Efficiently Stabbing Convex Polygons and Variants of the Hadwiger-Debrunner (p, q)-Theorem

Justin Dallant
Department of Computer Science, Université libre de Bruxelles, Belgium

Patrick Schnider
Department of Mathematical Sciences, University of Copenhagen, Denmark

Abstract

Hadwiger and Debrunner showed that for families of convex sets in $\mathbb{R}^d$ with the property that among any p of them some q have a common point, the whole family can be stabbed with $p - q + 1$ points if $p \geq q \geq d + 1$ and $(d - 1)p < d(q - 1)$. This generalizes a classical result by Helly. We show how such a stabbing set can be computed for a family of convex polygons in the plane with a total of n vertices in $\mathcal{O}((p - q + 1)n^{4/3} \log^8 n(\log\log n)^{1/3} + np^2)$ expected time. For polyhedra in $\mathbb{R}^3$, we get an algorithm running in $\mathcal{O}((p - q + 1)n^{5/2} \log^{10} n(\log\log n)^{1/6} + np^3)$ expected time. We also investigate other conditions on convex polygons for which our algorithm can find a fixed number of points stabbing them. Finally, we show that analogous results of the Hadwiger and Debrunner (p, q)-theorem hold in other settings, such as convex sets in $\mathbb{R}^d \times \mathbb{Z}^k$ or abstract convex geometries.

Keywords and phrases computational geometry, Helly, (p,q)-Theorem, convex sets, stabbing, range searching

Digital Object Identifier 10.57717/cgt.v1i1.18

Related Version A preprint is available at `https://arxiv.org/abs/2002.06947`. An extended abstract [14] was presented at CCCG 2021, Halifax, Nova Scotia, Canada. A preliminary version was presented at EuroCG 2020, University of Würzburg.

Funding *Justin Dallant*: This work was partially supported by the French Community of Belgium via the funding of a FRIA grant.
Patrick Schnider: Has received funding from the European Research Council under the European Union's Seventh Framework Programme ERC Grant agreement ERC StG 716424 - CASe. Part of this work was done when the author was employed at ETH Zürich.

1 Introduction

A classical result in convex geometry by Helly [25] states that if a family of convex sets in $\mathbb{R}^d$ is such that any $d + 1$ sets have a common intersection, then all sets do. In 1957, Hadwiger and Debrunner [22] considered a generalization of this setting. Let $\mathcal{F}$ be a family of sets in $\mathbb{R}^d$ and let $p \geq q \geq d + 1$ be integers. We say that $\mathcal{F}$ has the (p, q)-property if $|\mathcal{F}| \geq p$ and for every choice of p sets in $\mathcal{F}$ there exist q among them which have a common intersection. We further say that a set of points S stabs $\mathcal{F}$ if every set in $\mathcal{F}$ contains at least one point from S. Then the following holds.

▶ **Theorem 1.** *[Hadwiger and Debrunner [22]] Let $d \geq 1$ be an integer. Let p and q be integers such that $p \geq q \geq d + 1$ and $(d - 1)p < d(q - 1)$, and let $\mathcal{F}$ be a finite family of convex sets in $\mathbb{R}^d$. Suppose that $\mathcal{F}$ has the (p, q)-property. Then there exist $p - q + 1$ points in $\mathbb{R}^d$ stabbing $\mathcal{F}$.*

Note that the bound on the number of points needed is tight. That is, for every $p \geq q \geq d + 1$ there exist families of convex sets with the (p, q)-property where at least $p - q + 1$ points are needed to stab the whole family. This is easily seen by considering any family of $p - q + 1$ disjoint convex sets where one of them is taken with multiplicity q. It is also known that whenever $q \leq d$, there exist families of convex sets with the (p, q)-property where arbitrary large number of points are needed. This can be seen by taking n hyperplanes in general position in $\mathbb{R}^d$ (meaning that no two hyperplanes are parallel and no $d+1$ hyperplanes intersect at the same point). Then any d hyperplanes intersect at some point (in other words, they have the (d, d)-property) and any single point stabs at most d hyperplanes. Thus, at least $\lfloor n/d \rfloor$ points are necessary to stab all hyperplanes.

Many related results have since been established. Among the most famous is one from Alon and Kleitman [5] who in 1992 proved that for any $p \geq q \geq d+1$, there exists a finite upper bound on the maximum number of points needed to stab a family of convex sets with the (p, q)-property. However, all the known upper bounds are probably far from being tight in the general case. As an example, for $(p, q, d) = (4, 3, 2)$, their proof yields an upper bound of 4032 (while the best known lower bound is 3). Since then, this number has been proven to lie between 3 and 13 (inclusive) [29]. A further improvement to the upper bound from 13 to 9 by McGinnis has recently appeared as a preprint [33]. Still, the only values of $p \geq q \geq d+1$ for which exact values are known are those corresponding to Theorem 1. There is a lot of work in this more general setting, both improving the bounds (e.g. [28]) as well as adapting to generalizations of convex sets (e.g. [27, 36]), and it is an interesting open problem to study algorithmic questions connected to these results.

Special cases where some further restrictions are imposed on the considered sets have also led to interesting results. One much studied example is obtained by considering only axis-aligned boxes in $\mathbb{R}^d$. In this case, we can already start by strengthening the result given by Helly's theorem, as for a family of axis-aligned boxes in $\mathbb{R}^d$, if all pairs intersect then the whole family intersect. As is expected, this additional structure leads to stronger (p, q) results. One early result by Hadwiger and Debrunner [23] is the following (notice the weaker conditions on p and q and the independence on d).

▶ **Theorem 2** ([23]). *Let $d \geq 1$ be an integer. Let p and q be integers such that $2q - 2 \geq p \geq q \geq 2$ and let $\mathcal{F}$ be a finite family of axis-aligned boxes in $\mathbb{R}^d$. Suppose that $\mathcal{F}$ has the (p, q)-property. Then there exist $p - q + 1$ points in $\mathbb{R}^d$ stabbing $\mathcal{F}$.*

Another example is when all sets are translations either with or without scaling of some convex set K. Here, strong results exist only for some very simple cases such as K being a d-dimensional cube or ball. For example the maximum number of points needed to stab families of discs in the plane with the $(p, 2)$-property lies between $4p-4$ and $7p-10$ inclusive [41]. These bounds are tight for $p = 2$, that is, in the case of pairwise intersecting discs (which was previously shown by Stachó [38] and Danzer [15]).

From an algorithmic point of view, little work seems to have been done towards computing these stabbing points. One instance which has recently received some attention is the aforementioned case of pairwise intersecting discs in the plane. Har-Peled et al. [24] showed how such a family can be stabbed with 5 points in linear time (which is one more point than the theoretical optimum). Shortly after another paper, yet to be formally published, claimed to find a linear time algorithm for stabbing such a family with only 4 points [10]. However, the computation of small stabbing sets for families of general convex polyhedra with the (p, q)-property seems to not have been studied and will constitute one part of this paper, in the setting of Theorem 1.

For a great overview of the studied questions and known results around (p, q) problems, we refer the interested reader to the 2003 survey by Eckhoff [18].

Before continuing, we would also like to mention that Helly's theorem has been generalized to many other settings, as this will come into play in the second part of this paper. In general, we say that a set system has *Helly number h* if the following holds: if any h sets in the set system have a common intersection, then the whole set system does. Helly numbers have been shown to exist for many set systems, such as convex sets in $\mathbb{R}^d \times \mathbb{Z}^k$ [6, 26] or abstract convex geometries (see the book by Edelman and Jamison [19] or Chapter III in the book by Korte et al. [30]), which include subtrees of trees and ideals of posets. In many of these cases, the proofs can be adapted to show a constant stabbing number analogous to the result by Alon and Kleitman. In this work, we will show that under some weak conditions, the existence of a Helly number implies a tight Hadwiger-Debrunner type result.

Our contributions

In Section 2, we give an algorithm for computing $p - q + 1$ stabbing points for families of convex polytopes in $\mathbb{R}^d$ which obey the conditions of the Hadwiger-Debrunner (p, q)-Theorem. This algorithm runs in time $\mathcal{O}((p - q + 1)n^d + np^d)$. We then show in Section 3, how to substantially improve the (expected) runtimes in dimensions 2 and 3.

In Section 4 we explore two other settings for convex sets related to the (p, q)-Theorem where our algorithms also apply. The first generalizes a formulation (due to Breen [9]) of Helly's theorem in terms of holes in the unions of certain sets. The second is a quantitative variant of the (p, q)-Theorem due to Montejano and Soberón [35].

In Sections 5 and 6 we explore the same themes in settings which abstract and generalize those of convex sets in Euclidean space. We derive (p, q)-theorems in these setting as well as algorithmic results in these settings. We then give concrete examples in Section 7. Some of the results obtained in these last 3 section are folklore.

2 Stabbing convex polytopes in any fixed dimension

2.1 A proof of the Hadwiger-Debrunner theorem

We will first consider a proof of Theorem 1 which will naturally lead to an algorithm for finding stabbing points. In a book by Matoušek [32], the proof of this theorem is left as an exercise, yet the hint suggests that the intended solution is close to the proof below. The main differences with other proofs for this theorem are that it is more constructive and does not make use of a separating hyperplane, which will make it easier to adapt to other settings later on.

We will make use of a lemma which can also be found in the same book. We include the proof as we will later use the same ideas to prove a similar lemma. For a non-empty compact set S, let $lexmin(S)$ denote its lexicographical minimum point. Then we have the following.

▶ **Lemma 3** ([32, Lemma 8.1.2]). *Let $\mathcal{F}$ be a family of at least $d + 1$ compact convex sets in $\mathbb{R}^d$, such that $I := \bigcap \mathcal{F}$ is non-empty. Let $x := lexmin(I)$. Then, there exist a subfamily $\mathcal{H} \subset \mathcal{F}$ of size d such that $x = lexmin(\bigcap \mathcal{H})$.*

Proof. Let $\mathcal{F}$, I and x be as specified in the statement. Let S_x denote the set of all points lexicographically smaller than x. This set is convex and is disjoint from I. By Helly's theorem, there exists a subfamily of $d + 1$ members of $\mathcal{F} \cup \{S_x\}$ with an empty common intersection. These members have to include S_x, as all members of $\mathcal{F}$ have a non-empty common intersection. Let $\mathcal{H} \subset \mathcal{F}$ be the family consisting of the remaining d sets and let $x_{\mathcal{H}}$ be the lexicographical minimum point of $I' := \bigcap \mathcal{H}$ (which is compact and non-empty). $x_{\mathcal{H}}$ can not be lexicographically larger than x because $\mathcal{H} \subset \mathcal{F}$ and it can not be lexicographically smaller than x because $I' \cap S_x = \emptyset$. Thus, $x_{\mathcal{H}} = x$. ◀

Recall the theorem we wish to prove:

▶ **Theorem 1.** *[Hadwiger and Debrunner [22]] Let $d \geq 1$ be an integer. Let p and q be integers such that $p \geq q \geq d+1$ and $(d-1)p < d(q-1)$, and let $\mathcal{F}$ be a finite family of convex sets in $\mathbb{R}^d$. Suppose that $\mathcal{F}$ has the (p, q)-property. Then there exist $p-q+1$ points in $\mathbb{R}^d$ stabbing $\mathcal{F}$.*

Proof. We will prove the theorem for families of compact convex sets, as we will only deal with such families later. One can however reduce the original theorem to this one (see Appendix A), so this is done without loss of generality.

Call a pair of integers (p, q) *admissible* if $p \geq q \geq d+1$ and $(d-1)p < d(q-1)$. Let (p, q) be an admissible pair, and let $\mathcal{F}$ be a family of compact convex sets of $\mathbb{R}^d$ with the (p, q)-property.

We reason by induction on p, the base case being $p = q = d+1$ which is Helly's theorem.

If $p = q > d+1$, then $\mathcal{F}$ also has the $(d+1, d+1)$ property (as having the (p, q)-property implies having the $(p-1, q-1)$-property) and the result again follows from Helly's theorem.

So suppose that $p > q$ and that the result is true for any admissible pair (p', q') with $p' < p$.

If $(d-1)p = d(q-1) - k - 1$ for $k \geq 1$, then notice that $(p-k, q-k)$ is an admissible pair, as in that case $(d-1)(p-k) = d(q-k-1) - 1$ which together with $p > q$ also implies that $q - k \geq d+1$. Thus the result follows from the induction hypothesis.

It now remains to consider the case where $p > q$ and $(d-1)p = d(q-1) - 1$.

To do so, let us construct a point $x^*(\mathcal{F})$ as follows:

- For every non-empty subfamily $\mathcal{S} \subset \mathcal{F}$ of d convex sets with non-empty intersection, let $x_{\mathcal{S}}$ be the lexicographical minimum of $I_{\mathcal{S}} = \bigcap \mathcal{S}$.
- Let $x^*(\mathcal{F})$ be the lexicographical maximum point among all such $x_{\mathcal{S}}$'s.

Let $\mathcal{G}$ be one of the families defining $x^*(\mathcal{F})$, that is, $\mathcal{G} \subset \mathcal{F}$ is a subfamily of d sets which have $x^*(\mathcal{F})$ as the lexicographical minimum of their intersection.

To establish the theorem, it is enough to show that by choosing $x^*(\mathcal{F})$ as one of our stabbing points, we can stab all the remaining sets (i.e. those which do not contain $x^*(\mathcal{F})$) with $p - q$ points. Let $\mathcal{R} = \{C \in \mathcal{F} \mid x^*(\mathcal{F}) \notin C\}$ be the set of remaining sets.

Let us argue that for any $S \in \mathcal{R}$, $S \cap (\bigcap \mathcal{G})$ is empty. To do so, suppose it was not, and let y be the lexicographical minimum of that intersection. By Lemma 3, y is the lexicographical minimum of the intersection of d sets in $\mathcal{F}$. Moreover, by definition of $\mathcal{R}$ and $\mathcal{G}$, y is lexicographically larger than $x^*(\mathcal{F})$. This contradicts the definition of $x^*(\mathcal{F})$. Thus, $S \cap (\bigcap \mathcal{G})$ is empty.

Two cases arise:

1. ($|\mathcal{R}| \geq p - d$) We show that $\mathcal{R}$ has the $(p-d, q-d+1)$-property. Indeed, choose any $p - d$ members from $\mathcal{R}$ together with the d members from $\mathcal{G}$. We know from the (p, q)-property of $\mathcal{F}$ that there exists a subfamily $\mathcal{E} \subset \mathcal{R} \cup \mathcal{G}$ of size q whose members have a non-empty common intersection. $\mathcal{E}$ cannot contain all elements of $\mathcal{G}$, as $q > d = |\mathcal{G}|$ and the intersection of all members of $\mathcal{G}$ together with any member of $\mathcal{R}$ is empty. Thus, $\mathcal{E}$ contains at least $q - d + 1$ members of $\mathcal{R}$. This shows that $\mathcal{R}$ has the $(p-d, q-d+1)$-property. Notice that with the assumptions $p > q$ and $(d-1)p = d(q-1) - 1$ which we are working under, $(p-d, q-d+1)$ is admissible. Thus, by the induction hypothesis, $\mathcal{R}$ can be stabbed with $p - d - (q - d + 1) + 1 = p - q$ points.
2. ($|\mathcal{R}| < p - d$) In this case, choose $\mathcal{R}$ as a whole together with $\mathcal{G}$ and $p - d - |\mathcal{R}|$ other members of $\mathcal{F}$. By the same reasoning as in case 1., there exists a subset of $q - (d - 1 + p - d - |\mathcal{R}|) = |\mathcal{R}| + 1 + q - p$ members of $\mathcal{R}$ which intersect and can thus be stabbed by a single point. The remaining $|\mathcal{R}| - (|\mathcal{R}| + 1 + q - p) = p - q - 1$ sets can trivially be stabbed by $p - q - 1$ points.

Thus, $\mathcal{R}$ can be stabbed by $p-q$ points, which implies that $\mathcal{F}$ can be stabbed by $p-q+1$ points and by induction, concludes the proof. ◀

2.2 A first algorithm

This proof naturally leads to an algorithm. Let $d > 0$ be some fixed dimension and let $\mathcal{F}$ be a family of compact convex polytopes with (p, q)-property, described as intersections of a total of n halfspaces in general position. For simplicity we assume that the common intersection of any d of these polytopes is either empty or contains a unique point with minimum x-coordinate (which is then also the lexicographically minimum point in the intersection).

The algorithm works as follows:

1. Reduce p and q (as done in the proof of Theorem 1) to reach the case where $p = q = d+1$ or the case where $p > q$ and $(d-1)p = d(q-1) - 1$.
2. Construct a point $x^*(\mathcal{F})$ defined as in the proof. We choose it as one of our stabbing points. Now, remove from $\mathcal{F}$ all the sets that are stabbed by this point. If there are any remaining sets then either $|\mathcal{F}| \geq p - d$ and $\mathcal{F}$ satisfies the $(p-d, q-d+1)$-property, where $(p-d, q-d+1)$ is admissible, or $\mathcal{F}$ consists of $p-q+k$ sets, $k < q-d$, where some $k+1$ of them have a common intersection.
3. In the first case, we can continue inductively.
4. In the second case we can trivially stab the remaining sets using $p-q$ points.

The correctness of the algorithm follows immediately from the proof of Theorem 1. The only detail that needs some additional scrutiny is the correctness for the base case $p = q = d+1$. Notice that in this case all sets have a common intersection and Lemma 3 ensures that $x^*(\mathcal{F})$ stabs the whole family $\mathcal{F}$.

Regarding the runtime of Step 2, the most natural way to compute $x^*(\mathcal{F})$ gives the following.

▶ **Lemma 4.** *We can compute $x^*(\mathcal{F})$ in $\mathcal{O}(n|\mathcal{F}|^{d-1})$ time.*

Proof. For every polytope $P \in \mathcal{F}$, we let $n(P)$ denote the total number of halfspaces describing P. For every subfamily $\mathcal{G}$ of $\mathcal{F}$, we let $n(\mathcal{G}) := \sum_{P\in\mathcal{G}} n(P)$.

We can compute $x^*(\mathcal{F})$ by testing for intersection in every subfamily $\mathcal{G}$ of $\mathcal{F}$ of size d and computing the point with minimum x-coordinate of that intersection if it is non-empty. We then take the lexicographically maximum point among all those computed.

If we consider some fixed subfamily $\mathcal{G}$, this computation can be done in $\mathcal{O}(n(\mathcal{G}))$ time using linear programming in constant dimension. Thus, the computation for that subfamily will cost at most $c \cdot n(\mathcal{G})$ for some constant c which does not depend on $\mathcal{G}$. Charge this cost to the polytopes $P \in \mathcal{G}$ by attributing a cost of $c \cdot n(P)$ to a polytope P.

Now, consider the cost charged to some fixed polytope P for the whole computation. As P appears in no more than $|\mathcal{F}|^{d-1}$ subfamilies of size d, its total cost charge is upper bounded by $c \cdot n(P) \cdot |\mathcal{F}|^{d-2}$. Summing across all polytopes $P \in \mathcal{F}$, we get a total cost of $\mathcal{O}(n \cdot |\mathcal{F}|^{d-1})$). ◀

This quantity needs to be computed at most $p-q+1$ times with the family $\mathcal{F}$ decreasing in size each time.

For Step 4, we have the following.

▶ **Lemma 5.** *For any $k > 0$, we can find a point stabbing k polytopes of $\mathcal{F}$ in $\mathcal{O}(n|\mathcal{F}|^{d})$ time, if such a point exists.*

Proof. If $k \leq d$, then we can test every subfamily of size $k+1$ for common intersection and compute a point in the intersection for a total cost of $\mathcal{O}(n|\mathcal{F}|^k) \leq \mathcal{O}(n|\mathcal{F}|^d)$.

If $k > h-1$, then we know from Lemma 3 that the lexicographical minimum of the intersection of k convex polytopes is also the the lexicographical minimum of the intersection of some d polytopes in $\mathcal{F}$ (which in our case is also the point with minimum x-coordinate in the intersection). Thus, one can find a point stabbing at least k sets by computing the point with minimum x-coordinate for each subfamily of size d (in $\mathcal{O}(n|\mathcal{F}|^d)$ time) and counting the number of sets intersected for each of the $\mathcal{O}(|\mathcal{F}|^d)$ computed points (in $\mathcal{O}(n|\mathcal{F}|^d)$ time as well). ◀

Because we know that when reaching Step 4 we have $|\mathcal{F}| < p$, it follows that Step 4 can be done in $\mathcal{O}(np^d)$ time.

Thus, we get a total runtime of

$$\mathcal{O}((p-q+1)n^d + np^d).$$

If p (and thus q) is small compared to n, the bottleneck in the computation time is the first term, which scales as $\mathcal{O}(n^d)$ with respect to n. The natural question that now comes to mind is: can we do better than $\mathcal{O}(n^d)$? We will see in the following section that we can indeed do better at least in dimensions 2 and 3, although at the cost of considering expected rather than worst-case runtime.

▶ Remark 6. If we further restrict the problem to only consider convex polytopes described by at most a constant number of halfspaces each, then the second term in the runtime becomes $\mathcal{O}(p^{d+1})$. In the plane, this term can further be improved from $\mathcal{O}(p^3)$ to $\mathcal{O}(p^2 \log p)$ by adapting the Bentley-Ottmann sweep line algorithm [8] (see Appendix B for more details). On the other hand, one can easily reduce the problem of finding a point stabbing at least three lines among p lines to the problem of Step 4 in the above algorithm (for $k > 1$) in linear time if we allow for infinitesimally thin polygons. This problem is 3-SUM hard (see the seminal paper by Gejentaan and Overmars [21], where the concept of 3-SUM hardness was first introduced). There is a strong belief that such problems can not be solved in $O(p^{2-\epsilon})$ time, which means that Step 4 can probably not be solved in $O(p^{2-\epsilon})$ time either, even for constant-size polygons.

3 Faster algorithms for 2D and 3D polytopes

In what follows we deal with the cases $d = 2$ and $d = 3$. Note that for $d \leq 3$ we can get the vertex representation of our polytopes as well as the faces of all dimensions from the halfspace representation in $O(n \log n)$ time by computing the convex hulls of the dual point sets. Thus we will assume that we have access to the vertices and edges and faces of our polygons and polyhedra as the $O(n \log n)$ overhead will be dominated by the rest of our algorithms.

3.1 The planar case

In this whole section, the family $\mathcal{F}$ consists of compact convex polygons with a total of n (distinct) vertices in the plane and has the (p, q)-property, for some admissible pair (p, q). For the sake of simplicity, we will assume that the lines defining the polygon edges are in general position, non-vertical and that all points defined as the lexicographical minimum in the intersection of a pair of sets have different x-coordinates. Under these assumptions the lexicographical minimum in a polygon (or intersection of polygons) is simply the leftmost point.

We break down the computation of $x^*(\mathcal{F})$ into two parts. Consider two intersecting polygons P_1 and P_2. The point x which is the leftmost of $P_1 \cap P_2$ can be of one of two types. Either (case 1) x is the leftmost point of P_1 (resp. P_2) and is contained in the interior of P_2 (resp. P_1) or (case 2) x is the proper intersection of an upper-hull edge e_U of P_1 (resp. P_2) and a lower-hull edge e_L of P_2 (resp. P_1) with the following property: the outward facing normal vectors of e_U and e_L form a (counter-clockwise orientated) angle of less than 180 degrees. Reciprocally, an upper-hull and a lower-hull edge which intersect with this property define the leftmost point of an intersection of two polygons.

We define $x_1^*(\mathcal{F})$ to be the rightmost point among all pairs of intersecting polygons in $\mathcal{F}$ corresponding to the first case (or $x_1^*(\mathcal{F}) = (-\infty, \infty)$ if there is no such pair), and similarly for $x_2^*(\mathcal{F})$ and the second case. It is clear that $x^*(\mathcal{F})$ is the rightmost point of $\{x_1^*(\mathcal{F}), x_2^*(\mathcal{F})\}$.

We will use the following result, which can be obtained by an adaptation of a proof by Matoušek [31] with the halfspace partition tree construction from Chan [12] (see Appendix C).

▶ **Theorem 7.** *Let S be a set of n objects, k a constant, and $\phi_1, \phi_2, \dots, \phi_k$ mappings from S to $\mathbb{R}^d$. Let ϕ_S be the function which maps k-tuples of halfspaces $H_1, H_2, \dots, H_k$ of $\mathbb{R}^d$ to the set*

$$\phi_S(H_1, H_2, \dots, H_k) := \{s \in S \mid \phi_1(s) \in H_1, \phi_2(s) \in H_2 \dots, \phi_k(s) \in H_k\}.$$

Suppose we have computed the point sets $\phi_1(S), \dots, \phi_k(S)$ and let $n \le m \le n/\log^{\omega(1)} n$. Then we can preprocess the point sets in $\mathcal{O}(n \log^k n + m)$ time such that $|\phi_S(H_1, H_2, \dots, H_k)|$ can be computed in $\mathcal{O}((n/m^{1/d})(\log n)^{2(k+(k-d-1)/d)}(\log\log n)^{1/d})$ expected time for any k-tuple of halfspaces.

Note that we have made no big effort in minimizing the polylog factor in the query runtime. It is thus conceivable that a more careful use of the tools in Matoušek's and Chan's papers could make this factor smaller.

We can use this result to prove the following.

▶ **Lemma 8.** *We can compute $x_1^*(\mathcal{F})$ in $\mathcal{O}(n^{4/3} \log^4 n (\log\log n)^{1/3})$ expected time.*

Proof. To compute $x_1^*(\mathcal{F})$, we can test for each polygon if its leftmost point is contained in the interior of another, and keep the rightmost point among those which are. We triangulate all polygons, so that this reduces to testing, for each of the $\mathcal{O}(n)$ leftmost points, if it is in the interior of one of the $\mathcal{O}(n)$ triangles. In the dual plane, this can be expressed as the composition of three half-plane range queries. Using Theorem 7 with $d = 2$, $k = 3$, and $m = n^{4/3} \log^4 n (\log\log n)^{1/3}$, we can thus preprocess the $\mathcal{O}(n)$ triangles in $\mathcal{O}(m)$ time such that counting how many triangles contain a particular point can be done in $\mathcal{O}(n^{1/3} \log^4 n (\log\log n)^{1/3})$ expected time. By querying all points we get the result. ◀

It remains to see how to compute $x_2^*(\mathcal{F})$ in subquadratic time. For this we use a simple but remarkably powerful technique discovered by Chan [11], which reduces many optimization problems to the corresponding decision problem, with no blow-up in expected runtime.

▶ **Lemma 9.** *Let $\alpha < 1$ and r be fixed constants. Let $f : \mathcal{P} \to \mathcal{Q}$ be a function that maps inputs to values in a totally ordered set (where elements can be compared in constant time) with the following properties.*

1. *For any input $P \in \mathcal{P}$ of constant size, $f(P)$ can be computed in constant time.*
2. *For any input $P \in \mathcal{P}$ of size n, we can construct inputs $P_1, \dots, P_r \in \mathcal{P}$ each of size at most $\lceil \alpha n \rceil$ in time $T(n)$, such that $f(P) = \max\{f(P_1), \dots, f(P_r)\}$.*
3. *For any input $P \in \mathcal{P}$ of size n and any $t \in \mathcal{Q}$, we can decide $f(P) \leq t$ in time $T(n)$.*

Then for any input $P \in \mathcal{P}$, we can compute $f(P)$ in $\mathcal{O}(T(n))$ expected time, assuming that $T(n)/n^\epsilon$ is monotone increasing for some constant $\epsilon > 0$.

We can apply this technique to the computation of x_2^*. Here, each $P \in \mathcal{P}$ is a set of edges which are oriented depending on which side the polygon it bounds lies on, $\mathcal{Q}$ is the plane with lexicographical order, and $f(P)$ is x_2^* (we abuse notation slightly by using x_2^* both for sets of oriented edges and sets of polygons). We make the following observations.

- For any constant-size set $\mathcal{E}$ of oriented edges, $x_2^*(\mathcal{E})$ can be computed in constant time. This verifies Property 1.
- For any family $\mathcal{E}$ of n oriented edges, we can partition it into 3 disjoint subfamilies S_1, S_2, S_3 of size between $\lfloor n/3 \rfloor$ and $\lceil n/3 \rceil$ each. Then, let $\mathcal{E}_1 := S_2 \cup S_3$, $\mathcal{E}_2 := S_1 \cup S_3$ and $\mathcal{E}_3 := S_1 \cup S_2$. Every set $\mathcal{E}_i$ is of size $|\mathcal{E}_i| \leq \lceil 2n/3 \rceil$. Thus, $x_2^*(\mathcal{E})$ is the rightmost point among $\{x_2^*(\mathcal{E}_1), x_2^*(\mathcal{E}_2), x_2^*(\mathcal{E}_3)\}$. These families can be constructed in $\mathcal{O}(n)$ time. This verifies Property 2, assuming $T(n) \geq \Omega(n)$ (which it will be).

Thus, in order to apply Chan's framework, it remains to decide $x_2^*(\mathcal{E}) \leq_{lex} t$ quickly.

▶ **Lemma 10.** *For any point t in the plane and a set of n oriented edges $\mathcal{E}$, we can decide $x_2^*(\mathcal{E}) \leq_{lex} t$ in $\mathcal{O}(n^{4/3} \log^8 n (\log\log n)^{1/3})$ expected time.*

Proof. We can rephrase $x_2^*(\mathcal{E}) \leq_{lex} t$ as deciding whether there exist two oriented edges in $\mathcal{E}$ which intersect at an appropriate angle to the right of the vertical line ℓ passing through t. Thus we start by discarding all the (parts of) segments in $\mathcal{E}$ which lie to the left of ℓ. We then want to preprocess the $\mathcal{O}(n)$ segments corresponding to upper-hull edges (i.e. those with an outward facing normal pointing up) such that for any lower-hull edge e_L we can detect if there is an upper-hull edge which intersects it at an appropriate angle quickly.

Map each upper-hull edge e_U to its endpoints $a(e_U)$, $b(e_U)$ and to the point $p^*(e_U)$ dual to the line supporting it. Now for a lower-hull edge e_L, let $R(e_L)$ denote the region of the plane corresponding to all points whose dual lines intersect e_L at an appropriate angle. This region is a convex polygon with at most 4 edges. Thus it can be partitioned into two triangles $R_1(e_L)$ and $R_2(e_L)$. Call ℓ the line supporting e_L. Now, all upper-hull edges e_U intersecting e_L at an appropriate angle fall into exactly one of the following categories:

- $a(e_U)$ lies to the left of ℓ, $b(e_U)$ lies to the right of ℓ and $p^*(e_U) \in R_1(e_L)$,
- $a(e_U)$ lies to the left of ℓ, $b(e_U)$ lies to the right of ℓ and $p^*(e_U) \in R_2(e_L)$,
- $a(e_U)$ lies to the right of ℓ, $b(e_U)$ lies to the left of ℓ and $p^*(e_U) \in R_1(e_L)$,
- or $a(e_U)$ lies to the right of ℓ, $b(e_U)$ lies to the left of ℓ and $p^*(e_U) \in R_2(e_L)$.

The number of upper-hull edges corresponding to each category can be counted by a range query which is the composition of 5 half-plane queries on the 3 liftings defined above.

We can again use Theorem 7 as we did for x_1^*, this time with $k = 5$, to query all lower-hull edges in $\mathcal{O}(n^{4/3} \log^8 n (\log\log n)^{1/3})$ expected total time. ◀

We can thus use Lemma 9 to compute $x_2^*(\mathcal{F})$ in the same asymptotic expected time. Note that Hopcroft's problem reduces to computing $x_2^*(\mathcal{E})$ for a general set of oriented edges $\mathcal{E}$, and thus this runtime is likely close to optimal (see Erickson's paper [20] for a lower bound in a somewhat general model of computation).

Putting everything together we get the following.

▶ **Theorem 11.** *Let (p, q) be an admissible pair for $d = 2$ and let $\mathcal{F}$ be a family of compact convex polygons in the plane with a total of n vertices and the (p, q)-property. Then we can compute a set of at most $p - q + 1$ points stabbing $\mathcal{F}$ in $\mathcal{O}((p-q+1)n^{4/3}\log^8 n(\log\log n)^{1/3} + np^2)$ expected time.*

3.2 Constant-size polygons

If we restrict all polygons to have at most a constant number of vertices, then a simpler proof using a Theorem by Agarwal et al. [1, Theorem 2.7] yields a slightly faster algorithm. Indeed, this theorem states the following.

▶ **Theorem 12** ([1, Theorem 2.7]). *Let $\mathcal{F}$ be a family of compact convex polygons in the plane with a total of n vertices. Then, we can count the number of pairs of polygons in $\mathcal{F}$ which intersect in $\mathcal{O}(n^{4/3}\log^{2+\epsilon} n)$ time, for any constant $\epsilon > 0$.*

Using this, it is easy to prove the following.

▶ **Theorem 13.** *Given a vertical line ℓ and a family $\mathcal{F}$ of compact convex polygons in the plane with a total of n vertices, we can decide whether $x^*(\mathcal{F})$ lies to the right of ℓ in $O(n^{4/3}\log^{2+\epsilon} n)$ time, for any constant $\epsilon > 0$.*

Proof. We start by cutting all polygons along the vertical line ℓ and discarding the parts lying on the left of ℓ in linear time.

The point $x^*(\mathcal{F})$ lies to the right of ℓ if and only if there are two polygons which have a non-empty intersection but do not intersect on ℓ. This can be decided by counting the number of pairwise intersecting polygons in $\mathcal{O}(n^{4/3}\log^{2+\epsilon} n)$ time, counting the number of pairwise intersecting polygons on ℓ (this can be done in $\mathcal{O}(n\log(n))$ time, or one can use the same algorithm again), and then comparing these numbers. They differ if and only if some pair of polygons intersect exclusively to the right of ℓ.

This whole procedure leads to an algorithm with a $\mathcal{O}(n^{4/3}\log^{2+\epsilon} n)$ runtime. ◀

Together with the modification mentioned in Remark 6 and a straightforward application of Lemma 9, this yields the following.

▶ **Theorem 14.** *Let (p, q) be an admissible pair for $d = 2$ and let $\mathcal{F}$ be a family n of compact convex polygons in the plane with at most a constant number of vertices each and the (p, q)-property. Then we can compute a set of at most $p - q + 1$ points stabbing $\mathcal{F}$ in $\mathcal{O}((p-q+1)n^{4/3}\log^{2+\epsilon} n + p^2\log p)$ expected time, for any constant $\epsilon > 0$.*

▶ Remark 15. At first glance it might seem like the constant-size assumption plays no essential role here for the $n^{4/3}\log^{2+\epsilon} n$ term in the runtime, and that nothing stops us from using the same approach in the general case. The trouble in the general case however comes in enforcing point 3 in Lemma 9. When the polygons are not restricted in size, it might be impossible to create subproblems of appropriate sizes. This is not a problem in our proof as we are working with sets of oriented edges instead of sets of polygons.

We finish this part by proving a lower bound on a restricted case of the problem solved in Theorem 13 (thus, a lower bound on the general case also).

▶ **Theorem 16.** *Given a vertical line ℓ and a family $\mathcal{F}$ of n closed triangles in the plane which all intersect ℓ, detecting whether two triangles intersect exclusively to the right of ℓ requires $\Omega(n\log n)$ time for the worst case in the algebraic decision tree model.*

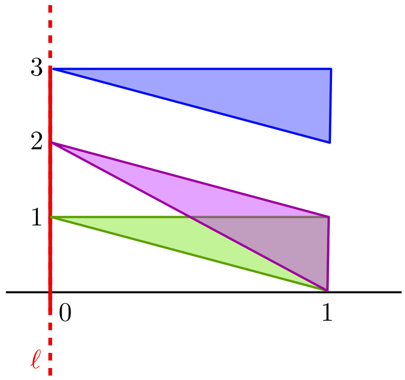

Figure 1 Right Intersection instance corresponding to the Element Uniqueness instance $[1, 0, 0]$

Proof. We prove the claim by a reduction from the Element Uniqueness Problem, which is known to have $\Theta(n \log n)$ time complexity in this model [7]. The Element Uniqueness Problem is the following: given an array of n integers, test if they are all distinct.

Let a be an array of length n representing an instance of the Element Distinctness Problem. Construct an instance of our problem in $O(n)$ time in the following way.

- Let ℓ be the y-axis.
- For every $k \in \{0, \ldots, n-1\}$, create a triangle with vertex coordinates $(0, n \cdot a[k] + k)$, $(1, 2 \cdot a[k])$ and $(1, 2 \cdot a[k] + 1)$.

All triangles trivially intersect the vertical line ℓ as they have a vertex lying on the y-axis. Let $k, k' \in \{0, \ldots, n-1\}$, such that $k \neq k'$.

Suppose that $a[k] = a[k']$. Because the corresponding triangles share two vertices with each other they have a non-empty intersection. Moreover, this intersection lies entirely to the right of the y-axis, as the respective leftmost vertices of the triangles lie on this line and are distinct.

Now suppose that $a[k] < a[k']$. Then it is easy to see that both triangles lie strictly on a different side of the line passing through the points with coordinates $(0, a[k] + n + 1/2)$ and $(1, 2 \cdot a[k] + 3/2)$, and are thus disjoint.

Thus, one gets a positive response to the Element Distinctness instance if and only if one gets a positive response to this constructed instance. Coupled with the $O(n)$ runtime of the reduction, this concludes the proof. ◀

The relevance of the requirement that all polygons intersect ℓ is that in the case of polygons with at most a constant number of vertices and having the (p, q) property, we can reduce the problem in Theorem 13 to this case in $\mathcal{O}(np)$ time.

3.3 The 3D case

Here we deal with the analogous case for the 3D polyhedra. In this case the Helly number becomes 4, and $x^*(\mathcal{F})$ is defined in terms of triplets of convex polyhedra with non-empty common intersection.

▶ **Theorem 17.** *Let (p, q) be an admissible pair for $d = 3$ and let $\mathcal{F}$ be a family of compact convex polyhedra in $\mathbb{R}^3$ with a total of n vertices and the (p, q)-property. We can compute a set of at most $p - q + 1$ points stabbing $\mathcal{F}$ in $\mathcal{O}((p - q + 1)n^{5/2} \log^{10} n (\log \log n)^{1/6} + np^3)$ expected time.*

Proof. First, we compute the family $\mathcal{F}_2$ of all polyhedra obtained as pairwise intersections of polyhedra in $\mathcal{F}$. This can be done in $\mathcal{O}(n^2)$ time using a linear-time algorithm to compute each intersection (for example the one by Chan [13]). Assuming the planes defining the polyhedra are in general position and none of the edges lie in a plane parallel to the yz-plane, the leftmost point in the intersection of three polyhedra in $\mathcal{F}$ is either

1. the leftmost point of a polyhedron in $\mathcal{F}_2$ contained in the interior of a polyhedron of $\mathcal{F}$,
2. the leftmost point of a polyhedron in $\mathcal{F}$ contained in the interior of a polyhedron of $\mathcal{F}_2$,
3. the intersection of an edge of a polyhedron in $\mathcal{F}_2$ and the interior of a facet of a polyhedron in $\mathcal{F}$ (note that not all such intersections define the leftmost point of the intersection of three polyhedra in $\mathcal{F}$).

The rightmost point corresponding to the first two cases can be found in expected time $\mathcal{O}(n^{9/4}\log^{\mathcal{O}(1)} n)$ by triangulating all polyhedra and then using the same methods as for the 2D case. We now focus on the third case. In what follows, we only deal with triangular facets, as the general case reduces to this one by triangulating all facets in $\mathcal{O}(n^2)$ total time. We will preprocess the triangles and then query each edge to count the number of triangles which intersect it and define the leftmost point of a three-wise intersection of polyhedra. Testing if a segment intersects a triangle in $\mathbb{R}^3$ can be done by comparing the signs of three polynomial functions of degree three on the coordinates of the points [37]. If $e = (p, q)$ is the segment we are testing against a triangular facet f, these polynomials take the form of the following determinant, where (a, b) is one of the three edges of f:

$$D(e,f) = \begin{vmatrix} p_x & p_y & p_z & 1 \\ q_x & q_y & q_z & 1 \\ a_x & a_y & a_z & 1 \\ b_x & b_y & b_z & 1 \end{vmatrix}.$$

It can be checked that testing $D(e, f) \geq 0$ can be expressed as testing if $P(f) \in H(e)$, where $P(f)$ is a point in $\mathbb{R}^5$ depending only on f and $H(e)$ is a halfspace of $\mathbb{R}^5$ depending only on e. The most convenient way is perhaps to use the algorithm described in a survey by Agarwal and Erickson [2, Section 5.2] (which is a very slight variant of an algorithm by Agarwal and Matoušek [3]), which computes a linearization of smallest dimension and simply involves computing the rank of a matrix whose coefficients depend on those of the polynomial to linearize.

Agarwal et al. [1] further show that when given an edge e of a polyhedron and a facet f of another polyhedron such that e and f intersect, testing if $e \cap f$ is the leftmost point of the intersection of the corresponding polyhedra can be expressed as testing if the outward normal vector of f lies in the intersection of three halfspaces (depending on e and the faces which support it).

Using again Theorem 7, this time in dimension $d = 5$ and with $k = 6$, we can preprocess the $\mathcal{O}(n^2)$ facets of $\mathcal{F}_2$ in $\mathcal{O}(n^{5/2}\log^{10} n(\log\log n)^{1/6})$ time such that we can query any oriented edge in $\mathcal{O}(n^{3/2}\log^{10} n(\log\log n)^{1/6})$ expected time. By querying all $\mathcal{O}(n)$ edges in $\mathcal{F}$ we can decide if a leftmost point in the intersection of three polyhedra in $\mathcal{F}$ corresponding to the third case exists in $\mathcal{O}(n^{5/2}\log^{10} n(\log\log n)^{1/6})$ total expected time. By applying Lemma 9 as we did in the planar case, we thus get the result. ◀

4 Other conditions

In this part we return to families of convex sets in the plane and investigate further conditions that are sufficient for the family to be stabbed by a fixed number of points. In the whole part, we will consider the planar case, but we expect that with some more care the results can be extended to higher dimensions.

The main reason why we consider the planar case is the following: all proofs below use the algorithm given above, all we do is showing that the algorithm is also correct under some assumptions other than the (p, q)-condition. In particular, we immediately get efficient algorithms for the below results.

4.1 Holes

The first condition we investigate considers *holes* in the union of sets. Let $\mathcal{F}$ be a finite family of convex sets in the plane and let $A \subset \mathbb{R}^2$ be the union of the sets in $\mathcal{F}$. A hole is a bounded connected component of $\mathbb{R}^2 \setminus A$.

There is an equivalent formulation of Helly's theorem due to Breen, which in the plane can be stated as follows: Let $\mathcal{F}$ be a finite family of pairwise intersecting convex sets in the plane with the property that the union of any three of them has no hole, then $\mathcal{F}$ can be stabbed by a single point [9, 34]. We prove the following generalization of this result.

▶ **Theorem 18.** *Let $\mathcal{F}$ be a finite family of pairwise intersecting convex sets in the plane with the property that the union of any $k + 3$ of them has at most k holes, then $\mathcal{F}$ can be stabbed by $k + 1$ points. Further, the $k + 1$ stabbing points can be chosen to lie on a single line.*

Proof. Let $x^*(\mathcal{F})$ be as above, that is, the lexicographical maximum among any lexicographical minimums in the intersection of two sets in $\mathcal{F}$, and let F_1 and F_2 be the sets in $\mathcal{F}$ that define $x^*(\mathcal{F})$. Consider the vertical line v through $x^*(\mathcal{F})$ and let F_1' and F_2' be the parts of F_1 and F_2, respectively, that lie to the left of v. Let now ℓ be a line through $x^*(\mathcal{F})$ which separates F_1' and F_2'. Such a line exists as otherwise $x^*(\mathcal{F})$ would not be the lexicographical minimums in the intersection of F_1' and F_2'. Further note that any set in $\mathcal{F}$ that is not stabbed by $x^*(\mathcal{F})$ must intersect ℓ to the left of $x^*(\mathcal{F})$: there cannot be intersections exclusively to the right of $x^*(\mathcal{F})$ by its definition, and as any set intersects F_1 and F_2, it follows from convexity that it must also intersect ℓ.

Let now $\mathcal{R}$ be the family of remaining sets, that is, the sets not stabbed by $x^*(\mathcal{F})$. We claim that among any $k + 1$ of them, some two intersect along ℓ. Indeed, if there were $k + 1$ sets whose intersections with ℓ are pairwise disjoint, the union of these sets with F_1 and F_2 would have $k + 1$ holes, which is excluded by the assumptions of the theorem. We can thus apply the Hadwiger-Debrunner (p, q)-theorem on ℓ to stab $\mathcal{R}$ with k points, so in total we have stabbed $\mathcal{F}$ with $k + 1$ collinear points. ◀

Note that opposed to the proof of the Hadwiger-Debrunner (p, q)-theorem, we only compute $x^*(\mathcal{F})$ once. After this, we only need the 1D-variant, where stabbing points of n intervals can easily be computed in time $O(n \log n)$. We thus get the following.

▶ **Proposition 19.** *Let $\mathcal{F}$ be a family of compact convex polygons in the plane with a total of n vertices and with the property that the union of any $k + 3$ of them has at most k holes. We can compute a set of at most $k + 1$ collinear points stabbing $\mathcal{F}$ in $\mathcal{O}(n^{4/3} \log^8 n (\log \log n)^{1/3})$ expected time.*

4.2 Number of intersections

Another result to which our algorithm can be applied is the following, due to Montejano and Soberón.

▶ **Theorem 20** (Montejano and Soberón [35]). *Let p, q, r, d be integers with $p > q > d$ and $r > \binom{p}{q} - \binom{p+1-d}{q+1-d}$. Let $\mathcal{F}$ be a family of convex sets in $\mathbb{R}^d$ with the property that for any p of them at least r of their q-tuples intersect. Then $\mathcal{F}$ can be stabbed by $p - q + 1$ points.*

In the plane, their proof is analogous to our proof of Theorem 18, that is, after removing all sets already stabbed by $x^*(\mathcal{F})$, the assumptions on the set can be used to show that the intersections of the remaining sets with the dividing line ℓ satisfy the $(p-2, q-1)$-property. In particular, the same algorithm to compute the stabbing points is correct, and we get the following.

▶ **Proposition 21.** *Let p, q, r be integers with $p > q > 2$ and $r > \binom{p}{q} - \binom{p-1}{q-1}$. Let $\mathcal{F}$ be a family of compact convex polygons in the plane with a total of n vertices and with the property that for any p of them at least r of their q-tuples intersect. We can compute a set of at most $p - q + 1$ collinear points stabbing $\mathcal{F}$ in $\mathcal{O}(n^{4/3} \log^8 n (\log \log n)^{1/3})$ expected time.*

5 A more general approach

We have seen in the last chapter an approach that uses Helly's theorem to prove the Hadwiger-Debrunner theorem. A natural path forward is to try adapting the method to other contexts where Helly-type theorems exist and prove corresponding (p, q) versions. By taking a close look at our proof for the Hadwiger-Debrunner (p, q)-theorem, we can observe that it relies almost exclusively on Lemma 3. This is strongly related with the notion of d-collapsibility, which is a property of the nerve of set-systems introduced by Wegner [40]. Lemma 3 is easily seen to imply that finite families of convex sets in $\mathbb{R}^d$ have d-collapsible nerves (this was already shown in a somewhat similar manner in Wegner's original paper) and it is a folklore result that this is a sufficient condition to prove Hadwiger and Debrunner's (p, q)-Theorem. Here we define *Ordered-d-collapsible systems* as set systems with a property analogous to Lemma 3 and give a proof of the folklore result that this implies such a (p, q)-Theorem. This proof has algorithmic consequences, some of which are given in the following section.

▶ **Definition 22** (Ordered-d-collapsible system).
Let $\mathcal{B}$ be a set, $\mathcal{D}$ be a family of subsets of $\mathcal{B}$, $d \geq 1$ be an integer and $\preceq$ be a total order on $\mathcal{B}$. We say that the tuple $(\mathcal{B}, \mathcal{D}, d, \preceq)$ is an Ordered-d-collapsible system if the following is true.

- *Let $\mathcal{F} \subset \mathcal{D}$ be a family of $n > d$ sets in $\mathcal{D}$ such that $I := \bigcap \mathcal{F}$ is non-empty. Let x be the $\preceq$-min of I (i.e. the smallest element of I with respect to the order $\preceq$, which we suppose exists). Then, there exists a subfamily $\mathcal{G} \subset \mathcal{F}$ of size d such that x is the $\preceq$-min of $\mathcal{G}$.*

The elements of $\mathcal{B}$ are called $\mathfrak{S}$-compact sets

As was stated earlier, this structure is enough to carry out a similar proof as the one we saw for the Hadwiger-Debrunner theorem. Call a pair (p, q) of integers d-admissible if $p \geq q \geq h$ and $(d-1)p < d(q-1)$. Then we have the following.

▶ **Theorem 23.** *Let $\mathfrak{S} = (\mathcal{B}, \mathcal{D}, d, \preceq)$ be an Ordered-d-collapsible system. Let (p, q) be a d-admissible pair of integers. Let $\mathcal{F}$ be a finite family of non-empty sets of $\mathcal{D}$. Suppose that $\mathcal{F}$ has the (p, q)-property. Then there exist $p - q + 1$ elements of $\mathcal{B}$ stabbing $\mathcal{F}$.*

Proof. Let $\mathfrak{S} = (\mathcal{B}, \mathcal{D}, d, \preceq)$ be an Ordered-d-collapsible system.

Let (p, q) be a d-admissible pair, and let $\mathcal{F}$ be a family of sets of $\mathcal{D}$ with the (p, q)-property.

We will reason by induction on p, the base case being $p = q = h$ which is true by the Helly property of the system. If $p = q > h$, then $\mathcal{F}$ also has the (h, h) property (as having the (p, q)-property implies having the $(p-1, q-1)$-property) and the results again follows from the Helly property of the system.

So suppose that $p > q$ and that the result is true for any d-admissible pair (p', q') with $p' < p$.

If $(h-2)p = (h-1)(q-1) - k - 1$ for $k \geq 1$, then notice that $(p-k, q-k)$ is an h-admissible pair, as in that case $(h-2)(p-k) = (h-1)(q-k-1) - 1$ which together with $p > q$ also implies that $q - k \geq h$. Thus the result follows from the induction hypothesis.

It now remains to consider the case where $p > q$ and $(h-2)p = (h-1)(q-1) - 1$.

To do so, let us construct an element $b^*(\mathcal{F})$ as follows:

- For every non-empty subfamily $\mathcal{S} \subset \mathcal{F}$ of $h-1$ sets with non-empty common intersection, let $b_{\mathcal{S}}$ be the $\preceq$-min of $I_{\mathcal{S}} = \bigcap_{C \in \mathcal{S}} C$, which exists by definition of an Ordered-d-collapsible system.
- Let $b^*(\mathcal{F})$ be the $\preceq$-max element among all such $b_{\mathcal{S}}$'s.

Let $\mathcal{G}$ be one of the families defining $b^*(\mathcal{F})$, that is, $\mathcal{G} \subset \mathcal{F}'$ is a subfamily of $h-1$ sets which have $b^*(\mathcal{F})$ as the $\preceq$-min of their intersection.

To establish the theorem, it is enough to show that by choosing $b^*(\mathcal{F})$ as one of our stabbing elements, we can stab all the remaining sets (i.e. those which do not contain $b^*(\mathcal{F})$) with $p - q$ elements. Let $\mathcal{R} = \{S \in \mathcal{F} \mid b^*(\mathcal{F}) \notin S\}$ be the set of remaining sets.

Let us argue that for any $S \in \mathcal{R}$, $S \cap (\bigcap \mathcal{G})$ is empty. To do so, suppose it was not, and let y be the $\preceq$-min of that intersection. By Lemma 25, y is the $\preceq$-min of the intersection of $h-1$ sets in $\mathcal{F}$. Moreover, by definition of $\mathcal{R}$ and $\mathcal{G}$, $b^*(\mathcal{F}) \prec y$. This contradicts the definition of $b^*(\mathcal{F})$. Thus, $S \cap (\bigcap \mathcal{G})$ is empty.

Two cases arise:

1. ($|\mathcal{R}| \geq p - h + 1$) We show that $\mathcal{R}$ has the $(p-h+1, q-h+2)$-property. Indeed, choose any $p - h + 1$ members from $\mathcal{R}$ together with the $h-1$ members from $\mathcal{G}$. We know from the (p, q)-property of $\mathcal{F}$ that there exists a subfamily $\mathcal{E} \subset \mathcal{R} \cup \mathcal{G}$ of size q whose members have a non-empty common intersection. $\mathcal{E}$ cannot contain all elements of $\mathcal{G}$, as $q > h - 1 = |\mathcal{G}|$ and the intersection of all members of $\mathcal{G}$ together with any member of $\mathcal{R}$ is empty. Thus, $\mathcal{E}$ contains at least $q - h + 2$ members of $\mathcal{R}$. This shows that $\mathcal{R}$ has the $(p-h+1, q-h+2)$-property. Notice that with the assumptions $p > q$ and $(h-2)p = (h-1)(q-1) - 1$ which we are working under, $(p-h+1, q-h+2)$ is admissible. Thus, by the induction hypothesis, $\mathcal{R}$ can be stabbed with $p-h+1-(q-h+2)+1 = p-q$ elements of $\mathcal{B}$.
2. ($|\mathcal{R}| < p - h + 1$) In this case, choose $\mathcal{R}$ as a whole together with $\mathcal{G}$ and $p - h + 1 - |\mathcal{R}|$ other members of $\mathcal{F}$. By the same reasoning as in case 1., there exists a subset of $q - (h - 2 + p - h + 1 - |\mathcal{R}|) = |\mathcal{R}| + 1 + q - p$ members of $\mathcal{R}$ which intersect and can thus be stabbed by a single element of $\mathcal{B}$. The remaining $|\mathcal{R}| - (|\mathcal{R}| + 1 + q - p) = p - q - 1$ sets can trivially be stabbed by $p - q - 1$ elements.

Thus, $\mathcal{R}$ can be stabbed by $p - q$ elements, which implies that $\mathcal{F}$ can be stabbed by $p - q + 1$ elements and by induction, concludes the proof. ◀

An almost identical proof shows the folklore result that under the weaker condition that the nerve of $\mathcal{F}$ is a d-collapsible system and $\mathcal{F}$ has the (p, q) condition the same conclusion holds.

In proving Lemma 3, we made use of relatively few properties of compact convex sets. These properties are (i) closure under intersection, (ii) existence of a lexicographically minimum point, (iii) Helly's theorem as well as (iv) the fact that the set of all points lexicographically smaller than some point y is convex. In what follows we will actually always consider families of sets with analogous properties, which we call *Ordered-Helly systems.*

▶ **Definition 24** (Ordered-Helly system).
An Ordered-Helly system $\mathfrak{S}$ is a tuple $(\mathcal{B}, \mathcal{C}, \mathcal{D}, d, \preceq)$ consisting of

- *a set $\mathcal{B}$, called the* base-set*;*
- *a family $\mathcal{C}$ of subsets of $\mathcal{B}$, whose members are called* convex sets *or* $\mathfrak{S}$-convex sets*;*
- *a family $\mathcal{D} \subset \mathcal{C}$ whose members are called* compact sets *or* $\mathfrak{S}$-compact sets*;*
- *a total order $\preceq$ on $\mathcal{B}$;*
- *and an integer $h \geq 2$, called the* Helly-number *of* $\mathfrak{S}$

with the following properties.

1. *(Intersection closure)*
   $\mathcal{D}$ is closed under intersection, i.e. for all $S_1, S_2 \in \mathcal{D}$ we have $S_1 \cap S_2 \in \mathcal{D}$.
2. *(Attainable minimum)*
   For all non-empty $S \in \mathcal{D}$, there exists $x \in S$ such that for all $y \in S$, $x \preceq y$. This x is necessarily unique and we call x the $\preceq$-min of S. We define the $\preceq$-max of a set similarly, if it exists.
3. *(Convex order)*
   For all $t \in \mathcal{B}$, we have $\{x \in \mathcal{B} \mid x \preceq t \text{ and } x \neq t\} \in \mathcal{C}$.
4. *(Helly property)*
   If $\mathcal{F} \subset \mathcal{C}$ is a finite subset of $n \geq d+1$ sets of $\mathcal{C}$ such that every subfamily of $d+1$ members of $\mathcal{F}$ has a non-empty common intersection, then all members of $\mathcal{F}$ have a non-empty common intersection.

Let us show that this is indeed a stronger set of conditions, in a proof analogous to that of Lemma 3.

▶ **Lemma 25.** *Let $\mathfrak{S} = (\mathcal{B}, \mathcal{C}, \mathcal{D}, d, \preceq)$ be an Ordered-Helly system. Then $(\mathcal{B}, \mathcal{D}, d, \preceq)$ is an Ordered-d-collapsible system.*

Proof. Let $\mathcal{F} \subset \mathcal{D}$ be a family of $n \geq h$ sets in $\mathcal{D}$ such that $I := \bigcap \mathcal{F}$ is non-empty. Let x be the $\preceq$-min of I (which exists by the properties of intersection closure and attainable minimum).
Let S_x denote $\{y \in \mathcal{B} \mid y \preceq x \text{ and } y \neq x\}$, which is a $\mathfrak{S}$-convex set by the property of convex order. It is also disjoint from I as $\preceq$ is a total order. By the Helly property, there exists a subfamily of $d+1$ members of $\mathcal{F} \cup \{S_x\}$ with an empty common intersection. These members have to include S_x, as all members of $\mathcal{F}$ have a non-empty common intersection. Let $\mathcal{G} \subset \mathcal{F}$ be the family consisting of the remaining d sets and let $x_{\mathcal{G}}$ be the $\preceq$-min of $I' := \bigcap \mathcal{G}$ (which is a non-empty $\mathfrak{S}$-compact set). We know that $x_{\mathcal{G}} \preceq x$ because $x \in I'$. If we now suppose $x_{\mathcal{G}} \neq x$ this implies that $x_{\mathcal{G}} \in S_x$ and contradicts the fact that $I' \cap S_x = \emptyset$. Thus, $x_{\mathcal{G}} = x$ and $(\mathcal{B}, \mathcal{D}, d, \preceq)$ is an Ordered-d-collapsible system. ◀

We immediately have the following corollary.

▶ **Corollary 26.** *Let $\mathfrak{S} = (\mathcal{B}, \mathcal{C}, \mathcal{D}, h, \preceq)$ be an Ordered-Helly system. Let (p, q) be a d-admissible pair of integers. Let $\mathcal{F}$ be a finite family of non-empty sets of $\mathcal{D}$. Suppose that $\mathcal{F}$ has the (p, q)-property. Then there exist $p - q + 1$ elements of $\mathcal{B}$ stabbing $\mathcal{F}$.*

It should be mentioned that the existence of a Helly number alone is not enough to show such a result. Alon et al. [4] give an example of a set system with Helly number 2 but no general (p, q)-theorem. Note also that many other classical results about the intersection patterns of convex sets, including Alon and Kleitmann's (p, q)-Theorem, also follow from the notion of d-collapsibility and thus hold for Ordered-d-collapsible and Ordered-Helly systems (see the survey by Tancer [39]).

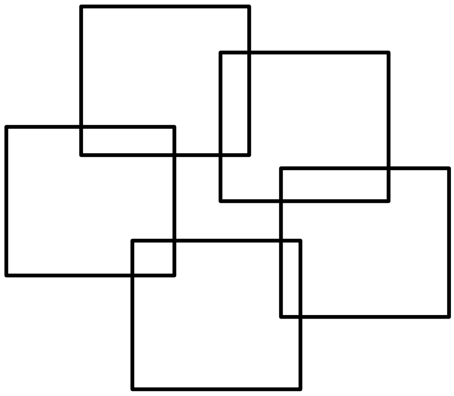

Figure 2 Counterexample to the claim that axis aligned rectangles in the plane with the $(3, 2)$-property can always be stabbed with two points.

▶ Remark. For such a result to hold, none of the conditions 1. to 3. of Ordered-Helly systems are necessary in the sense that there exist families of sets violating all three for which a Hadwiger-Debrunner type theorem does hold. Consider for example the family of all open disks in the plane with the lexicographical order on points. Neither intersection closure, attainable minimum nor convex order holds in this case, but of course the Hadwiger-Debrunner theorem still applies as these are a special case of convex sets.

However, condition 3. (convex order) is in fact necessary in the sense that dropping it while maintaining the others unchanged would make Corollary 26 false. Otherwise, we could for example prove that a family of axis aligned rectangles in the plane with the $(3, 2)$ property can be stabbed with two points. This is false, as Figure 2 illustrates.

6 Computing stabbing points in an Ordered-d-collapsible system

The proof we saw once again leads to an algorithm computing stabbing elements of a family of $\mathfrak{S}$-compact sets with the (p, q)-property for an admissible pair (p, q), given we have access to some oracles. We will write the run-times in terms of the description complexity of a set, which depends on the exact context. Thus, for a $\mathfrak{S}$-compact set S, let $\#S$ denote this complexity (of at least 1), and for a family $\mathcal{F}$ of $\mathfrak{S}$-compact sets, let $\#\mathcal{F} := \sum_{S \in \mathcal{F}} \#S$.

Consider an Ordered-d-collapsible system $\mathfrak{S} = (\mathcal{B}, \mathcal{D}, d, \preceq)$ (for a constant d) and suppose we have access to the following oracles.

- For two elements $b_1, b_2 \in \mathcal{B}$, we can test $b_1 \preceq b_2$ in constant time.
- For a family of at most d $\mathfrak{S}$-compact sets $\mathcal{F} \subset \mathcal{D}$, we can test if the sets in $\mathcal{F}$ have a common intersection and compute the $\preceq$-min of that intersection if it is non-empty in $\mathcal{O}(\#\mathcal{F})$ time.
- For a $\mathfrak{S}$-compact set $S \in \mathcal{D}$ and a point $b \in \mathcal{B}$ we can test if $b \in S$ in $\mathcal{O}(\#S)$ time.

We could naturally consider other run-times for these oracles. We only specify them in order to showcase an example of run-time analysis which is tighter than if we had worked with general run-times and swapped in concrete functions afterwards (and matches the case of convex polytopes in $\mathbb{R}^d$ for small d). Other run-times might require other specialized forms of analysis.

Now, let $\mathcal{F} \subset \mathcal{D}$ be a family of $\mathfrak{S}$-compact sets. Among all points in $\mathcal{B}$ defined as the $\preceq$-min of the intersection of $h - 1$ sets in $\mathcal{F}$, let $b^*(\mathcal{F})$ be the $\preceq$-max of those.

Let us state two lemmas which will be useful for our algorithm. The proofs are identical to those of Lemma 4 and Lemma 5 and we repeat them only for the reader's convenience.

▶ **Lemma 27.** *We can compute $b^*(\mathcal{F})$ in $\mathcal{O}(\#\mathcal{F} \cdot |\mathcal{F}|^{d-1})$ time.*

Proof. We can compute $b^*(\mathcal{F})$ by testing for intersection in every subfamily $\mathcal{G}$ of $\mathcal{F}$ of size d and computing the $\preceq$-min of that intersection if it is non-empty.

If we consider some fixed subfamily $\mathcal{G}$, the computation for that subfamily will cost at most $c \cdot \#\mathcal{G}$ for some constant c which doesn't depend on $\mathcal{G}$. Charge this cost to the sets $S \in \mathcal{G}$ by attributing a cost of $c \cdot \#S$ to a set S.

Now, consider the cost charged to some fixed set S for the whole computation. As S appears in no more than $|\mathcal{F}|^{d-1}$ subfamilies of size $h-1$, its total cost charge is upper bounded by $c \cdot \#S \cdot |\mathcal{F}|^{d-1}$. Summing across all sets $S \in \mathcal{F}$, we get a total cost of $\mathcal{O}(\#\mathcal{F} \cdot |\mathcal{F}|^{d-1})$. ◀

▶ **Lemma 28.** *Suppose there exists some subfamily $\mathcal{G} \subset \mathcal{F}$ of size $k+1$ such that all sets in $\mathcal{G}$ have a common intersection, where k is a known parameter. We can compute $|\mathcal{F}| - k$ points in $\mathcal{B}$ stabbing $\mathcal{F}$ in $\mathcal{O}(\#\mathcal{F} \cdot |\mathcal{F}|^d)$ time.*

Proof. If $k+1 \leq d$, then we can test every subfamily of size $k+1$ for common intersection and compute its $\preceq$-min for a total cost of $\mathcal{O}(\#\mathcal{F} \cdot |\mathcal{F}|^{k+1}) \leq \mathcal{O}(\#\mathcal{F} \cdot |\mathcal{F}|^d)$.

If $k+1 > d$, then we know from Lemma 25 that the $\preceq$-min of the intersection of all sets in $\mathcal{G}$ is also the $\preceq$-min of the intersection of some d sets in $\mathcal{F}$. Thus, one can find a point stabbing at least $k+1$ sets by computing the $\preceq$-min point for each subfamily of size d (in $\mathcal{O}(\#\mathcal{F}|\mathcal{F}|^d)$ time) and counting the number of sets intersected for each of the $\mathcal{O}(|\mathcal{F}|^d)$ computed points (in $\mathcal{O}(\#\mathcal{F}|\mathcal{F}|^d)$ time as well).

As soon as we find a point b stabbing at least $k+1$ sets, we return b along with the $\preceq$-min of every set in $\mathcal{F}$ which is not stabbed by b. ◀

With these algorithms, we can now prove the following.

▶ **Theorem 29.** *Let $\mathcal{F}$ be a family of $\mathfrak{S}$-compact sets with the (p, q)-property. Suppose we have access to the relevant oracles described above. We can compute a set of at most $p-q+1$ elements stabbing $\mathcal{F}$ in time*

$$\mathcal{O}((p-q+1)(\#\mathcal{F})^d + (\#\mathcal{F})p^d).$$

Proof.
Consider the following algorithm.

1. Reduce p and q (as done in the proof of Corollary 26) to reach the case where $p = q = d+1$ or the case where $p > q$ and $(d-1)p = d(q-1) - 1$.
2. Compute $b^*(\mathcal{F})$ and choose it as one of the stabbing points. Now, remove from $\mathcal{F}$ all the sets which are stabbed by this point. If there are any remaining sets then either $|\mathcal{F}| \geq p-d$ and $\mathcal{F}$ satisfies the $(p-d, q-d+1)$-property, where $(p-d, q-d+1)$ is d-admissible, or $\mathcal{F}$ consists of $p-q+k$ sets, $k < q-d$, where some $k+1$ of them have a common intersection.
3. In the first case, we can continue inductively.
4. In the second case we can trivially stab the remaining sets using $p-q$ elements.

- *Correctness:*

The correctness of the algorithm follows from the proof of Corollary 26. The only detail that needs some additional scrutiny is the correctness for the base case $p = q = d+1$. Notice that in this case all sets have a common intersection and the definition of an Ordered-d-collapsible system ensures that $b^*(\mathcal{F})$ stabs the whole family $\mathcal{F}$.

- *Runtime:*

We know that Step 1 can be done in $\mathcal{O}((\#\mathcal{F})^d)$ time and has to be done at most $(p-q+1)$ times. We also know from Lemma 28 that Step 4 can be done in $\mathcal{O}((\#\mathcal{F})^d + (\#\mathcal{F})p^d)$ time. This step is only done once.

Thus, we get a total runtime of

$$\mathcal{O}((p-q+1)(\#\mathcal{F})^d + (\#\mathcal{F})p^d).$$ ◀

With access to the right oracle, we could for example apply Lemma 9 analogously to what we did for convex polytopes in the Euclidean setting and get the corresponding speedup.

7 Examples of Ordered-Helly systems

Until now, the only Ordered-Helly system we have seen is the one corresponding to compact convex sets in $\mathbb{R}^d$. We will see that this structure does have some other interesting representatives and is not restricted to this single example (in which case the usefulness of introducing it would have been doubtful).

7.1 Hadwiger-Debrunner type results for subsets of $\mathbb{R}^d$

Let us start by stating and proving some Hadwiger-Debrunner type results for sets which are defined as the intersection of a compact convex set in $\mathbb{R}^d$ with a subset $S \in \mathbb{R}$. De Loera et al. [16] introduced the notion of S-Helly number, which can be defined as follows.

▶ **Definition 30.** *Let S be a subset of $\mathbb{R}^d$. The S-Helly number, denoted by $h(S)$, is the smallest integer $k > 0$ such that the following holds:*
Given a finite family $\mathcal{F}$ of convex sets in $\mathbb{R}^d$, if in every subfamily of $\mathcal{F}$ of size k all sets share a point in S, then all sets in $\mathcal{F}$ share a point in S.
If no such k exists, then $h(S) = \infty$.

One of the first results concerning S-Helly numbers was discovered by Doignon [17], and is the case $S = \mathbb{Z}^d$.

▶ **Theorem 31** (Doignon). *Let $\mathcal{F}$ be a family of $n \geq 2^d$ convex sets in $\mathbb{R}^d$. If in every subfamily of $\mathcal{F}$ of size 2^d all sets share a point in $\mathbb{Z}^d$, then all sets in $\mathcal{F}$ share a point in $\mathbb{Z}^d$.*

Alon et al. [4] showed that the nerve of such families is $2^d - 1$-collapsible, which among other things implies a Hadwiger-Debrunner type (p,q)-theorem.

In a paper by Hoffman [26] a mixed-integer version of this theorem is stated, which generalizes both Helly's theorem and Doignon's version. It was later rediscovered and proved in detail by Averkov and Weismantel [6].

▶ **Theorem 32** (Mixed-Integer Helly). *Let $\mathcal{F}$ be a family of $n \geq (d+1)2^k$ convex sets in $\mathbb{R}^{d+k}$, where $d,k \geq 0$ and $d+k \geq 1$. If in every subfamily of $\mathcal{F}$ of size $(d+1)2^k$ all sets share a point in $\mathbb{R}^d \times \mathbb{Z}^k$, then all sets in $\mathcal{F}$ share a point in $\mathbb{R}^d \times \mathbb{Z}^k$.*

Let us show that a corresponding Hadwiger-Debrunner-type theorem holds.

▶ **Theorem 33** (Mixed-Integer Hadwiger-Debrunner).
Let $d,k \geq 0$ be integers such that $d+k \geq 1$. Let (p,q) be a $((d+1)2^k - 1)$-admissible pair.

Let $\mathcal{F}$ be a finite family of sets obtained as the intersection of $\mathbb{R}^d \times \mathbb{Z}^k$ with a compact convex set in $\mathbb{R}^{d+k}$. Suppose that $\mathcal{F}$ has the (p,q) property. Then there exist $p-q+1$ points in $\mathbb{R}^d \times \mathbb{Z}^k$ stabbing $\mathcal{F}$.

Proof. Let $\mathcal{B} = \mathbb{R}^d \times \mathbb{Z}^k$, let $\mathcal{C}$ be the family of all sets obtained as the intersection of $\mathcal{B}$ with a convex set in R^{d+k} and let $\mathcal{D}$ be the family of all sets obtained as the intersection of $\mathcal{B}$ with a compact convex set in R^{d+k}. Then, using Theorem 32, it is easy to check that $(\mathcal{B}, \mathcal{C}, \mathcal{D}, (d+1)2^k - 1, \leq_{lex})$ is an Ordered-Helly system. Thus, using Corollary 26, we get the result. ◀

More generally, every upper bound on an S-Helly number leads to a corresponding Hadwiger-Debrunner version if S is closed in $\mathbb{R}^d$. In particular, this applies to some of the bounds obtained by De Loera et al. [16] for several other families of sets S. The corresponding algorithmic results also follow, provided we have access to the required oracles.

7.2 Abstract convex geometries

Let us now explore how the structure of Ordered-Helly systems relates to the structure of abstract convex geometries as introduced by Edelman and Jamison [19]. Convex geometries are an abstraction capturing the basic combinatorial structure of classical convexity in a similar manner to matroids capturing the basic combinatorial properties of linear independence. Convex geometries appear in many contexts outside of convex sets such as graph theory or order theory. We refer the interested reader to the book of Edelman and Jamison [19] or to Chapter III in the book of Korte et al. [30] for an in-depth overview. We will only go over the basic definitions and theorems needed for our purpose, which can all be found in these two sources.

Some background

For the following definitions, it is useful to imagine the operator τ as analogous to the convex hull operator on a point set.

▶ **Definition 34.** *Consider some finite set E and a family $\mathcal{N}$ of subsets of E. Let τ be the operator defined on subsets of E as $\tau(A) = \bigcap\{X \mid A \subset X,\ X \in \mathcal{N}\}$. We say that $(E, \mathcal{N})$ is a convex geometry if it has the following properties.*

1. *$\emptyset \in \mathcal{N}$, $E \in \mathcal{N}$.*
2. *$X, Y \in \mathcal{N}$ implies $X \cap Y \in \mathcal{N}$.*
3. *If $y, z \notin \tau(X)$ and $z \in \tau(X \cup \{y\})$ then $y \notin \tau(X \cup \{z\})$.*

The sets in $\mathcal{N}$ are called convex.

Extreme points are defined in a similar way way as in the Euclidean setting:

▶ **Definition 35.** *For a set $A \subset E$, we say that $x \in A$ is an extreme point of A if $x \notin \tau(A \backslash \{x\})$. The set of extreme points of A is denoted by $ex(A)$.*

A set $X \subset E$ is called free if $X = ex(X)$.

We will use the following notion.

▶ **Definition 36.** *A sequence $x_1, \ldots, x_k$ of points of E is called a shelling sequence if for all $1 \leq i \leq k$, x_i is an extreme point of $E \setminus \{x_1, \ldots, x_{i-1}\}$.*

A shelling sequence can be thought of as a way to reach a convex set by starting with the whole set E and stripping away points one after the other in such a way that the set remains convex at each step. A useful characterisation of convex sets for our purpose is the following, where we describe a convex set via a shelling process.

▶ **Proposition 37** ([19]). *A set $X \subset E$ is convex if and only if there exists a shelling sequence $x_1, \ldots, x_k$ such that $X = E \setminus \{x_1, \ldots, x_k\}$.*

The final ingredient we need is the following Helly-type theorem for convex geometries.

▶ **Theorem 38** ([19]). *Let $h(\mathcal{N})$ denote the smallest integer k such that the following holds:*

For a family $\mathcal{F}$ of convex sets, if every subfamily of size at most k has a non-empty intersection, then $\mathcal{F}$ has a non-empty intersection.

Then the Helly number of $\mathcal{N}$, $h(\mathcal{N})$, is equal to the maximum size of a free convex set.

Hadwiger-Debrunner theorem for convex geometries

We are now ready to state and prove a Hadwiger-Debrunner-type theorem for convex geometries.

▶ **Theorem 39.** *Consider a convex geometry $(E, \mathcal{N})$. Let h be the size of a maximal free convex set and let (p, q) be an $(h-1)$-admissible pair. Let $\mathcal{F} \subset \mathcal{N}$ be a family of $n \geq p$ non-empty convex sets. If $\mathcal{F}$ has the (p, q)-property then there exist $p - q + 1$ elements of E stabbing $\mathcal{F}$.*

Proof. We know that $\emptyset$ is convex, thus there exist a shelling sequence $S = \{x_1, \ldots, x_k\}$ such that $\emptyset = E \setminus S$, i.e. $S = E$. Let for $1 \leq i, j \leq k$, let us say that $x_i \preceq x_j$ if and only if $i \geq j$. Let $1 \leq t \leq k$ be an integer. Because $\{x_1, x_2, \ldots, x_{t-1}\}$ is a valid shelling sequence, $\{x \in E \mid x \leq x_t\}$ is a convex set. Thus, $\preceq$ has the convex order property.

Let h be the maximum size of a free convex set. Then, it is easy to verify that $\mathfrak{S} = (E, \mathcal{N}, \mathcal{N}, h-1, \preceq)$ also has the intersection closure and attainable minimum properties. The Helly property (for Helly number h) is given by Theorem 38.

Thus, $\mathfrak{S}$ is an Ordered-Helly system and we get the result from Corollary 26. ◀

Two examples of convex geometries

We will now give two illustrative examples of abstract convex geometries and the resulting Hadwiger-Debrunner type results we obtain for them. One such convex geometry (arguably the most natural) is the one obtained by taking convex hulls of subsets on a finite point set in $\mathbb{R}^d$. This is conceptually similar to the case of polytopes in Euclidean space which we have already discussed. The following examples are perhaps not so immediately related.

We first consider the convex geometry obtained by taking subtrees of a tree.

▶ **Proposition 40** ([19]). *Let T be a tree on a set of vertices V. Let $\mathcal{N}$ be the family of all sets of vertices corresponding to subtrees of T. Then $(V, \mathcal{N})$ is a convex geometry with Helly number $h(\mathcal{N}) = 2$.*

This means that for a given tree T and a given family of subtrees of T, if all pairs of subtrees intersect at some vertex, then all subtrees share a vertex. Using Theorem 39 we can thus get the following result.

▶ **Corollary 41.** *Let T be a tree and let $\mathcal{F}$ be a family of subtrees of T (represented as sets of vertices). Let (p, q) be a 1-admissible pair. Let $\mathcal{F} \subset \mathcal{N}$ be a family of non-empty subtrees of T with the (p, q)-property. Then $\mathcal{F}$ can be stabbed with $p - q + 1$ vertices.*

From an algorithmic point of view, let us suppose that the tree is represented as a conventional pointer structure and that the subtrees in $\mathcal{F}$ are themselves represented in full as trees. One can compute a shelling sequence of the empty tree (and thus $\preceq$) by starting with the whole tree T and choosing leaves to cut off until we reach the empty tree. This amounts to $O(|V|)$ preprocessing time. We can trivially find the $\preceq$-min of a subtree or test if a subtree S contains a vertex in $O(|V(S)|)$ time. Using Theorem 29, this leads to an algorithm finding stabbing vertices in $\mathcal{O}(|V| + p(\#\mathcal{F}))$ where $\#\mathcal{F}$ is the sum of the number of vertices over all subtrees in $\mathcal{F}$.

Another example of convex geometry is the one obtained by taking so-called ideals of a poset. For a poset $(E, \leq)$, we say that a set $S \subset E$ is an ideal of E if for all $x \in S$ and all $y \in E$, $y \leq w \Rightarrow y \in S$. Let $width(E)$ denote the maximum size of an antichain in E. Then the following holds.

▶ **Proposition 42** ([19]). *Let $(E, \leq)$ be a finite poset. Let $\mathcal{F} = \{S \subset E \mid S \text{ is an ideal}\}$. Then $(E, \mathcal{F})$ is a convex geometry with Helly number $width(E)$.*

Using Theorem 39 we can thus get the following result.

▶ **Corollary 43.** *Let $(E, \leq)$ be a finite poset. Let (p, q) be a $(width(E) - 1)$-admissible pair and let $\mathcal{F}$ be a family of non-empty ideals of E with the (p, q)-property. Then $\mathcal{F}$ can be stabbed by $p - q + 1$ elements of E.*

From an algorithmic point of view, the situation is similar to the one for subtrees of a tree if we choose to represent ideals as the sets of their elements.

Conclusion

We have shown how to stab convex polygons with a total of n vertices and the (p, q)-property (for admissible (p, q)) in expected $\tilde{\mathcal{O}}(n^{4/3})$ time with respect to n. As an intermediate step, we compute a certain quantity x_2^*, which is a Hopcroft-Hard problem, in $\tilde{\mathcal{O}}(n^{4/3})$ expected time. While this is believed to be near optimal, finding a non-trivial lower-bound for the original problem remains open. For the 3D case, we have an algorithm running in expected $\tilde{\mathcal{O}}(n^{5/2})$ time with respect to n.

We have also considered other conditions which allow us to conclude that a set of polygons can be stabbed with a fixed number of points, and applied our algorithm to those. One of these conditions is a new generalization of Helly's theorem in the plane in terms of holes in the union of convex sets.

Finally, we have derived (p, q)-theorems (some of which are folklore results) along with algorithms in other settings where Helly-type theorems are known.

A natural next step in the Euclidean setting would be to drop the restriction $(d - 1)p < d(q - 1)$ and find efficient algorithms for the Alon-Kleitman (p, q)-theorem. Their proof of existence of stabbing sets of constant size uses the fractional Helly theorem, whose proof is similar to the above proof of the Hadwiger-Debrunner (p, q)-theorem. It is thus conceivable that similar ideas could be applied to this more general case. Another natural question is to find an efficient algorithm to test whether a set of convex polygons in the plane has the (p, q) property for given p and q. Ideally, such an algorithm would run in polynomial time, with an exponent which does not depend on p or q.

References

1 Pankaj K. Agarwal, Mark de Berg, Sariel Har-Peled, Mark H. Overmars, Micha Sharir, and Jan Vahrenhold. Reporting intersecting pairs of convex polytopes in two and three dimensions. *Comput. Geom.*, 23(2):195–207, 2002. doi:10.1016/S0925-7721(02)00049-4.

2 Pankaj K. Agarwal and Jeff Erickson. Geometric range searching and its relatives. In *Advances in discrete and computational geometry. Proceedings of the 1996 AMS-IMS-SIAM joint summer research conference on discrete and computational geometry: ten years later, South Hadley, MA, USA, July 14–18, 1996*, pages 1–56. Providence, RI: American Mathematical Society, 1999.

3 Pankaj K. Agarwal and Jirí Matousek. On range searching with semialgebraic sets. *Discret. Comput. Geom.*, 11:393–418, 1994. doi:10.1007/BF02574015.

4 Noga Alon, Gil Kalai, Jirí Matousek, and Roy Meshulam. Transversal numbers for hypergraphs arising in geometry. *Adv. Appl. Math.*, 29(1):79–101, 2002. doi:10.1016/S0196-8858(02)00003-9.

5 Noga Alon and Daniel J. Kleitman. Piercing convex sets. *Bulletin of the American Mathematical Society*, 27(2):252–257, August 1992. doi:10.1090/s0273-0979-1992-00304-x.

6 G. Averkov and R. Weismantel. Transversal numbers over subsets of linear spaces. *Adv. Geom.*, 12(1):19–28, 2012. doi:10.1515/advgeom.2011.028.

7 Michael Ben-Or. Lower bounds for algebraic computation trees (preliminary report). In David S. Johnson, Ronald Fagin, Michael L. Fredman, David Harel, Richard M. Karp, Nancy A. Lynch, Christos H. Papadimitriou, Ronald L. Rivest, Walter L. Ruzzo, and Joel I. Seiferas, editors, *Proceedings of the 15th Annual ACM Symposium on Theory of Computing, 25-27 April, 1983, Boston, Massachusetts, USA*, pages 80–86. ACM, 1983. doi:10.1145/800061.808735.

8 Jon Louis Bentley and Thomas Ottmann. Algorithms for reporting and counting geometric intersections. *IEEE Trans. Computers*, 28(9):643–647, 1979. doi:10.1109/TC.1979.1675432.

9 Marilyn Breen. Starshaped unions and nonempty intersections of convex sets in $\mathbb{R}^d$. *Proc. Am. Math. Soc.*, 108(3):817–820, 1990. doi:10.2307/2047806.

10 Paz Carmi, Matthew J. Katz, and Pat Morin. Stabbing Pairwise Intersecting Disks by Four Points. *arXiv e-prints*, December 2018. arXiv:1812.06907.

11 Timothy M. Chan. Geometric applications of a randomized optimization technique. *Discret. Comput. Geom.*, 22(4):547–567, 1999. doi:10.1007/PL00009478.

12 Timothy M. Chan. Optimal partition trees. *Discret. Comput. Geom.*, 47(4):661–690, 2012. doi:10.1007/s00454-012-9410-z.

13 Timothy M. Chan. A simpler linear-time algorithm for intersecting two convex polyhedra in three dimensions. *Discret. Comput. Geom.*, 56(4):860–865, 2016. doi:10.1007/s00454-016-9785-3.

14 Justin Dallant and Patrick Schnider. Effciently stabbing convex polygons and variants of the hadwiger-debrunner (p, q)-theorem. In Meng He and Don Sheehy, editors, *Proceedings of the 33rd Canadian Conference on Computational Geometry, CCCG 2021, August 10-12, 2021, Dalhousie University, Halifax, Nova Scotia, Canada*, pages 157–166, 2021.

15 L. Danzer. Zur Lösung des Gallaischen Problems über Kreisscheiben in der Euklidischen Ebene. *Stud. Sci. Math. Hung.*, 21:111–134, 1986.

16 J. A. De Loera, R. N. La Haye, D. Oliveros, and E. Roldán-Pensado. Helly numbers of algebraic subsets of $\mathbb{R}^d$ and an extension of Doignon's theorem. *Adv. Geom.*, 17(4):473–482, 2017. doi:10.1515/advgeom-2017-0028.

17 Jean-Paul Doignon. Convexity in cristallographical lattices. *J. Geom.*, 3:71–85, 1973. doi:10.1007/BF01949705.

18 Jürgen Eckhoff. A survey of the Hadwiger-Debrunner (p, q)-problem. In *Discrete and Computational Geometry: The Goodman-Pollack Festschrift*, pages 347–377. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003. doi:10.1007/978-3-642-55566-4_16.

19 Paul H. Edelman and Robert E. Jamison. The theory of convex geometries. *Geometriae Dedicata*, 19(3), December 1985. doi:10.1007/bf00149365.

20 Jeff Erickson. New lower bounds for Hopcroft's problem. *Discret. Comput. Geom.*, 16(4):389–418, 1996. doi:10.1007/BF02712875.

21 Anka Gajentaan and Mark H. Overmars. On a class of o(n2) problems in computational geometry. *Comput. Geom.*, 5:165–185, 1995. doi:10.1016/0925-7721(95)00022-2.

22 Hugo Hadwiger and H. Debrunner. Über eine Variante zum Hellyschen Satz. *Arch. Math.*, 8:309–313, 1957. doi:10.1007/BF01898794.

23 Hugo Hadwiger, Hans Emil Debrunner, and Victor Klee. *Combinatorial Geometry in the Plane*. New York (N.Y.) : Holt, 1964.

24 Sariel Har-Peled, Haim Kaplan, Wolfgang Mulzer, Liam Roditty, Paul Seiferth, Micha Sharir, and Max Willert. Stabbing pairwise intersecting disks by five points. *Discret. Math.*, 344(7):112403, 2021. doi:10.1016/j.disc.2021.112403.

25 Ed. Helly. Über Mengen konvexer Körper mit gemeinschaftlichen Punkten. *Jahresber. Dtsch. Math.-Ver.*, 32:175–176, 1923. URL: http://eudml.org/doc/145659.

26 Alan J. Hoffman. Binding constraints and Helly numbers. *Annals of the New York Academy of Sciences*, 319:284 – 288, 12 2006. doi:10.1111/j.1749-6632.1979.tb32803.x.

27 Andreas F Holmsen and Dong-Gyu Lee. Radon numbers and the fractional Helly theorem. *arXiv preprint arXiv:1903.01068*, 2019.

28 Chaya Keller, Shakhar Smorodinsky, and Gábor Tardos. Improved bounds on the Hadwiger–Debrunner numbers. *Israel Journal of Mathematics*, 225(2):925–945, 2018.

29 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. doi:10.1007/s004930100020.

30 Bernhard Korte, László Lovász, and Rainer Schrader. *Greedoids*, volume 4. Springer, Berlin, 1991. doi:10.1007/978-3-642-58191-5.

31 Jirí Matousek. Range searching with efficient hierarchical cuttings. In David Avis, editor, *Proceedings of the Eighth Annual Symposium on Computational Geometry, Berlin, Germany, June 10-12, 1992*, pages 276–285. ACM, 1992. doi:10.1145/142675.142732.

32 Jiří Matoušek. *Lectures on discrete geometry*, volume 212. New York, NY: Springer, 2002. doi:10.1007/978-1-4613-0039-7.

33 Daniel McGinnis. A family of convex sets in the plane satisfying the $(4,3)$-property can be pierced by nine points. *arXiv preprint arXiv:2010.13195*, 2020. arXiv:2010.13195.

34 Luis Montejano. A new topological helly theorem and some transversal results. *Discret. Comput. Geom.*, 52(2):390–398, 2014. doi:10.1007/s00454-014-9613-6.

35 Luis Montejano and Pablo Soberón. Piercing numbers for balanced and unbalanced families. *Discret. Comput. Geom.*, 45(2):358–364, 2011. doi:10.1007/s00454-010-9295-7.

36 Shay Moran and Amir Yehudayoff. On weak epsilon-nets and the radon number. In Gill Barequet and Yusu Wang, editors, *35th International Symposium on Computational Geometry, SoCG 2019, June 18-21, 2019, Portland, Oregon, USA*, volume 129 of *LIPIcs*, pages 51:1–51:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.SoCG.2019.51.

37 Rafael Jesús Segura and Francisco R. Feito. An algorithm for determining intersection segment-polygon in 3d. *Comput. Graph.*, 22(5):587–592, 1998. doi:10.1016/S0097-8493(98)00064-8.

38 Lajos Stachó. A solution of Gallai's problem on pinning down circles. *Mat. Lapok*, 32:19–47, 1984.

39 Martin Tancer. Intersection patterns of convex sets via simplicial complexes: A survey. In János Pach, editor, *Thirty Essays on Geometric Graph Theory*, pages 521–540. Springer New York, New York, NY, 2013. doi:10.1007/978-1-4614-0110-0_28.

40 Gerd Wegner. d-collapsing and nerves of families of convex sets. *Archiv der Mathematik*, 26(1):317–321, December 1975. doi:10.1007/bf01229745.

41 Gerd Wegner. Über Helly-Gallaische Stichzahlprobleme. *3. Kolloquium über Diskrete Geometrie*, pages 277–282, 1985.

Appendices

A From compact convex sets to general convex sets in the Hadwiger-Debrunner theorem

In the main body of this paper, we have given a proof of the Hadwiger-Debrunner theorem for compact convex sets only, claiming that this was done without loss of generality. To see this let us describe how to reduce the general case to this one.

Let $\mathcal{F} = \{S_1, S_2, \ldots, S_n\}$ be a finite family of convex sets in $\mathbb{R}^d$ with the (p, q)-property for some admissible pair (p, q). Let us construct a new family of compact convex sets $\mathcal{F}' = \{S'_1, S'_2, \ldots, S'_n\}$ as follows:

- For every subfamily $\mathcal{G} \subset \mathcal{F}$ of sets with non-empty common intersection, choose $p_{\mathcal{G}}$ to be a point in that intersection, and let $\mathcal{P}$ denote the set of all such chosen points.
- For $i \in \{1, \ldots, n\}$, let S'_i be the convex hull of $S_i \cap \mathcal{P}$.

It is clear that $\mathcal{F}'$ consists of compact convex sets and that whenever some subfamily $\mathcal{G} \subset \mathcal{F}$ of sets have a common intersection, the corresponding sets in $\mathcal{F}'$ also do. This means that $\mathcal{F}'$ has the (p, q)-property. Moreover, for $i \in \{1, \ldots, n\}$ we know that $S'_i \subset S_i$. If we can stab $\mathcal{F}'$ with $p - q + 1$ points, the same holds for $\mathcal{F}$. Thus, if the Hadwiger-Debrunner theorem holds for compact convex sets, it also holds for general convex sets.

B Adapting the Bentley-Ottmann sweep-line algorithm

Here we briefly describe how to adapt the Bentley-Ottmann sweep-line algorithm [8] to solve the following problem in $\mathcal{O}(n^2 \log(n))$ time.

▶ **Problem 44.** *Given a family of convex polygons in the plane with a total of n vertices, compute a point p stabbing as many polygons as possible.*

Imagine a vertical line sweeping through the plane, stopping each time it reaches the beginning of an edge, the end of an edge or the intersection of two edges, which we call an event (for simplicity, assume the events are separated along the horizontal axis by shifting them infinitesimally). During the whole sweep, we keep track of the edges crossing our sweep line ordered according to the y coordinate of their intersection with the sweep line. This forms the general idea behind the Bentley-Ottmann algorithm.

More specifically, the algorithm maintains a self-balancing Binary Search Tree (BST) of the segments intersecting the sweep line (this line is conceptual only and is not explicitly stored in any manner) as well as a priority queue of events to come. At each stage neighbouring edges on the sweep line are tested for future intersection and this intersection is added to the priority queue if it exists. Then the next event in the priority queue is considered and the imaginary sweep line is moved to that event. If this is the beginning or the end of an edge, this edge is respectively added to or deleted from the BST. If this event is the intersection of two edges, their positions in the BST are swapped. All these operations can be done in $\mathcal{O}(\log(n))$ time with the usual data structures for priority queues and self-balancing BSTs. Thus, as there are at most $\mathcal{O}(n^2)$ events, the whole sweep takes $\mathcal{O}(n^2 \log(n))$ time.

Now, we can partition the edges of the polygons into two classes: those corresponding to the upper hull and those corresponding to the lower hull. We can then augment each node in the BST with the following information: for each node v store the number of leaves which are lower hull edges and upper hull edges in the subtree rooted at v. These quantities can easily be maintained in $\mathcal{O}(\log(n))$ time per operation performed on the BST.

When considering a new event, we can now easily compute the number of polygons stabbed by the corresponding point by taking the number of upper hull edges above it and subtracting the number of lower hull edges above. Both of these latter quantities can be computed in $\mathcal{O}(\log(n))$ time by travelling up the BST from the vertex of interest to the root.

To compute a point p stabbing as many polygons as possible, it is enough to consider only the points which we have defined as events. Thus, with this modified algorithm we can find the event point stabbing the largest number of polygons and thus solve the considered problem in $\mathcal{O}(n^2 \log(n))$ time.

C Proof of Theorem 7

Here we give pointers to prove Theorem 7. The methods used here are standard in the field of geometric range queries and are variations on proofs by Matoušek [31] using results by Chan [12]. Although we recommend the lecture of these papers for a better understanding of these approaches, we give some details here for the sake of completeness.

Recall the statement of the theorem.

▶ **Theorem 7.** *Let S be a set of n objects, k a constant, and $\phi_1, \phi_2, \ldots, \phi_k$ mappings from S to $\mathbb{R}^d$. Let ϕ_S be the function which maps k-tuples of halfspaces $H_1, H_2, \ldots, H_k$ of $\mathbb{R}^d$ to the set*

$$\phi_S(H_1, H_2, \ldots, H_k) := \{s \in S \mid \phi_1(s) \in H_1, \phi_2(s) \in H_2 \ldots, \phi_k(s) \in H_k\}.$$

Suppose we have computed the point sets $\phi_1(S), \ldots, \phi_k(S)$ and let $n \leq m \leq n/\log^{\omega(1)} n$. Then we can preprocess the point sets in $\mathcal{O}(n \log^k n + m)$ time such that $|\phi_S(H_1, H_2, \ldots, H_k)|$ can be computed in $\mathcal{O}((n/m^{1/d})(\log n)^{2(k+(k-d-1)/d)}(\log\log n)^{1/d})$ expected time for any k-tuple of halfspaces.

From now on, d, k and the mappings $\phi_1, \ldots, \phi_k$ are fixed. For the sake of simplicity we suppose that all objects $s \in S$ we consider are of constant size and that $\phi_i(s)$ can be computed in constant time (but we could precompute all these points as a preprocessing step and then work only with the points $\phi_i(s)$). For any finite set of objects S and any $1 \leq k' \leq k'' \leq k$, we let $\phi_S^{k',k''}$ denote the function which maps $(k'' - k' + 1)$-tuples of halfspaces $H_{k'}, \ldots, H_{k''}$ of $\mathbb{R}^d$ to

$$\phi_S^{k',k''}(H_{k'}, \ldots, H_{k''}) := \{s \in S \mid \phi_{k'}(s) \in H_{k'}, \ldots, \phi_{k''}(s) \in H_{k''}\}.$$

When $k' = k''$, we also use the notation $\phi_S^{k'} := \phi_S^{k',k''}$.

We need the following result by Matoušek [31].

▶ **Theorem 45** ([31, Theorem 5.1]). *For any $1 \leq k' \leq k$, any set of n objects S and any parameter $r < n$, we can build a datastructure in $\mathcal{O}(nr^{d-1})$ time with the following properties.*

- *There are $t \in \mathcal{O}(\log r)$ collections of subsets of S, $\mathcal{C}_1, \mathcal{C}_1, \ldots, \mathcal{C}_t$ such that $\mathcal{C}_i$ contains $\mathcal{O}(\rho^i)$ subsets of size at most n/ρ^i (where $\rho > 1$ is a constant dependent on r). We call all the subsets in these collections the inner sets.*
- *There are an additional $\mathcal{O}(r^d)$ subsets of S, each of size at most n/r, called the remainder sets. We denote the collection of these subsets as $\mathcal{R}$.*

- *For any halfspace H, we can in $\mathcal{O}(\log r)$ time return pointers to $t+1$ of these subsets, $S_1 \in \mathcal{C}_1, S_2 \in \mathcal{C}_2, \ldots, S_t \in \mathcal{C}_t$ and $S_\mathcal{R} \in \mathcal{R}$, all disjoint, such that*

$$\phi_S^{k'}(H) = S_1 \cup S_2 \cup \cdots \cup S_t \cup \phi_{S_\mathcal{R}}^{k'}(H).$$

Note that the inner and remainder sets are not stored individually in full but implicitly as a partition tree structure.

We also need this other result by Chan [12].

▶ **Theorem 46** ([12]). *For any $1 \leq k' \leq k$, any set of n objects S, we can build a datastructure in $\mathcal{O}(n \log^{k'} n)$ time which can then compute $|\phi_S^{1,k'}(H_1, \ldots, H_{k'})|$ in $\mathcal{O}(n^{(d-1)/d} \log^{k'-1} n)$ expected time for any k'-tuple of halfspaces.*

We can now prove the following, by adapting a proof by Matoušek [31, Theorem 6.1].

▶ **Theorem 47.** *For any, $1 \leq p \leq k$, any set of n objects S, we can build a datastructure in time $\mathcal{O}(n^d \log^{p-d-1} n \log\log n)$ which can then compute $|\phi_S^{1,p}(H_1, \ldots, H_p)|$ in $\mathcal{O}(\log^p n)$ expected time for any p-tuple of halfspaces.*

Proof. The proof is by induction. Let us consider the base case $p = 1$. We use the datastructure from Theorem 45 with $k' = 1$ and $r = n/\log^{d/(d-1)} n$. To each inner set we add an attribute representing the size of the set. To each remainder set we attach the corresponding datastructure from Theorem 46 with $k' = 1$. We can then compute $|\phi_S^1(H)|$ by using the primary datastructure to find the decomposition into inner sets and remaining set $S_\mathcal{R}$ given by Theorem 45, and use the secondary attached datastructure to compute $|\phi_{S_\mathcal{R}}^1(H)|$. Because there are $\mathcal{O}(r^d)$ remainder sets, all of size at most n/r, the total preprocessing time is $\mathcal{O}(nr^{d-1} + r^d(n/r)\log(n/r)) \subset \mathcal{O}(n^d \log^{1-d-1} n \log\log n)$. Because in each query there is only one remainder set to consider, the total expected query time is $\mathcal{O}(\log r + (n/r)^{(d-1)/d}) \subset \mathcal{O}(\log n)$. Thus the claim holds for $p = 1$.

Let us now suppose it holds for $1 \leq p-1 < k$ and show it holds for p. We again use the datastructure from *Theorem* 45 with $r = n/\log^{d/(d-1)} n$, this time with $k' = p$. To each inner set we attach the corresponding datastructure we inductively suppose exists for $p-1$. To each remainder set we attach the corresponding datastructure from Theorem 46 with $k' = p$. Using the knowledge about the distribution of sizes of the inner sets in Theorem 45, the total preprocessing time will be

$$\begin{aligned}
&\mathcal{O}\left(nr^{d-1} + r^d(n/r)\log^p(n/r) + \sum_{i=1}^{\mathcal{O}(\log r)} \rho^i \left(\frac{n}{\rho^i}\right)^d \log^{p-1-d-1} n \log\log n\right) \\
&\subset \mathcal{O}\left(n^d \log^{-1} n + n^d \log^{p-d-2} n (\log\log n)^{p+1} + n^d \log^{p-d-2} \log\log n \sum_{i=1}^{\mathcal{O}(\log r)} (\rho^{1-d})^i\right) \\
&\subset \mathcal{O}\left(n^d \log^{p-d-1} n \log\log n\right).
\end{aligned}$$

The total expected query time will be

$$\begin{aligned}
&\mathcal{O}\left(\log r + \log r \log^{p-1} n + (n/r)^{(d-1)/d} \log^{p-1}(n/r)\right) \\
&\subset \mathcal{O}\left(\log^p n\right).
\end{aligned}$$

Thus by induction, the claim holds. ◀

For the final proof we need some other results by Chan [12], which we summarize in the following theorem.

▶ **Theorem 48** ([12]). *For any set of n objects S and any $1 \leq B \leq n/\log^{\omega(1)} n$, we can build a datastructure in $\mathcal{O}(n \log^k n)$ time with the following properties.*

- *There are $\mathcal{O}((n/B)\log^{k-1} n)$ subsets of S. We denote their collection $\mathcal{C}$. The subset of $\mathcal{C}$ consisting of sets of size at most B is denoted as $\mathcal{B}$.*
- *For any k-tuple of halfspaces $H_1, \dots, H_k$, we can compute in $\mathcal{O}((n/B)^{(d-1)/d}\log^{k-1} n)$ expected time pointers to $t \in \mathcal{O}((n/B)^{(d-1)/d}\log^{k-1} n)$ sets of $\mathcal{C}$, denoted as $C_1, C_2 \dots C_t$ and another $t' \in \mathcal{O}((n/B)^{(d-1)/d}\log^{k-1} n)$ to sets of $\mathcal{B}$, denoted as $S_1, S_2 \dots S_{t'}$, all disjoint. Moreover they are such that*

$$\phi_S(H_1, \dots, H_k) = (C_1 \cup C_2 \cup \dots \cup C_t) \cup \left(\phi_{S_1}(H_1, \dots, H_k) \cup \dots \cup \phi_{S_{t'}}(H_1, \dots, H_k)\right).$$

We are now ready for the final proof, which is again a minor adaptation of a proof by Matoušek [31, Theorem 6.2].

Proof of Theorem 7. We use the datastructure from Theorem 48, for some unspecified B. To each set in $\mathcal{C}$, we add an attribute representing the size of the set. To each set in $\mathcal{B}$, we also attach the datastucture we get from Theorem 47 with $p = k$. Because each set in $\mathcal{B}$ is of size at most B and there are $\mathcal{O}((n/B)\log^{k-1} n)$ such sets, the total preprocessing time will be $\mathcal{O}(n \log^k n + m)$, where

$$\begin{aligned} m &:= (n/B)\log^{k-1} n + (n/B)(\log^{k-1} n)B^d(\log^{k-d-1} B)\log\log B \\ &\in \mathcal{O}\left(B^{d-1} n (\log^{2k-d-2} n)\log\log n\right). \end{aligned}$$

We can make m vary between $n\log^{k-1} n$ and $n^d/\log^{\omega(1)} n$ by having B vary between 1 and $n/\log^{\omega(1)} n$ (but we can still choose $m < n\log^{k-1} n$ in the statement of the theorem, as the $n\log^k n$ term then dominates the preprocessing). The total expected query time will be

$$\begin{aligned} &\mathcal{O}\left((n/B)^{(d-1)/d} + (n/B)^{(d-1)/d}(\log^{k-1} n)\log^k B\right) \\ &\subset \mathcal{O}\left((n/B)^{(d-1)/d})\log^{k-1} n \log^{2k-1} n\right). \end{aligned}$$

By rewriting this in terms of m instead of B we get an expected query time of

$$\mathcal{O}\left((n/m^{1/d})\log^{2(k+(k-d-1)/d))} n (\log\log n)^{1/d}\right).$$

◀