. If p and q lie on different arcs then without loss of generality, assume that p lies on l_i and q lies on l_j where $j = (i+1) \mod 3$. Then H_r, where r is any point on l_i other than p, contains both p and q.

Now, we argue that no two halfspaces in $\mathcal{H}$ cover all points in P. To see this consider any pair of halfspaces H_p and H_q. If both p and q lie on the same arc l_i then none of the points in l_k where $k = (i-1) \mod 3$ are covered by $H_p \cup H_q$. If p and q lie on different arcs, we assume without loss of generality that p lies on an arc l_i and q lies on l_j where $j = (i+1) \mod 3$. Then $H_p \cup H_q$ does not cover p.

Remark. The above example also improves a result from [11]. Lemma 17 of [11] shows that given a set P of n points in the plane, it is not always possible to hit all halfspaces in the plane containing more than ϵn points of P with just two points in P if $\epsilon < 3/5$. Our construction improves the bound $3/5$ to $2/3$ since all the halfspaces in our example contain $2n/3 - 1$ points and it is easily seen that no two points hit all halfspaces.

In our example, two halfspaces barely fail to cover all the points when every pair of points in P is covered by a halfspace. It seems intuitive that three halfspaces

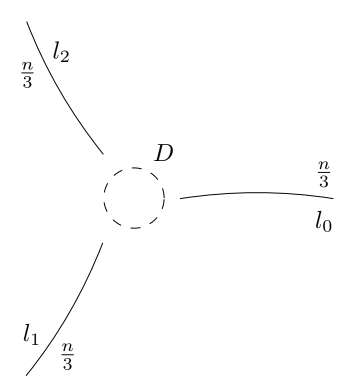

Figure 2: Counterexample that shows the tightness of Theorem 1.

should suffice to cover all points under this constraint. Indeed this is true and we show this later in Theorem 7.

3 Convex pseudodisks in the plane

A set of simply connected regions in the plane form a family of pseudodisks if the boundaries of any pair of regions either do not intersect or intersect at exactly two points. Furthermore, there are no tangential intersections, i.e., at each intersection the boundaries properly cross. Examples of families of pseudodisks include disks, squares, unit height rectangles, and homothets of a convex region.

Let $P \subset \mathbb{R}^2$ be a set of points and $\mathcal{D}$ a family of convex pseudodisks in the plane. We will show that if every three pseudodisks in $\mathcal{D}$ intersect at a point in P, then there exist two points in P that together hit all the pseudodisks in $\mathcal{D}$.

We first need a definition and some lemmas.

Definition 1 *For any disk $D \in \mathcal{D}$ we define its core* $\text{core}(D)$ *with respect to P as the convex hull of $D \cap P$. See Figure 3.*

Even though the core of a pseudodisk D is defined with respect to the point set P, we will skip the reference to P when it is clear from the context.

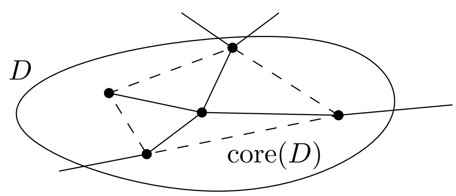

Figure 3: The dashed line outlines $\text{core}(D)$.

Note that any pseudodisk $D \in \mathcal{D}$ contains its own core since D is convex. Let $\mathcal{C}$ be the set of cores of the disks $D \in \mathcal{D}$.

Lemma 2 *The intersection of all cores in $\mathcal{C}$ is non-empty.*

Proof. Since every triple of convex pseudodisks in $\mathcal{D}$ intersect at some point $p \in P$ and such a point is contained in the cores of all three of them, all triples of cores in $\mathcal{C}$ have a non-empty intersection. Thus, by Helly's theorem all cores in $\mathcal{C}$ intersect at a point in $\mathbb{R}^2$ (which may not be a point in P). □

The next lemma follows from Lemma 5 in [12] by using an empty set as the set of *compulsory edges*.

Lemma 3 *There exists a straight-edge plane triangulation on P denoted $\mathcal{T}$ such that the points and edges inside any pseudodisk in $\mathcal{D}$ form a connected subgraph of $\mathcal{T}$.*

Lemma 5 in [12] is for arbitrary non-convex k-admissible regions (which includes pseudodisks) and therefore allows the edges to be curved. However, for convex pseudodisks, it does yield a straight-edge drawing. In fact, it shows that any maximal subset of the $\binom{n}{2}$ edges defined by P which are pairwise non-crossing and which do not cut across any of the cores of the regions form such a triangulation. Note that if $\mathcal{D}$ is a set of circular disks in the plane, then the Delaunay triangulation of P provides the triangulation claimed in the lemma above.

Lemma 4 *If the core of some $D \in \mathcal{D}$ intersects an edge $e \in \mathcal{T}$, then D must contain at least one of the endpoints of e.*

Proof. Since the edges of $\mathcal{T}$ are straight line segments, the points and edges inside any core in $\mathcal{C}$ also form a connected sub-graph of $\mathcal{T}$. If the core of a disk $D \in \mathcal{D}$ intersects an edge $e \in \mathcal{T}$ but does not contain either endpoint of e, then we obtain a contradiction since the edges lying in $\text{core}(D)$ cannot form a connected subgraph of $\mathcal{T}$. Thus $\text{core}(D)$ must contain an endpoint of e. □

Lemma 5 *There exist two points p, $q \in P$ that hit all disks in $\mathcal{D}$.*

Proof. Let $x \in \mathbb{R}^2$ be a point in the common intersection $\cap_{D \in \mathcal{D}} \text{core}(D)$ of all cores. By lemma 2, such a point x exists. Let T be the triangle in the triangulation $\mathcal{T}$ containing x. Since all the cores contain x, all cores intersect the edges of T and thus, by Lemma 4 contain at least one of the three corners of T. In other words, the three corners of T hit all pseudodisks.

We now show that one of the corners is redundant and can be dropped. Assume, for contradiction that all three corners of T are necessary i.e. for each corner there exists a pseudodisk that is hit only by that corner among the three corners. These three pseudodisks intersect inside T at x and, by definition, at some point in $p \in P$ which must lie outside T. Since the intersection of the three convex pseudodisks is convex, all three disks contain the segment joining x and p. Therefore, they all intersect some edge $e \in T$. However, by Lemma 4 this means that all three disks are hit by the two endpoints of e contradicting the assumption that all three corners are necessary for hitting the three disks. □

Thus we have proved the following theorem.

Theorem 6 *Given a set of points P and set of convex pseudodisks $\mathcal{D}$ in the plane s.t. every triple of pseudodisks in $\mathcal{D}$ intersects at a point in P, there exists a set of two points $\{p, q\} \subseteq P$ which intersects each $D \in \mathcal{D}$.*

Remark. The above theorem implies that given a set of n points and a set $\mathcal{D}$ of convex pseudodisks in the plane, all pseudodisks in $\mathcal{D}$ containing more than $2n/3$ points can be hit by two points. This is because any three sets containing more than $2n/3$ elements of a set of size n must have a common element. This generalizes Theorem 18 of [11] which proves this for disks.

Note that Theorem 6 is true for halfspaces in the plane too since we can think of halfspaces as disks of infinite radius. This can be used to prove the following theorem.

Theorem 7 *Let $\mathcal{H}$ be a set of halfspaces and let P be a set of points in the plane such that for each pair of points in P, there is a halfspace in $\mathcal{H}$ covering both points. Then, three of the halfspaces in $\mathcal{H}$ cover all points in P.*

Proof. Let $\mathcal{H}'$ be a set of halfspaces consisting of the complements of the halfspaces in $\mathcal{H}$. Assume for contradiction that no three halfspaces in $\mathcal{H}$ cover all points in P. This means that every triple of halfspaces in $\mathcal{H}'$ intersects at a point in P. Since the halfspaces in $\mathcal{H}'$ can be thought of as disks of infinite radius, we can apply Theorem 6 to $\mathcal{H}'$ and P to conclude that two points $p, q \in P$ hit all halfspaces in $\mathcal{H}'$. This implies that each halfspace in $\mathcal{H}$ avoids at least one of the points in $\{p, q\}$. This contradicts the assumption that for every pair of points in P, there is some halfspace in $\mathcal{H}$ that contains both. □

4 Axis-parallel rectangles in $\mathbb{R}^2$

In this section, we prove a couple of simple discrete Helly-type theorems for axis-parallel rectangles in the plane.

Theorem 8 *Let P be a set of points and let $\mathcal{R}$ be a set of axis-parallel rectangles in the plane s.t. for every triple of points $p, q, r \in P$, there exists a rectangle $R \in \mathcal{R}$ containing all three. Then, 2 of the rectangles in $\mathcal{R}$ together cover all points in P. Furthermore, the number 2 here is tight.*

Proof. Let l, r, t and b be the leftmost, rightmost, topmost and bottom-most points respectively in P. Note that these need not be distinct. Then, the rectangle containing l, r and t and the rectangle containing l, r and b together cover the convex hull of the points in P and therefore all points in P. To see that the number 2 is tight, consider the four points $(1, 0), (-1, 0), (0, 1)$ and $(0, -1)$. For every three of these points, we can add to our family of rectangles, a rectangle containing these three but not the fourth. □

Theorem 9 *Let P be a set of points and let $\mathcal{R}$ be a set of axis-parallel rectangles in the plane s.t. for every pair of points $p, q \in P$, there exists a rectangle $R \in \mathcal{R}$ containing both p and q. Then, there are 5 rectangles in $\mathcal{R}$ whose union covers all points in P. Furthermore, the constant 5 is tight.*

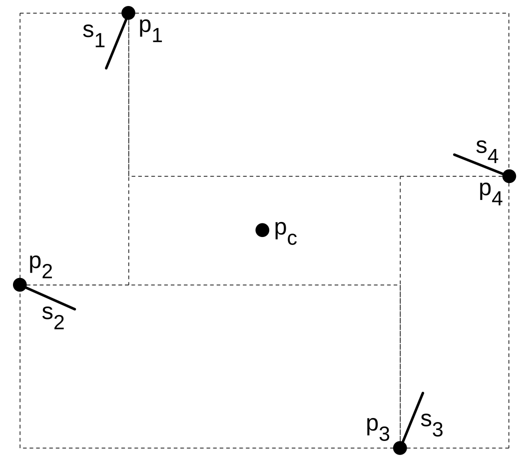

Figure 4: Four axis-parallel rectangles do not cover all points

Proof. Let l, r, t and b be the leftmost, rightmost, topmost and bottom-most points respectively in P. For any pair $x, y \in \{l, r, t, b\}$ of points, denote by $R_{x,y}$ any of the rectangles in $\mathbb{R}$ containing both x and y. We claim the the following five rectangles cover all the points in P: $R_{l,r}, R_{l,t}, R_{t,r}, R_{l,b}$ and $R_{b,r}$. To see this note that if a point $p \in P$ is not covered by $R_{l,r}$ then it lies either above or below $R_{l,r}$. If it lies above, it lies in the union of $R_{l,t}$ and $R_{t,r}$. Similarly, if it lies below, it lies in the union of $R_{l,b}$ and $R_{b,r}$.

We now give an example of a set of points and a set of axis-parallel rectangles satisfying the conditions of the theorem but in which no four of the rectangles cover all points. The point set is shown in Figure 4. Each line segment s_i in the figure extends a short distance from p_i in the direction of $p_{(i+1) mod 4}$, so that no two of them are intersected by the same vertical or horizontal line. The point p_c has an x-coordinate between those of p_1 and p_3, and a y-coordinate between those of p_2 and p_3.

Our point set includes the points p_1, p_2, p_3, p_4 and p_c. In addition, on each segment s_i, we place four more (distinct) points. Each segment s_i thus has five points including p_i.

The family of rectangles is the following. Let $i \in \{1, 2, 3, 4\}$ and let $j = (i + 1) \mod 4$. For each point $q \in s_j$, we add the axis-parallel rectangle with corners p_i and q to our set. We also add the rectangle with corners x and y for every pair x, y where x lies in s_1 and y lies in s_3. Similarly, we add the rectangle with corners x and y for every pair x, y where x lies in s_2 and y lies in s_4. Observe that for each pair of points in our point set, we have added a rectangle containing both, thus satisfying the conditions of the theorem.

We now show that no four rectangles in our family cover all points in our point set. To see this, assume to the contrary that it is possible to cover all points with four of the rectangles. Consider the segment s_i, $i \in \{1, 2, 3, 4\}$. Since s_i has five points, at least one of the four rectangles that together cover all points must contain two of the points in s_i. This means that we are *forced* to include one of the rectangles which has one corner at p_i and another corner in s_j where $j = (i + 1) \mod 4$. Since none of these four *forced* rectangles cover p_c, we derive a contradiction. The theorem follows. □

5 Open Problems

The following are some of the simple cases for which we are currently unable to obtain a precise answer.

Open problem 1. Let $\mathcal{D}$ be a set disks and let P be a set of points in $\mathbb{R}^2$ s.t. for every triple of points in P, there is a disk in $\mathcal{D}$ covering the three points. How many disks of $\mathcal{D}$ suffice to cover all points (in the worst case)? We believe that two disks in $\mathcal{D}$ suffice. Note that 2 is a lower bound even for halfplanes, as shown in this paper.

Open Problem 2. Let $\mathcal{R}$ be a set of axis-parallel rectangles and let P be a set of points in the plane s.t. every triple of rectangles in $\mathcal{R}$ intersects at a point in P. How many points from P suffice to hit all rectangles in $\mathcal{R}$. There is a simple example showing that the answer is at least 2 [9]. We are currently unable to prove that two points suffice, nor do we have an example in which three points are necessary.

References

[1] Ed. Helly. Über mengen konvexer körper mit gemeinschaftlichen punkte. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 32:175–176, 1923.

[2] Jiri Matousek. *Lectures on Discrete Geometry*. Springer-Verlag, Berlin, Heidelberg, 2002.

[3] N. Alon and D. Kleitman. Piercing convex sets and the Hadwiger-Debrunner (p,q)-problem. *Adv. Math.*, 96(1):103–112, 1992.

[4] Chaya Keller and Shakhar Smorodinsky. A new lower bound on hadwiger-debrunner numbers in the plane. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 1155–1169. SIAM, 2020.

[5] Chaya Keller and Shakhar Smorodinsky. On piercing numbers of families satisfying the $(p, q)_r$ property. *Comput. Geom.*, 72:11–18, 2018.

[6] Chaya Keller and Shakhar Smorodinsky. From a (p, 2)-theorem to a tight (p, q)-theorem. In Bettina Speckmann and Csaba D. Tóth, editors, *34th International Symposium on Computational Geometry, SoCG 2018, June 11-14, 2018, Budapest, Hungary*, volume 99 of *LIPIcs*, pages 51:1–51:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

[7] Daniel J. Kleitman, András Gyárfás, and Géza Tóth. Convex sets in the plane with three of every four meeting. *Combinatorica*, 21(2):221–232, 2001.

[8] L. M. Kelly. Covering problems. *National Mathematics Magazine*, 19(3):123–130, 1944.

[9] Nir Halman. Discrete and lexicographic helly-type theorems. *Discrete and Computational Geometry*, 39(4):690–719, Sep 2007.

[10] Jürgen Eckhoff. Chapter 2.1 - helly, radon, and carathéodory type theorems. In P.M. Gruber and J.M. Wills, editors, *Handbook of Convex Geometry*, pages 389 – 448. North-Holland, Amsterdam, 1993.

[11] Pradeesha Ashok, Umair Azmi, and Sathish Govindarajan. Small strong epsilon nets. *Comput. Geom.*, 47(9):899–909, 2014.

[12] Evangelia Pyrga and Saurabh Ray. New existence proofs epsilon-nets. In Monique Teillaud, editor, *Proceedings of the 24th ACM Symposium on Computational Geometry, College Park, MD, USA, June 9-11, 2008*, pages 199–207. ACM, 2008.