Stabbing Pairwise Intersecting Disks by Five Points*

Sariel Har-Peled[1], Haim Kaplan[2], Wolfgang Mulzer[3], Liam Roditty[4], Paul Seiferth[3], Micha Sharir[2], and Max Willert[3]

[1]Department of Computer Science, University of Illinois, Urbana, IL 61801, USA, `sariel@illinois.edu`
[2]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel, `{haimk,michas}@tau.ac.il`
[3]Institut für Informatik, Freie Universität Berlin, 14195 Berlin, Germany, `{mulzer,pseiferth,willerma}@inf.fu-berlin.de`
[4]Department of Computer Science, Bar Ilan University, Ramat Gan 5290002, Israel, `liamr@macs.biu.ac.il`

Abstract

Suppose we are given a set $\mathcal{D}$ of n pairwise intersecting disks in the plane. A planar point set P *stabs* $\mathcal{D}$ if and only if each disk in $\mathcal{D}$ contains at least one point from P. We present a deterministic algorithm that takes $O(n)$ time to find five points that stab $\mathcal{D}$. Furthermore, we give a simple example of 13 pairwise intersecting disks that cannot be stabbed by three points. Moreover, we present a simple argument showing that eight disks can be stabbed by at most three points.

This provides a simple—albeit slightly weaker—algorithmic version of a classical result by Danzer that such a set $\mathcal{D}$ can always be stabbed by four points.

1 Introduction

The *maximum clique problem* is a classic problem in combinatorial optimization [15]: given a simple graph $G = (V, E)$, find a maximum-cardinality set $C \subseteq V$ of vertices such that any two distinct vertices in C are adjacent. In 1972, Karp proved that the maximum clique problem is NP-hard [15]. Even worse, a subsequent line of research showed that the maximum clique problem is hard to approximate. In particular, we now know that for any fixed $\varepsilon > 0$, if there is a polynomial-time algorithm that approximates maximum clique in an n-vertex graph up to a factor of $n^{1-\varepsilon}$, then P = NP [22]

However, if the input graph has additional structure, the problem can become easier. For example, if the input is the intersection graph of a set of disks in the plane, the maximum clique problem admits efficient (approximation) algorithms: for unit disk graphs, it can be solved in polynomial time [8], while for general disk intersection graphs, there is a randomized EPTAS [3]. Earlier, Ambühl and Wagner [2] presented a polynomial-time algorithm that computes a $\tau/2$-approximation for the maximum clique in a general disk intersection graph, where τ is the minimum *stabbing number* of any arrangement of pairwise intersecting disks in the plane, i.e., the minimum number of points that are needed to stab every disk in such an arrangement. Motivated by this application, our goal here is to understand this stabbing number better.

Let $\mathcal{D}$ be a set of n disks in the plane. If every *three* disks in $\mathcal{D}$ intersect, then Helly's theorem shows that the whole intersection $\bigcap \mathcal{D}$ of $\mathcal{D}$ is nonempty [13, 14, 17]. In other words, there is a single point p that lies in all disks of $\mathcal{D}$, that is, p *stabs* $\mathcal{D}$. More generally, when we know only that every *pair* of disks

*A preliminary version appeared as S. Har-Peled, H. Kaplan, W. Mulzer, L. Roditty, P. Seiferth, M. Sharir, and M. Willert. *Stabbing Pairwise Intersecting Disks by Five Points*. Proc. 29th ISAAC, pp. 50:1–50:12. SHP was supported by a NSF AF awards CCF-1421231, and CCF-1217462. WM was supported by DFG grant MU/3501/1 and ERC STG 757609. PS was supported by DFG grant MU/3501/1. MS was supported by ISF grant 892/13 and 260/18, by the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11), and by the Blavatnik Research Fund in Computer Science at Tel Aviv University. HK was supported by ISF grant 1595-19 and the Blavatnik Family Foundation. Work on this paper was supported in part by grant 1367/2016 and 1161/2011 from the German-Israeli Science Foundation (GIF).

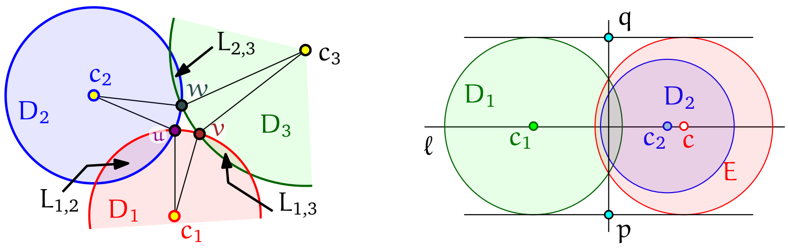

Figure 1: Left: At least one lens angle is large. Right: D_1 and E have the same radii and lens angle $2\pi/3$. By Lemma 2.2, D_2 is a subset of E. $\{c_1, c, p, q\}$ is the set P from Lemma 2.4.

in $\mathcal{D}$ intersect, there must be a point set P of constant size such that each disk in $\mathcal{D}$ contains at least one point in P – the minimum cardinality of P is the *stabbing number* of $\mathcal{D}$. It is indeed not surprising that $\mathcal{D}$ can be stabbed by a constant number of points, but for some time, the exact bound remained elusive. Eventually, in July 1956 at an Oberwolfach seminar, Danzer presented the answer: four points are always sufficient and sometimes necessary to stab any finite set of pairwise intersecting disks in the plane. Danzer was not satisfied with his original argument, so he never formally published it. In 1986, he presented a new proof [9]. Previously, in 1981, Stachó had already given an alternative proof [21], building on a previous construction of five stabbing points [20]. This line of work was motivated by a result of Hadwiger and Debrunner, who showed that three points suffice to stab any finite set of pairwise intersecting *unit* disks [12]. In later work, these results were significantly generalized and extended, culminating in the celebrated (p, q)-theorem that was proven by Alon and Kleitman in 1992 [1]. See also a recent paper by Dumitrescu and Jiang that studies generalizations of the stabbing problem for translates and homothets of a convex body [10].

Danzer's published proof [9] is fairly involved. It uses a compactness argument that does not seem to be constructive, and one part of the argument relies on an underspecified verification by computer. Therefore, it is quite challenging to check the correctness of the argument, let alone to derive any intuition from it. There seems to be no obvious way to turn it into an efficient algorithm for finding a stabbing set of size four. The proof of Stachó [21] is simpler, but it is obtained through a lengthy case analysis that requires a very disciplined and focused reader. Here, we present a new argument that yields five stabbing points. Our proof is constructive, and it lets us find the stabbing set in deterministic linear time. Following the conference version of this paper, Carmi, Katz, and Morin published a manuscript in which they present an algorithm that can find four stabbing points in linear time [4].

As for lower bounds, Grünbaum gave an example of 21 pairwise intersecting disks that cannot be stabbed by three points [11]. Later, Danzer reduced the number of disks to ten [9]. This example is close to optimal, because every set of eight disks can be stabbed by three points, as mentioned by Stachó [20] and formally proved in Section 5 below. However, it is hard to verify Danzer's lower bound example—even with dynamic geometry software, the positions of the disks cannot be visualized easily.

We present a new and simple proof that shows that the stabbing number of $\mathcal{D}$ is upper bounded by 5. Moreover, we obtain a linear time algorithm that can find these 5 stabbing points. Finally, we present a simple construction of 13 pairwise intersecting disks that cannot be stabbed by 3 points, and work out a proof of Stachó's eight-disk claim.

2 The Geometry of Pairwise Intersecting Disks

Let $\mathcal{D}$ be a set of n pairwise intersecting disks in the plane. A disk $D_i \in \mathcal{D}$ is given by its center c_i and its radius r_i. To simplify the analysis, we make the following assumptions: (i) the radii of the disks are pairwise distinct; (ii) the intersection of any two disks has a nonempty interior; and (iii) the intersection of any three disks is either empty or has a nonempty interior. A simple perturbation argument can then handle the degenerate cases.

The *lens* of two disks $D_i, D_j \in \mathcal{D}$ is the set $L_{i,j} = D_i \cap D_j$. Let u be any of the two intersection points of the boundary of D_i and the boundary of D_j. The angle $\angle c_i u c_j$ is called the *lens angle* of D_i and D_j. It is at most π. A finite set $\mathcal{C}$ of disks is *Helly* if their common intersection $\bigcap \mathcal{C}$ is nonempty. Otherwise, $\mathcal{C}$ is *non-Helly*. We present some useful geometric lemmas.

Lemma 2.1. *Let $\{D_1, D_2, D_3\}$ be a set of three pairwise intersecting disks that is non-Helly. Then, the set contains two disks with lens angle larger than $2\pi/3$.*

Proof. Since $\{D_1, D_2, D_3\}$ is non-Helly, the lenses $L_{1,2}$, $L_{1,3}$ and $L_{2,3}$ are pairwise disjoint. Let u be the vertex of $L_{1,2}$ nearer to D_3, and let v, w be the analogous vertices of $L_{1,3}$ and $L_{2,3}$ (see Figure 1, left). Consider the simple hexagon $c_1 u c_2 w c_3 v$, and write $\angle u$, $\angle v$, and $\angle w$ for its interior angles at u, v, and w. The sum of all interior angles is 4π. Thus, $\angle u + \angle v + \angle w < 4\pi$, so at least one angle is less than $4\pi/3$. It follows that the corresponding lens angle, which is the exterior angle at u, v, or w must be larger than $2\pi/3$. $\square$

Lemma 2.2. *Let D_1 and D_2 be two intersecting disks with $r_1 \geq r_2$ and lens angle at least $2\pi/3$. Let E be the unique disk with radius r_1 and center c, such that*

- *(i) the centers c_1, c_2, and c are collinear and c lies on the same side of c_1 as c_2; and*
- *(ii) the lens angle of D_1 and E is exactly $2\pi/3$ (see Figure 1, right).*

Then, if c_2 lies between c_1 and c, we have $D_2 \subseteq E$.

Proof. Let $x \in D_2$. Since c_2 lies between c_1 and c, the triangle inequality gives

$$|xc| \leq |xc_2| + |c_2 c| = |xc_2| + |c_1 c| - |c_1 c_2|. \tag{1}$$

Since $x \in D_2$, we get $|xc_2| \leq r_2$. Also, since D_1 and E have radius r_1 each and lens angle $2\pi/3$, it follows that $|c_1 c| = \sqrt{3}\, r_1$. Finally, $|c_1 c_2| = \sqrt{r_1^2 + r_2^2 - 2r_1 r_2 \cos\alpha}$, by the law of cosines, where α is the lens angle of D_1 and D_2. As $\alpha \geq 2\pi/3$ and $r_1 \geq r_2$, we get $\cos\alpha \leq -1/2 = (\sqrt{3} - 3/2) - \sqrt{3} + 1 \leq (\sqrt{3} - 3/2) r_1/r_2 - \sqrt{3} + 1$, As such, we have

$$\begin{aligned} |c_1 c_2|^2 &= r_1^2 + r_2^2 - 2r_1 r_2 \cos\alpha \geq r_1^2 + r_2^2 - 2r_1 r_2 \Big((\sqrt{3} - 3/2)\frac{r_1}{r_2} - \sqrt{3} + 1 \Big) \\ &= r_1^2 - 2\big(\sqrt{3} - 3/2\big) r_1^2 + 2(-\sqrt{3} + 1) r_1 r_2 + r_2^2 \\ &= (1 - 2\sqrt{3} + 3) r_1^2 + 2(-\sqrt{3} + 1) r_1 r_2 + r_2^2 = \big(r_1(\sqrt{3} - 1) + r_2\big)^2. \end{aligned}$$

Plugging this into Equation 1 gives $|xc| \leq r_2 + \sqrt{3} r_1 - \left(r_1\left(\sqrt{3} - 1\right) + r_2\right) = r_1$, i.e., $x \in E$. $\square$

Lemma 2.3. *Let D_1 and D_2 be two intersecting disks with equal radius r and lens angle $2\pi/3$. There is a set P of four points so that any disk F of radius at least r that intersects both D_1 and D_2 contains a point of P.*

Proof. Consider the two tangent lines of D_1 and D_2, and let p and q be the midpoints on these lines between the respective two tangency points. We set $P = \{c_1, c_2, p, q\}$; see Figure 2.

Given the disk F that intersects both D_1 and D_2, we shrink its radius, keeping its center fixed, until either the radius becomes r or until F is tangent to D_1 or D_2. Suppose the latter case holds and F is tangent to D_1. We move the center of F continuously along the line spanned by the center of F and c_1 towards c_1, decreasing the radius of F to maintain the tangency. We stop when either the radius of F reaches r or F becomes tangent to D_2. We obtain a disk $G \subseteq F$ with center $c = (c_x, c_y)$ so that either: (i) radius$(G) = r$ and G intersects both D_1 and D_2; or (ii) radius$(G) \geq r$ and G is tangent to both D_1 and D_2. Since $G \subseteq F$, it suffices to show that $G \cap P \neq \emptyset$.

We introduce a coordinate system, setting the origin o midway between c_1 and c_2, so that the y-axis passes through p and q. Then, as in Figure 2, we have $c_1 = (-\sqrt{3}\, r/2, 0)$, $c_2 = (\sqrt{3}\, r/2, 0)$, $q = (0, r)$, and $p = (0, -r)$.

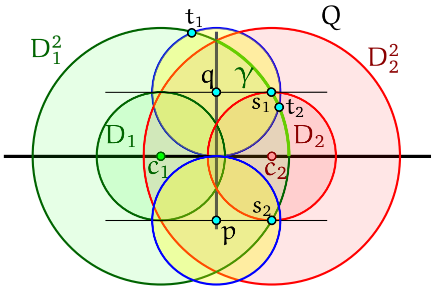

Figure 2: Left: $P = \{c_1, c_2, p, q\}$ is the stabbing set. The green arc $\gamma = \partial D_1^2 \cap Q$ is covered by $D_2 \cup D_q$.

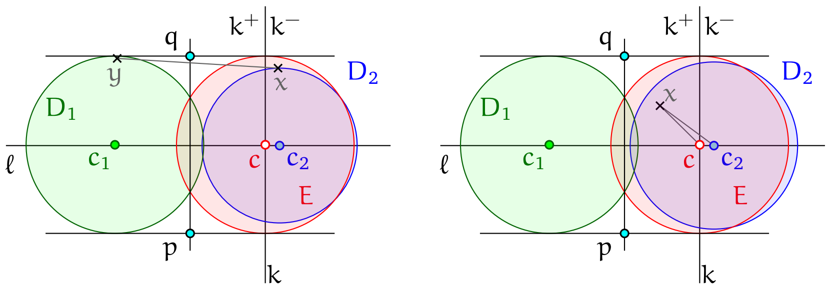

Figure 3: Proof of Lemma 2.4. Left (Case (i)): x is an arbitrary point in $D_2 \cap F \setminus k^+$ and y is an arbitrary point in $D_1 \cap F$. Right (Case (ii)): x is an arbitrary point in $D_2 \cap F \cap k^+$. The angle at c in the triangle Δxcc_2 is $\geq \pi/2$.

For case (i), let D_1^2 be the disk of radius $2r$ centered at c_1, and D_2^2 the disk of radius $2r$ centered at c_2. Since G has radius r and intersects both D_1 and D_2, its center c has distance at most $2r$ from both c_1 and c_2, i.e., $c \in D_1^2 \cap D_2^2$. Let D_p and D_q be the two disks of radius r centered at p and q. We will show that $D_1^2 \cap D_2^2 \subseteq D_1 \cup D_2 \cup D_p \cup D_q$. Then it is immediate that $G \cap P \neq \emptyset$. By symmetry, it is enough to focus on the upper-right quadrant $Q = \{(x, y) \mid x \geq 0, y \geq 0\}$. We show that all points in $D_1^2 \cap Q$ are covered by $D_2 \cup D_q$. Without loss of generality, we assume that $r = 1$. Then, the two intersection points of D_1^2 and D_q are $t_1 = (\frac{5\sqrt{3}-2\sqrt{87}}{28}, \frac{38+3\sqrt{29}}{28}) \approx (-0.36, 1.93)$ and $t_2 = (\frac{5\sqrt{3}+2\sqrt{87}}{28}, \frac{38-3\sqrt{29}}{28}) \approx (0.98, 0.78)$, and the two intersection points of D_1^2 and D_2 are $s_1 = (\frac{\sqrt{3}}{2}, 1) \approx (0.87, 1)$ and $s_2 = (\frac{\sqrt{3}}{2}, -1) \approx (0.87, -1)$. Let γ be the boundary curve of D_1^2 in Q. Since $t_1, s_2 \notin Q$ and since $t_2 \in D_2$ and $s_1 \in D_q$, it follows that γ does not intersect the boundary of $D_2 \cup D_q$ and hence $\gamma \subset D_2 \cup D_q$. Furthermore, the subsegment of the y-axis from o to the start point of γ is contained in D_q, and the subsegment of the x-axis from o to the endpoint of γ is contained in D_2. Hence, the boundary of $D_1^2 \cap Q$ lies completely in $D_2 \cup D_q$, and since $D_2 \cup D_q$ is simply connected, it follows that $D_1^2 \cap Q \subseteq D_2 \cup D_q$, as desired.

For case (ii), since G is tangent to D_1 and D_2, the center c of G is on the perpendicular bisector of c_1 and c_2, so the points p, o, q and c are collinear. Suppose without loss of generality that $c_y \geq 0$. Then, it is easily checked that c lies above q, and radius$(G) + r = |c_1c| \geq |oc| = r + |qc|$, so $q \in G$. □

Lemma 2.4. *Consider two intersecting disks D_1 and D_2 with $r_1 \geq r_2$ and lens angle at least $2\pi/3$. Then, there is a set P of four points such that any disk F of radius at least r_1 that intersects both D_1 and D_2 contains a point of P.*

Proof. Let ℓ be the line through c_1 and c_2. Let E be the disk of radius r_1 and center $c \in \ell$ that satisfies the conditions (i) and (ii) of Lemma 2.2. Let $P = \{c_1, c, p, q\}$ as in the proof of Lemma 2.3, with respect to D_1 and E (see Figure 1, right). We claim that

$$D_1 \cap F \neq \emptyset \ \wedge \ D_2 \cap F \neq \emptyset \ \Rightarrow \ E \cap F \neq \emptyset. \qquad (*)$$

Once (*) is established, we are done by Lemma 2.3. If $D_2 \subseteq E$, then (*) is immediate, so assume that $D_2 \nsubseteq E$. By Lemma 2.2, c lies between c_1 and c_2. Let k be the line through c perpendicular to ℓ, and let k^+ be the open halfplane bounded by k with $c_1 \in k^+$ and k^- the open halfplane bounded by k with $c_1 \notin k^-$. Since $|c_1c| = \sqrt{3}\, r_1 > r_1$, we have $D_1 \subset k^+$; see Figure 3. Recall that F has radius at least r_1 and intersects D_1 and D_2. We distinguish two cases: (i) there is no intersection of F and D_2 in k^+, and (ii) there is an intersection of F and D_2 in k^+; see Figure 3 for the two cases.

For case (i), let x be any point in $D_1 \cap F$. Since we know that $D_1 \subset k^+$, we have $x \in k^+$. Moreover, let y be any point in $D_2 \cap F$. By assumption, y is not in k^+, but it must be in the infinite strip defined by the two tangents of D_1 and E. Thus, the line segment $\overline{xy}$ intersects the diameter segment $k \cap E$. Since F is convex, the intersection of $\overline{xy}$ and $k \cap E$ is in F, so $E \cap F \neq \emptyset$.

For case (ii), fix $x \in D_2 \cap F \cap k^+$ arbitrarily. Consider the triangle Δxcc_2. Since $x \in k^+$, the angle at c is at least $\pi/2$. Thus, $|xc| \leq |xc_2|$. Also, since $x \in D_2$, we know that $|xc_2| \leq r_2 \leq r_1$. Hence, $|xc| \leq r_1$, so $x \in E$ and (*) follows, as $x \in E \cap F$. □

3 Existence of Five Stabbing Points

With these tools we can now show that there is a stabbing set with five points.

Theorem 3.1. *Let $\mathcal{D}$ be a set of n pairwise intersecting disks in the plane. There is a set P of five points such that each disk in $\mathcal{D}$ contains at least one point from P.*

Proof. If $\mathcal{D}$ is Helly, there is a single point that lies in all disks of $\mathcal{D}$. Thus, assume that $\mathcal{D}$ is non-Helly, and let $D_1, D_2, \ldots, D_n$ be the disks in $\mathcal{D}$ ordered by increasing radius. Let i^* be the smallest index with $\bigcap_{i \leq i^*} D_i = \emptyset$. By Helly's theorem [13, 14, 17], there are indices $j, k < i^*$ such that $\{D_{i^*}, D_j, D_k\}$ is non-Helly. By Lemma 2.1, two disks in $\{D_{i^*}, D_j, D_k\}$ have lens angle at least $2\pi/3$. Applying Lemma 2.4 to these two disks, we obtain a set P' of four points so that every disk D_i with $i \geq i^*$ contains at least one point from P'. Furthermore, by definition of i^*, we have $\bigcap_{i<i^*} D_i \neq \emptyset$, so there is a point q that stabs every disk D_i with $i < i^*$. Thus, $P = P' \cup \{q\}$ is a set of five points that stabs every disk in $\mathcal{D}$, as desired. □

Remark. A weakness in our proof is that it combines two different stages, one of finding the point q that stabs all the small disks, and one of constructing the four points of Lemma 2.4 that stab all the larger disks. It is an intriguing challenge to merge the two arguments so that altogether they only require four points. The proof of Carmi et al. [4] uses a different approach.

4 Algorithmic Considerations

The proof of Theorem 3.1 leads to a simple $O(n \log n)$ time algorithm for finding a stabbing set of size five. For this, we need an oracle that decides whether a given set of disks is Helly. This has already been done by Löffler and van Kreveld [16], in a more general context:

Lemma 4.1 (Theorem 6 in [16])**.** *Given a set of n disks, the problem of choosing a point in each disk such that the smallest enclosing circle of the resulting point set has minimum radius can be solved in $O(n)$ deterministic time.*

Now, an $O(n \log n)$-time algorithm for finding the five stabbing points is based on the analysis in the proof of Theorem 3.1. It works as follows: first, we sort the disks in $\mathcal{D}$ by increasing radius. This takes $O(n \log n)$ time. Let $\mathcal{D} = \langle D_1, \ldots, D_n \rangle$ be the resulting order. Next, we use binary search with the oracle from Lemma 4.1 to determine the smallest index i^* such that the prefix $\{D_1, \ldots, D_{i^*}\}$ is non-Helly. This yields the disk D_{i^*}. We have to invoke the oracle $O(\log n)$ times, which gives a total time of $O(n \log n)$ for this step. After that, we use another binary search with the oracle from Lemma 4.1 to determine the smallest index $k < i^*$ such that $\{D_{i^*}, D_1, \ldots, D_k\}$ is non-Helly. This costs $O(n \log n)$ time as well. Then, we perform a linear search to find an index $j < k$ such that $\{D_j, D_k, D_{i^*}\}$ is a non-Helly triple. This step works in $O(n)$ time. Finally, we use Lemma 4.1 to obtain in $O(n)$ time a stabbing point q for the Helly set $\{D_1, \dots, D_{i^*-1}\}$ and the method from the proof of Theorem 3.1 to extend q to a stabbing set for the whole set $\mathcal{D}$. This last step works in $O(1)$ time since the result depends solely on $\{D_j, D_k, D_{i^*}\}$. Hence, we can state our claimed theorem.

Theorem 4.2. *Given a set $\mathcal{D}$ of n pairwise intersecting disks in the plane, we can find in $O(n \log n)$ time a set P of five points such that every disk of $\mathcal{D}$ contains at least one point of P.*

The proof of Lemma 4.1 uses the *LP-type framework* by Sharir and Welzl [6,19]. As we will see next, a more sophisticated application of the framework directly leads to a deterministic linear time algorithm to find a stabbing set with five points.

The LP-type framework. An *LP-type problem* $(\mathcal{H}, w, \leq)$ is an abstract generalization of a low-dimensional linear program. It consists of a finite set of *constraints* $\mathcal{H}$, a *weight function* $w : 2^{\mathcal{H}} \to \mathcal{W}$, and a *total order* $(\mathcal{W}, \leq)$ on the weights. The weight function w assigns a weight to each subset of constraints. It must fulfill the following two axioms:

- **Monotonicity**: for any $\mathcal{H}' \subseteq \mathcal{H}$ and $H \in \mathcal{H}$, we have $w\big(\mathcal{H}' \cup \{H\}\big) \leq w(\mathcal{H}')$;
- **Locality:** for any $\mathcal{B} \subseteq \mathcal{H}' \subseteq \mathcal{H}$ with $w(\mathcal{B}) = w(\mathcal{H}')$ and for any $H \in \mathcal{H}$, we have that if $w\big(\mathcal{B} \cup \{H\}\big) = w(\mathcal{B})$, then also $w\big(\mathcal{H}' \cup \{H\}\big) = w(\mathcal{H}')$.

Given a subset $\mathcal{H}' \subseteq \mathcal{H}$, a *basis* for $\mathcal{H}'$ is an inclusion-minimal set $\mathcal{B} \subseteq \mathcal{H}'$ with $w(\mathcal{B}) = w(\mathcal{H}')$. The *combinatorial dimension* of $(\mathcal{H}, w, \leq)$ is the maximum size of any basis of any subset of H. The goal in an LP-type problem is to determine $w(\mathcal{H})$ and a corresponding basis $\mathcal{B}$ for $\mathcal{H}$. Next, given a set $\mathcal{B} \subseteq \mathcal{H}$ and a constraint $H \in \mathcal{H}$, we say that H *violates* $\mathcal{B}$ if $w\big(\mathcal{B} \cup \{H\}\big) < w(\mathcal{B})$.

A generalization of Seidel's algorithm for low-dimensional linear programming [18,19] shows that we can solve an LP-type problem in $O(|\mathcal{H}|)$ expected time, provided that a constant time algorithm for the following problem is available. Here and below, the constant factor in the O-notation may depend on the combinatorial dimension.

- **Violation test:** Given a basis $\mathcal{B}$ and a constraint $H \in \mathcal{H}$, determine whether H violates $\mathcal{B}$ and return an error message if $\mathcal{B}$ is not a basis for any $\mathcal{H}' \subseteq \mathcal{H}$.[1]

For a deterministic solution, we need an additional computational assumption. Let $\mathcal{B} \subseteq \mathcal{H}$ be a basis of any subset $\mathcal{H}' \subseteq \mathcal{H}$, we use $\text{vio}(\mathcal{B})$ to denote the set of all constraints $H \in \mathcal{H}$ that violate $\mathcal{B}$, i.e., that have $w(\mathcal{B} \cup \{H\}) < w(\mathcal{B})$. Consider the *range space* $(\mathcal{H}, \mathcal{R} = \{\text{vio}(\mathcal{B}) \mid \mathcal{B} \text{ is a basis for some } \mathcal{H}' \subseteq \mathcal{H}\})$. For a subset $\mathcal{Y} \subseteq \mathcal{H}$, we let $(\mathcal{Y}, \mathcal{R}_{\mathcal{Y}})$ be the *induced range space*, that is, $\mathcal{R}_{\mathcal{Y}} = \{\mathcal{Y} \cap R \mid R \in \mathcal{R}\}$. Chazelle and Matoušek [7] have shown that an LP-type problem can be solved in $O(|\mathcal{H}|)$ *deterministic* time if there is a constant-time violation test as stated above and the following computational assumption holds:

- **Oracle:** Given a subset $\mathcal{Y} \subseteq \mathcal{H}$, we can compute some superset $\mathcal{R}' \supseteq \mathcal{R}_{\mathcal{Y}}$ in time $|\mathcal{Y}|^{O(1)}$.

During the following discussion, we will show that the problem of finding a non-Helly triple as in Theorem 3.1 is LP-type and fulfills the four requirements for the algorithm of Chazelle and Matoušek.

Remark. Löffler and van Kreveld provide proofs that the underlying problem in Lemma 4.1 is of LP-type, but they do not give arguments for the two computational assumptions, see [16]. However, it is not difficult to also verify the two missing statements.

[1] Here, we follow the presentation of Chazelle and Matoušek [7]. Sharir and Welzl [19] use a violation test without the error message. Instead, they need an additional *basis computation* primitive: given a basis $\mathcal{B}$ and a constraint $H \in \mathcal{H}$, find a basis for $\mathcal{B} \cup \{H\}$. If a violation test with error message exists and if the combinatorial dimension is a constant, a basis computation primitive can easily be implemented by brute-force enumeration.

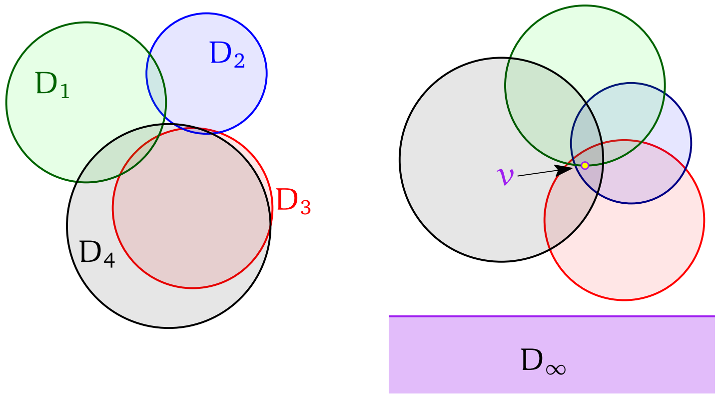

Figure 4: Left: The disks D_3 and D_4 are destroyers of the Helly set $\{D_1, D_2\}$. Moreover, D_3 is the smallest destroyer of the whole set $\{D_1, D_2, D_3, D_4\}$. Right: The disks without D_∞ form a Helly set $\mathcal{C}$. The smallest destroyer of $\mathcal{C}$ is D_∞ and the point v is the extreme point for $\mathcal{C}$ and D_∞, i.e., $\mathrm{dist}(\mathcal{C}) = d(v, D_\infty)$.

Geometric observations. The *distance* between two closed sets $A, B \subseteq \mathbb{R}^2$ is defined as $d(A, B) = \min\{|ab| \mid a \in A, b \in B\}$. From now on, we assume that all points in $\bigcup \mathcal{D}$ have positive y-coordinates. This can be ensured with linear overhead by an appropriate translation of the input. We denote by D_∞ the closed halfplane below the x-axis. It is interpreted as a disk with radius ∞ and center at $(0, -\infty)$. First, observe that for any subsets $\mathcal{C}_1 \subseteq \mathcal{C}_2 \subseteq \mathcal{D} \cup \{D_\infty\}$, we have that if $\mathcal{C}_1$ is non-Helly, then $\mathcal{C}_2$ is non-Helly. For any $\mathcal{C} \subseteq \mathcal{D} \cup \{D_\infty\}$, we say that a disk D *destroys* $\mathcal{C}$ if $\mathcal{C} \cup \{D\}$ is non-Helly. Observe that D_∞ destroys every non-empty subset of $\mathcal{D}$. Moreover, if $\mathcal{C}$ is non-Helly, then every disk is a destroyer. See Figure 4 for an example. We can make the following two observations.

Lemma 4.3. *Let $\mathcal{C} \subseteq \mathcal{D}$ be Helly and D a destroyer of $\mathcal{C}$. Then, the point $v \in \bigcap \mathcal{C}$ with minimum distance to D is unique.*

Proof. Suppose there are two distinct points $v \neq w \in \bigcap \mathcal{C}$ with $d(v, D) = d(\bigcap \mathcal{C}, D) = d(w, D)$. Since $\bigcap \mathcal{C}$ is convex, the segment $\overline{vw}$ lies in $\bigcap \mathcal{C}$. Now, if $D \neq D_\infty$, then every point in the relative interior of $\overline{vw}$ is strictly closer to D than v and w. If $D = D_\infty$, then all points in $\overline{vw}$ have the same distance to D, but since $\bigcap \mathcal{C}$ is strictly convex, the relative interior of $\overline{vw}$ lies in the interior of $\bigcap \mathcal{C}$, so there must be a point in $\bigcap \mathcal{C}$ that is closer to D than v and w. In either case, we obtain a contradiction to the assumption $v \neq w$ and $d(v, D) = d(\bigcap \mathcal{C}, D) = d(w, D)$. The claim follows. □

Let $\mathcal{C} \subseteq \mathcal{D}$ be Helly and D a destroyer of $\mathcal{C}$. The unique point $v \in \bigcap \mathcal{C}$ with minimum distance to D is called the *extreme point* for $\mathcal{C}$ and D (see Figure 4, right).

Lemma 4.4. *Let $\mathcal{C}_1 \subseteq \mathcal{C}_2 \subseteq \mathcal{D}$ be two Helly sets and D a destroyer of $\mathcal{C}_1$ (and thus of $\mathcal{C}_2$). Let $v \in \bigcap \mathcal{C}_1$ be the extreme point for $\mathcal{C}_1$ and D. We have $d(\bigcap \mathcal{C}_1, D) \leq d(\bigcap \mathcal{C}_2, D)$. In particular, if $v \in \bigcap \mathcal{C}_2$, then $d(\bigcap \mathcal{C}_1, D) = d(\bigcap \mathcal{C}_2, D)$ and v is also the extreme point for $\mathcal{C}_2$ and D. If $v \notin \bigcap \mathcal{C}_2$, then $d(\bigcap \mathcal{C}_1, D) < d(\bigcap \mathcal{C}_2, D)$.*

Proof. The first claim holds trivially: let $w \in \bigcap \mathcal{C}_2$ be the extreme point for $\mathcal{C}_2$ and D. Since $\mathcal{C}_1 \subseteq \mathcal{C}_2$, it follows that $w \in \bigcap \mathcal{C}_1$, so $d(\bigcap \mathcal{C}_1, D) \leq d(w, D) = d(\bigcap \mathcal{C}_2, D)$. If $v \in \bigcap \mathcal{C}_2$, then $d(\bigcap \mathcal{C}_1, D) \leq d(\bigcap \mathcal{C}_2, D) \leq d(v, D) = d(\bigcap \mathcal{C}_1, D)$, so $v = w$, by Lemma 4.3. If $v \notin \bigcap \mathcal{C}_2$, then $d(\bigcap \mathcal{C}_1, D) < d(\bigcap \mathcal{C}_2, D)$, by Lemma 4.3 and the fact that $\mathcal{C}_1 \subseteq \mathcal{C}_2$. See Figure 5. □

Let $\mathcal{C}$ be a subset of $\mathcal{D}$. For $0 < r \leq \infty$ we define $\mathcal{C}_{<r}$ as the set of all disks in $\mathcal{C}$ with radius smaller than r. Recall that we assume that all the radii are pairwise distinct. A disk D with radius r, $0 < r \leq \infty$, is called *smallest destroyer* of $\mathcal{C}$ if (i) $D \in \mathcal{C}$ or $D = D_\infty$, (ii) D destroys $\mathcal{C}_{<r}$, and (iii) there is no disk

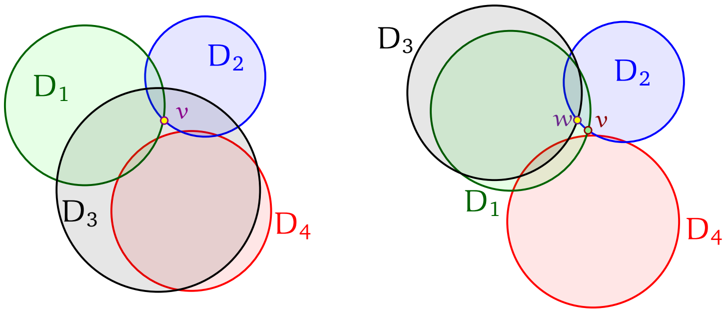

Figure 5: Left: The disk D_4 is a destroyer for the Helly sets $\{D_1, D_2\}$ and $\{D_1, D_2, D_3\}$. The extreme point v for $\{D_1, D_2\}$ is also the extreme point for $\{D_1, D_2, D_3\}$. Right: The disk D_4 is a destroyer for the Helly sets $\{D_1, D_2\}$ and $\{D_1, D_2, D_3\}$. The extreme point v for $\{D_1, D_2\}$ is not in D_3. The distance to D_4 increases.

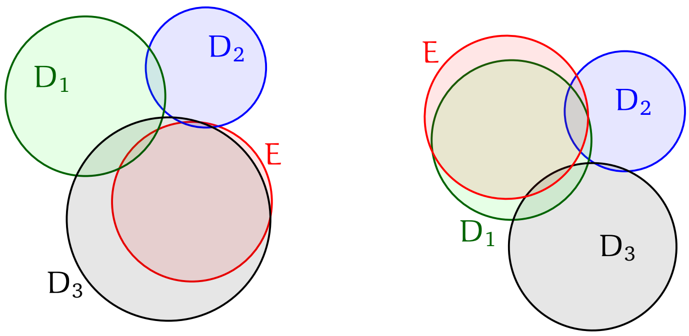

Figure 6: Monotonicity: In both cases, $\{D_1, D_2, D_3\}$ is non-Helly with smallest destroyer D_3. Adding a disk E either decreases the radius of the smallest destroyer (left) or increases the distance to the smallest destroyer (right).

$D' \in \mathcal{C}_{<r}$ that destroys $\mathcal{C}_{<r}$. Observe that Property (iii) is the same as saying that $\mathcal{C}_{<r}$ is Helly. See Figure 4 for an example.

Let $\mathcal{C}$ be a subset of $\mathcal{D}$ and D the smallest destroyer of $\mathcal{C}$. We write $\operatorname{rad}(\mathcal{C})$ for the radius of D and $\operatorname{dist}(\mathcal{C})$ for the distance between D and the set $\bigcap \mathcal{C}_{<\operatorname{rad}(\mathcal{C})}$, i.e., $\operatorname{dist}(\mathcal{C}) = d\big(\bigcap \mathcal{C}_{<\operatorname{rad}(\mathcal{C})}, D\big)$. Now, if $\mathcal{C}$ is Helly, then $D = D_\infty$ and thus $\operatorname{rad}(\mathcal{C}) = \infty$. If $\mathcal{C}$ is non-Helly, then $D \in \mathcal{C}$ and thus $\operatorname{rad}(\mathcal{C}) < \infty$. In both cases, $\operatorname{dist}(\mathcal{C})$ is the distance between D and the extreme point for $\mathcal{C}_{<\operatorname{rad}(\mathcal{C})}$ and D. We define the *weight of* $\mathcal{C}$ as $w(\mathcal{C}) = (\operatorname{rad}(\mathcal{C}), -\operatorname{dist}(\mathcal{C}))$, and we denote by $\leq$ the lexicographic order on $\mathbb{R}^2$. Chan observed, in a slightly different context, that $(\mathcal{D}, w, \leq)$ is LP-type [5]. However, Chan's paper does not contain a detailed proof for this fact. Thus, in the following lemmas, we show the two LP-type axioms, present a constant time violation test, and a polynomial-time oracle. We start with the monotonicity axiom followed by the locality axiom.

Lemma 4.5. *For any $\mathcal{C} \subseteq \mathcal{D}$ and $E \in \mathcal{D}$, we have $w\big(\mathcal{C} \cup \{E\}\big) \leq w(\mathcal{C})$.*

Proof. Set $\mathcal{C}^* = \mathcal{C} \cup \{E\}$. Let D be the smallest destroyer of $\mathcal{C}$, and let $r = \operatorname{rad}(\mathcal{C})$ be the radius of D. Since D destroys $\mathcal{C}_{<r}$, the set $\mathcal{C}_{<r} \cup \{D\}$ is non-Helly. Moreover, since $\mathcal{C}_{<r} \cup \{D\} \subseteq \mathcal{C}^*_{<r} \cup \{D\}$, we know that $\mathcal{C}^*_{<r} \cup \{D\}$ is also non-Helly. Therefore, D destroys $\mathcal{C}^*_{<r}$ and we can derive $\operatorname{rad}(\mathcal{C}^*) \leq \operatorname{rad}(\mathcal{C})$. If we have $\operatorname{rad}(\mathcal{C}^*) < \operatorname{rad}(\mathcal{C})$, we are done. Hence, assume that $\operatorname{rad}(\mathcal{C}^*) = \operatorname{rad}(\mathcal{C})$. Then D is the smallest destroyer of $\mathcal{C}^*$, and Lemma 4.4 gives $-\operatorname{dist}(\mathcal{C}^*) = -d\big(\bigcap \mathcal{C}^*_{<r}, D\big) \leq -d\big(\bigcap \mathcal{C}_{<r}, D\big) = -\operatorname{dist}(\mathcal{C})$. Hence, $w\big(\mathcal{C}^*\big) \leq w(\mathcal{C})$. See Figure 6 for an illustration. $\square$

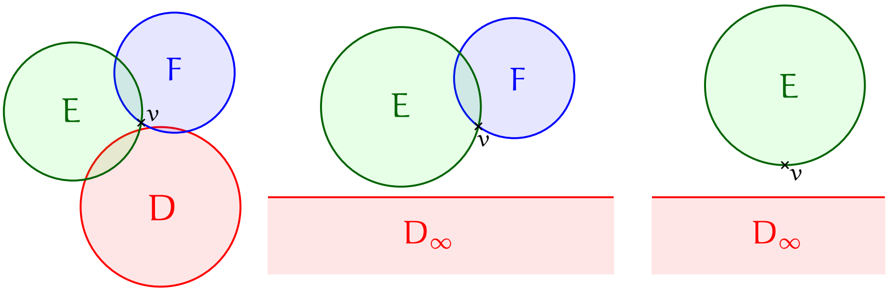

Figure 7: A basis can either be a non-Helly triple (left), a pair of intersecting disks E and F where the point of minimum y-coordinate in $E \cap F$ is a vertex (middle), or a single disk (right).

Lemma 4.6. *Let $\mathcal{B} \subseteq \mathcal{C} \subseteq \mathcal{D}$ with $w(\mathcal{B}) = w(\mathcal{C})$ and let $E \in \mathcal{D}$. Then, if $w\big(\mathcal{B} \cup \{E\}\big) = w(\mathcal{B})$, we also have $w\big(\mathcal{C} \cup \{E\}\big) = w(\mathcal{C})$.*

Proof. Set $\mathcal{C}^* = \mathcal{C} \cup \{E\}$, $\mathcal{B}^* = \mathcal{B} \cup \{E\}$. Let $r = \mathrm{rad}(\mathcal{C})$ and D be the smallest destroyer of $\mathcal{C}$. Since $w(\mathcal{C}) = w(\mathcal{B}) = w(\mathcal{B}^*)$, we have that D is also the smallest destroyer of $\mathcal{B}$ and of $\mathcal{B}^*$. If the radius of E is larger than r, then E cannot be the smallest destroyer of $\mathcal{C}^*$, so $w\big(\mathcal{C}^*\big) = w(\mathcal{C})$. Thus, assume that E has radius less than r. Let v be the extreme point of $\mathcal{C}_{<r}$ and D. Since $w(\mathcal{B}^*) = w(\mathcal{B})$, we know that $d\big(\bigcap \mathcal{B}_{<r}, D\big) = d\big(\bigcap \mathcal{B}^*_{<r}, D\big) = d(v, D)$. Now, Lemma 4.4 implies that $v \in E$, since $E \in \mathcal{B}^*_{<r}$. Thus, the set $\mathcal{C}^*_{<r} = \mathcal{C}_{<r} \cup \{E\}$ is Helly and therefore, there is no disk $D' \in \mathcal{C}^*_{<r}$ that destroys $\mathcal{C}^*_{<r}$. Furthermore, since D destroys $\mathcal{C}_{<r}$ and $\mathcal{C}_{<r} \subset \mathcal{C}^*_{<r}$, the disk D also destroys $\mathcal{C}^*_{<r}$. Therefore, D is also the smallest destroyer of $\mathcal{C}^*$, so $\mathrm{rad}(\mathcal{C}^*) = r = \mathrm{rad}(\mathcal{C})$. Finally, since $\mathcal{B}^*_{<r} \subseteq \mathcal{C}^*_{<r}$ we can use Lemma 4.4 to derive

$$d\Big(\bigcap \mathcal{C}_{<r}, D\Big) = d\Big(\bigcap \mathcal{B}^*_{<r}, D\Big) \leq d\Big(\bigcap \mathcal{C}^*_{<r}, D\Big) \leq d(v, D) = d\Big(\bigcap \mathcal{C}_{<r}, D\Big).$$

The claim follows. □

Next, we are going to describe the violation test for $(\mathcal{D}, w, \leq)$: given a basis $\mathcal{B} \subseteq \mathcal{D}$ and a disk $E \in \mathcal{D}$, check whether E violates $\mathcal{B}$, i.e., whether $w\big(\mathcal{B} \cup \{E\}\big) < w(\mathcal{B})$, and return an error message if $\mathcal{B}$ is not a basis. But first, we show that the combinatorial dimension of $(\mathcal{D}, w, \leq)$ is at most 3.

Lemma 4.7. *For each $\mathcal{C} \subseteq \mathcal{D}$, there is a set $\mathcal{B} \subseteq \mathcal{C}$ with $|\mathcal{B}| \leq 3$ and $w(\mathcal{B}) = w(\mathcal{C})$.*

Proof. Let D be the smallest destroyer of $\mathcal{C}$. Let $r = \mathrm{rad}(\mathcal{C})$ be the radius of D, and let $v \in \bigcap \mathcal{C}_{<r}$ be the extreme point for $\mathcal{C}_{<r}$ and D. First of all, we observe that v cannot be in the interior of $\bigcap \mathcal{C}_{<r}$, since v minimizes the distance to D. Thus, there has to be a non-empty subset $\mathcal{A} \subseteq \mathcal{C}_{<r}$ such that v lies on the boundary of each disk of $\mathcal{A}$. Let $\mathcal{A}$ be a minimal set such that $d(\bigcap \mathcal{A}, D) = d(v, D)$. It follows that $|\mathcal{A}| \leq 2$. See Figure 7 for an illustration.

First, assume that $\mathcal{A} = \{E\}$. Then, since $d(E, D) = d(v, D) > 0$, we know that $E \cap D = \emptyset$. As the disks in $\mathcal{C}$ intersect pairwise, we derive $D \notin \mathcal{C}$ and hence $D = D_\infty$. Setting $\mathcal{B} = \mathcal{A}$, we get $\mathrm{rad}(\mathcal{C}) = \infty = \mathrm{rad}(\mathcal{B})$ and $\mathrm{dist}(\mathcal{C}) = d(v, D) = d(E, D) = \mathrm{dist}(\mathcal{B})$. Thus, $|\mathcal{B}| \leq 3$ and $w(\mathcal{B}) = w(\mathcal{C})$.

Second, assume that $\mathcal{A} = \{E, F\}$. Then, v is one of the two vertices of the lens $L = E \cap F$. Next, we show that $d(L, D) \geq d(v, D)$. Assume for the sake of contradiction that there is a point $w \in L$ with $d(w, D) < d(v, D)$. By general position and since v is the intersection of two disk boundaries, there is a relatively open neighborhood N around v in $\bigcap \mathcal{C}_{<r}$ such that N is also relatively open in L. Since L is convex, there is a point $x \in N$ that also lies in the relative interior of the line segment $\overline{wv}$. Then, $d(x, D) < d(v, D)$ and $x \in \bigcap \mathcal{C}_{<r}$. This yields a contradiction, as v is the extreme point for $\mathcal{C}_{<r}$ and D. Thus, we have $d(L, D) \geq d(v, D)$ which also shows hat $D \cap E \cap F = \emptyset$.

We set $\mathcal{B} = \{E, F\}$, if $\mathcal{C}$ is Helly (i.e., $D = D_\infty$), and $\mathcal{B} = \{D, E, F\}$, if $\mathcal{C}$ is non-Helly (i.e., $D \in \mathcal{C}$). In both cases, we have $\mathcal{B} \subseteq \mathcal{C}$ and $|\mathcal{B}| \leq 3$. Moreover, we can conclude that D destroys $\mathcal{B}_{<r} = \{E, F\}$, and since $\mathcal{B}_{<r}$ is Helly, D is the smallest destroyer of $\mathcal{B}$. Hence, we have $\mathrm{rad}(\mathcal{C}) = r = \mathrm{rad}(\mathcal{B})$.

To obtain $\mathrm{dist}(\mathcal{B}) = \mathrm{dist}(\mathcal{C})$, it remains to show $d(\bigcap \mathcal{B}_{<r}, D) = d(\bigcap \mathcal{C}_{<r}, D)$. Since $\mathcal{B}_{<r} \subseteq \mathcal{C}_{<r}$, we can use Lemma 4.4 as well as $d(L, D) \geq d(v, D)$ to derive

$$d\Big(\bigcap \mathcal{C}_{<r}, D\Big) \geq d\Big(\bigcap \mathcal{B}_{<r}, D\Big), = d(L, D) \geq d(v, D) = d\Big(\bigcap \mathcal{C}_{<r}, D\Big)$$

as desired. We conclude that $w(\mathcal{B}) = w(\mathcal{C})$.

We remark that the set $\mathcal{B}$ is actually a basis for $\mathcal{C}$: if $\mathcal{B}$ is a non-Helly triple, then removing any disk from $\mathcal{B}$ creates a Helly set and increases the radius of the smallest destroyer to ∞. If $|\mathcal{B}| \leq 2$, then D_∞ is the smallest destroyer of $\mathcal{B}$ and the minimality follows directly from the definition. □

Following the argument of the last proof, the violation test is now immediate. We present pseudo-code in Algorithm 1. It obviously needs constant time. Finally, to apply the algorithm of Chazelle and Matoušek, we still need to check that there is a polynomial-time oracle that computes a superset of $\mathcal{R}_\mathcal{Y}$ for a given set of disks $\mathcal{Y}$.

Algorithm 1 The violation test.

procedure VIOLATES(set $\mathcal{B} \subseteq \mathcal{D}$, disk $E \in \mathcal{D}$ with radius r')
    if $|\mathcal{B}| > 3$ or $|\mathcal{B}| = 3$ and $\mathcal{B}$ is Helly **then return "*$\mathcal{B}$* is not a basis."**
    if $|\mathcal{B}| = 2$ and the y-minimum of $\bigcap \mathcal{B}$ is also the y-minimum of a single disk of $\mathcal{B}$ **then**
        return "*$\mathcal{B}$* is not a basis."
    if $\mathcal{B} = \{D_1\}$ **then**
        if the y-minimum in $E \cap D_1$ differs from the y-minimum in D_1 **then**
            return "*E* violates *$\mathcal{B}$*."
        else return "*E* does not violate *$\mathcal{B}$*."
    if $\mathcal{B} = \{D_1, D_2\}$ **then**
        $v = \mathrm{argmin}\ \{w_y \mid w \in D_1 \cap D_2\}$
        if $v \notin E$ **then return "*E* violates *$\mathcal{B}$*."**
        else return "*E* does not violate *$\mathcal{B}$*."
    else ▷ $\mathcal{B}$ is of size 3, non-Helly, and does not contain D_∞.
        $D =$ smallest destroyer of $\mathcal{B}$
        $\{D_1, D_2\} = \mathcal{B} \setminus \{D\}$
        $r = \mathrm{rad}(\mathcal{B})$
        if $r' > r$ **then return "*E* does not violate *$\mathcal{B}$*."**
        else
            $v = \mathrm{argmin}\ \{d(w, E) \mid w \in D_1 \cap D_2\}$
            if $v \notin E$ **then return "*E* violates *$\mathcal{B}$*."**
            else return "*E* does not violate *$\mathcal{B}$*."

Lemma 4.8. *Given a set $\mathcal{Y} \subseteq \mathcal{D}$ of disks, we can compute a superset of $\mathcal{R}_\mathcal{Y}$ in time $O(|\mathcal{Y}|^4)$.*

Proof. Let $v \in \mathbb{R}^2$ and $r > 0$. First, we let $R_v = \{D \in \mathcal{Y} \mid v \notin D\}$ be the range of all disks that do not contain v. Second, let $R_{v,r}$ be the range of all disks of diameter smaller than r that do not contain the point v, i.e., $R_{v,r} = \{D \in \mathcal{Y} \mid v \notin D \text{ and } r_D < r\}$. We define $\mathcal{R}'$ to be the set of all ranges R_v over all v and subsequently, we let $\mathcal{R}''$ be the set of all ranges $R_{v,r}$ over all v and r, that is, $\mathcal{R}'' = \{R_{v,r} \mid v \in \mathbb{R}^2 \text{ and } r > 0\}$.

The discussion from the previous lemmas shows that for any basis $\mathcal{B}$, there is a point $v_\mathcal{B} \in \mathbb{R}^2$ and a radius $r_\mathcal{B} > 0$ such that a disk $E \in \mathcal{D}$ with radius r_E violates $\mathcal{B}$ if and only if $v_\mathcal{B} \notin E$ and $r_E < r_\mathcal{B}$. Hence, we have $\mathcal{R}'' \supseteq \mathcal{R}_\mathcal{Y}$. We show how to compute $\mathcal{R}''$ in polynomial time. For this, we first construct $\mathcal{R}'$.

For the given set $\mathcal{Y}$ of disks, we compute the arrangement $A(\mathcal{Y})$ and then focus on the facets of $A(\mathcal{Y})$. Since the arrangement has $O(|\mathcal{Y}|^2)$ facets, we can compute $A(\mathcal{Y})$ in time $O(|\mathcal{Y}|^3)$ using a simple brute-force approach (faster algorithms exist, but are not needed here). Clearly, for two points v and w of the same facet of $A(\mathcal{Y})$, we have $R_v = R_w$. Therefore, for a given facet f, we pick an arbitrary point

$v \in f$, and we compute R_v by a linear scan of $\mathcal{Y}$. Summing over all facets, we can thus compute $\mathcal{R}'$ in time $O(|\mathcal{Y}|^3)$.

Finally, to compute $\mathcal{R}''$, we iterate over all $O(|\mathcal{Y}|^2)$ ranges in $\mathcal{R}'$. Given a range $R_v \in \mathcal{R}'$, we get all $R_{v,r}$ for $r > 0$ by first sorting R_v by increasing radii and then taking every prefix of the sorted list of disks. For a fixed v, this can be done in time $O(|\mathcal{Y}|^2)$. Hence, $\mathcal{R}''$ can be computed in $O(|\mathcal{Y}|^4)$ time. The claim follows. ◻

The following lemma summarizes the discussion so far.

Lemma 4.9. *Given a set $\mathcal{D}$ of n pairwise intersecting disks in the plane, we can decide in $O(n)$ deterministic time whether $\mathcal{D}$ is Helly. If so, we can compute a point in $\bigcap \mathcal{D}$ in $O(n)$ deterministic time. If not, we can compute the smallest destroyer D of $\mathcal{D}$ and two disks $E, F \in \mathcal{D}_{<r}$ that form a non-Helly triple with D. Here, r is the radius of D.*

Proof. Since (i) $(\mathcal{D}, w, \leq)$ is LP-type, (ii) the violation test needs constant time, and (iii) the oracle needs polynomial time, we can apply the deterministic algorithm of Chazelle and Matoušek [7] to compute $w(\mathcal{D}) = (\text{rad}(\mathcal{D}), -\text{dist}(\mathcal{D}))$ and a corresponding basis $\mathcal{B}$ in $O(n)$ time. Then, $\mathcal{D}$ is Helly if and only if $\text{rad}(\mathcal{D}) = \infty$. If $\mathcal{D}$ is Helly, then $|\mathcal{B}| \leq 2$. We compute the unique point $v \in \bigcap \mathcal{B}$ with $d(v, D_\infty) = d\big(\bigcap \mathcal{B}, D_\infty\big)$. Since $\mathcal{B} \subseteq \mathcal{D}$ and $d\big(\bigcap \mathcal{B}, D_\infty\big) = d\big(\bigcap \mathcal{D}, D_\infty\big)$, we have $v \in \bigcap \mathcal{D}$ by Lemma 4.4. We output v. If $\mathcal{D}$ is non-Helly, we simply output $\mathcal{B}$, because $\mathcal{B}$ is a non-Helly triple with the smallest destroyer D of $\mathcal{D}$ and two disks $E, F \in \mathcal{D}_{<r}$, where r is the radius of D. ◻

Theorem 4.10. *Given a set $\mathcal{D}$ of n pairwise intersecting disks in the plane, we can find in deterministic $O(n)$ time a set P of five points such that every disk of $\mathcal{D}$ contains at least one point of P.*

Proof. Using the algorithm from Lemma 4.9, we decide whether $\mathcal{D}$ is Helly. If so, we return the extreme point computed by the algorithm. Otherwise, the algorithm gives us a non-Helly triple $\{D, E, F\}$, where D is the smallest destroyer of $\mathcal{D}$ and $E, F \in \mathcal{D}_{<r}$, with r being the radius of D. Since $\mathcal{D}_{<r}$ is Helly, we can obtain in $O(n)$ time a stabbing point $q \in \bigcap \mathcal{D}_{<r}$ by using the algorithm from Lemma 4.9 again. Next, by Lemma 2.1, there are two disks in $\{D, E, F\}$ whose lens angle is at least $2\pi/3$. Let P' be the set of four points from the proof of Lemma 2.4. Then, $P = P' \cup \{q\}$ is a set of five points that stabs every disk in $\mathcal{D}$. ◻

5 Simple Bounds

We now provide some easy lower and upper bounds on the number of disks for which a certain number of stabbing points is necessary or sufficient.

Eight disks can be stabbed by three points. For the proof that any set of eight pair-wise intersecting disks can be stabbed by at most three points, we show the following lemma.

Lemma 5.1. *Let $\mathcal{D}$ be a set of at least 5 pairwise intersecting disks. Then, $\mathcal{D}$ contains a Helly-triple.*

Proof. Let $\mathcal{D}$ be a set of exactly 5 pairwise intersecting disks. We assume that no three centers of the disks are on a line, since otherwise these three disks are a Helly-triple. Since the complete graph K_5 does not have a planar embedding, there have to be four different disks $D_1, \dots, D_4 \in \mathcal{D}$ with centers $c_1, \dots, c_4$ and radii $r_1, \dots, r_4$ such that the line segments c_1c_3 and c_2c_4 intersect, see Figure 8. Let x be the intersection point. Moreover, let α (resp., β) be the intersection of the lens $L_{1,3}$ (resp., $L_{2,4}$) and the line segment c_1c_3 (resp., c_2c_4). If x is in α or β, we are done. Otherwise, let y be the point of α that is closest to x and let z be the point of β closest to x. We can assume without loss of generality that $|xy| \leq |xz|$ and $x \notin D_4$. Using the triangle inequality, We can derive

$$|c_2y| \leq |c_2x| + |xy| \leq |c_2x| + |xz| \leq r_2$$

to conclude that $y \in D_1 \cap D_2 \cap D_3$. ◻

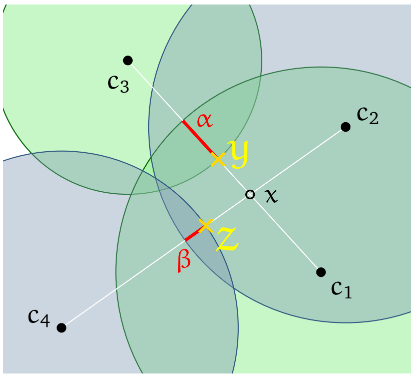

Figure 8: Proof of Lemma 5.1.

Now consider a set $\mathcal{D}$ of 8 pairwise intersecting disks. Using Lemma 5.1, we can find a Helly-triple in $\mathcal{D}$. Among the remaining 5 disks, we find a second Helly-triple. The remaining two disks can be stabbed by one point. This reasoning yields the following corollary, which was already mentioned by Stachó [20].

Corollary 5.2. *Every set $\mathcal{D}$ of at most 8 pairwise intersecting disks can be stabbed by 3 points.*

13 disks with 4 stabbing points. Danzer presented a set of 10 pairwise intersecting pseudo-disks with stabbing number four [9]. However, it is not clear to us how these 10 pseudo-disks can be realized as pairwise intersecting Euclidean disks achieving the same stabbing number. Moreover, it is another open problem whether 9 pairwise intersecting disks can be stabbed by three points. Instead, we want to describe a set of 13 pairwise intersecting disks in the plane such that no point set of size three can pierce all of them.

The construction begins with an inner disk A of radius 1 and three larger disks D_1, D_2, D_3 of equal radius, so that each pair of disks in $\{A, D_1, D_2, D_3\}$ is tangent. For $i = 1, 2, 3$, we denote the contact point of A and D_i by ξ_i.

We add six more disks as follows. For $i = 1, 2, 3$, we draw the two common outer tangents to A and D_i, and denote by T_i^- and T_i^+ the halfplanes that are bounded by these tangents and are openly disjoint from A. The labels T_i^- and T_i^+ are chosen such that the points of tangency between A and T_i^-, D_i, and T_i^+, appear along the boundary of A in this counterclockwise order. One can show that the nine points of tangency between A and the other disks and tangents are pairwise distinct (see Figure 9). We regard the six halfplanes T_i^-, T_i^+, for $i = 1, 2, 3$, as (very large) disks; in the end, we can apply a suitable inversion to turn the disks and halfplanes into actual disks, if so desired.

Finally, we construct three additional disks A_1, A_2, A_3. To construct A_i, we slightly expand A into a disk A_i' of radius $1 + \varepsilon_1$, while keeping the tangency with D_i at ξ_i. We then roll A_i' clockwise along D_i, by a tiny angle $\varepsilon_2 \ll \varepsilon_1$, to obtain A_i.

This gives a set of 13 disks. For sufficiently small ε_1 and ε_2, we can ensure the following properties for each A_i: (i) A_i intersects all other 12 disks; (ii) the nine intersection regions $A_i \cap D_j$, $A_i \cap T_j^-$, $A_i \cap T_j^+$, for $j = 1, 2, 3$, are pairwise disjoint; and (iii) $\xi_i \notin A_i$.

Theorem 5.3. *The construction yields a set of* 13 *disks that cannot be stabbed by* 3 *points.*

Proof. Consider any set P of three points. Set $A^* = A \cup A_1 \cup A_2 \cup A_3$. If $P \cap A^* = \emptyset$, we have unstabbed disks, so suppose that $P \cap A^* \neq \emptyset$. For $p \in P \cap A^*$, property (ii) implies that p stabs at most one of the nine remaining disks D_j, T_j^+ and T_j^-, for $j = 1, 2, 3$. Thus, if $P \subset A^*$, we would have unstabbed disks, so we may assume that $|P \cap A^*| \in \{1, 2\}$.

Suppose first that $|P \cap A^*| = 2$. As just argued, at most two of the remaining disks are stabbed by $P \cap A^*$. The following cases can then arise.

(a) None of D_1, D_2, D_3 is stabbed by $P \cap A^*$. Since $\{D_1, D_2, D_3\}$ is non-Helly and a non-Helly set must be stabbed by at least two points, at least one disk remains unstabbed.

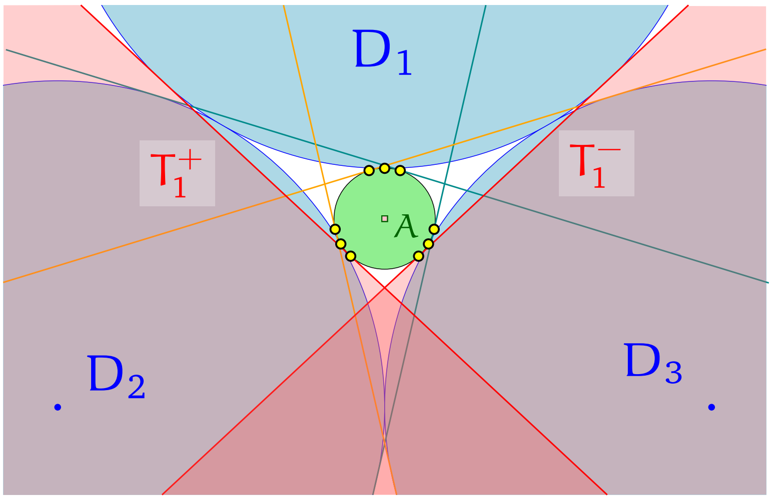

Figure 9: Each common tangent ℓ between A and D_i represents a very large disk, whose interior is disjoint from A. The nine points of tangency are pairwise distinct.

(b) Two disks among D_1, D_2, D_3 are stabbed by $P \cap A^*$. Then the six unstabbed halfplanes form many non-Helly triples, e.g., T_1^-, T_2^-, and T_3^-, and again, a disk remains unstabbed.

(c) The set $P \cap A^*$ stabs one disk in $\{D_1, D_2, D_3\}$ and one halfplane. Then, there is (at least) one disk D_i such that D_i and its two tangent halfplanes T_i^-, T_i^+ are all unstabbed by $P \cap A^*$. Then, $\{D_i, T_i^-, T_i^+\}$ is non-Helly, and at least 2 more points are needed to stab it.

Suppose now that $|P \cap A^*| = 1$, and let $P \cap A^* = \{p\}$. We may assume that p stabs all four disks A, A_1, A_2, A_3, since otherwise a disk would stay unstabbed. By property (iii), we can derive $p \notin \{\xi_1, \xi_2, \xi_3\}$. Now, since $p \in A \setminus \{\xi_1, \xi_2, \xi_3\}$, the point p does not stab any of D_1, D_2, D_3. Moreover, by property (ii), the point p can only stab at most one of the remaining halfplanes. Since $\{D_1, D_2, D_3\}$ is non-Helly, it requires two stabbing points. Moreover, since $|P \setminus \{p\}| = 2$, it must be the case that one point q of $P \setminus A^*$ is the point of tangency of two of these disks, say $q = D_2 \cap D_3$. Then, q stabs only two of the six halfplanes, say, T_1^- and T_1^+. But then, $\{D_1, T_2^+, T_3^-\}$ is non-Helly and does not contain any point from $\{p, q\}$. At least one disk remains unstabbed. □

6 Conclusion

We gave a simple linear-time algorithm, based on techniques for solving LP-type problems, to find five stabbing points for a set of pairwise intersecting disks in the plane. The arXiv manuscript by Carmi, Katz, and Morin [4] claims a similar linear-time algorithm for finding four stabbing points. It would now be interesting to see whether these results, the ones by Danzer, Stachó, and ours, could be used to find new deterministic approximation algorithms for computing large cliques in disk graphs; refer to [2, 3] for the known algorithms. On the lower-bound side, it is still not known whether nine disks can always be stabbed by three points or not. For eight disks, we provided a proof that three points always suffice, as already mentioned by Stachó [20]. The lower bound construction of Danzer with ten disks [9] can easily be verified for pseudo-disks. However, the example is not easy to draw, even with the help of geometry processing software. Until now, we were not able to check whether his pseudo-disk arrangement can be realized as a Euclidean disk arrangement.

References

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

[2] C. Ambühl and U. Wagner. The clique problem in intersection graphs of ellipses and triangles. *Theory of Computing Systems*, 38(3):279–292, 2005.

[3] M. Bonamy, E. Bonnet, N. Bousquet, P. Charbit, and S. Thomassé. EPTAS for max clique on disks and unit balls. In *Proc. 59th Annu. IEEE Sympos. Found. Comput. Sci. (FOCS)*, pages 568–579, 2018.

[4] P. Carmi, M. J. Katz, and P. Morin. Stabbing pairwise intersecting disks by four points. `arXiv:1812.06907`, 2018.

[5] T. M. Chan. An optimal randomized algorithm for maximum Tukey depth. In *Proc. 15th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 430–436, 2004.

[6] B. Chazelle. *The Discrepancy Method—Randomness and Complexity.* Cambridge University Press, Cambridge, 2001.

[7] B. Chazelle and J. Matoušek. On linear-time deterministic algorithms for optimization problems in fixed dimension. *J. Algorithms*, 21(3):579–597, 1996.

[8] B. N. Clark, C. J. Colbourn, and D. S. Johnson. Unit disk graphs. *Discrete Mathematics*, 86(1-3):165–177, 1990.

[9] L. Danzer. Zur Lösung des Gallaischen Problems über Kreisscheiben in der Euklidischen Ebene. *Studia Sci. Math. Hungar.*, 21(1-2):111–134, 1986.

[10] A. Dumitrescu and M. Jiang. Piercing translates and homothets of a convex body. *Algorithmica*, 61(1):94–115, 2011.

[11] B. Grünbaum. On intersections of similar sets. *Portugal. Math.*, 18:155–164, 1959.

[12] H. Hadwiger and H. Debrunner. Ausgewählte Einzelprobleme der kombinatorischen Geometrie in der Ebene. *Enseignement Math. (2)*, 1:56–89, 1955.

[13] E. Helly. Über Mengen konvexer Körper mit gemeinschaftlichen Punkten. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 32:175–176, 1923.

[14] E. Helly. Über Systeme von abgeschlossenen Mengen mit gemeinschaftlichen Punkten. *Monatshefte für Mathematik*, 37(1):281–302, 1930.

[15] R. M. Karp. Reducibility among combinatorial problems. In *Proceedings of a symposium on the Complexity of Computer Computations*, pages 85–103, 1972.

[16] M. Löffler and M. van Kreveld. Largest bounding box, smallest diameter, and related problems on imprecise points. *Comput. Geom.*, 43(4):419–433, 2010.

[17] J. Radon. Mengen konvexer Körper, die einen gemeinsamen Punkt enthalten. *Mathematische Annalen*, 83(1):113–115, 1921.

[18] R. Seidel. Small-dimensional linear programming and convex hulls made easy. *Discrete Comput. Geom.*, 6:423–434, 1991.

[19] M. Sharir and E. Welzl. A combinatorial bound for linear programming and related problems. *Proc. 9th Sympos. Theoret. Aspects Comput. Sci. (STACS)*, pages 567–579, 1992.

[20] L. Stachó. Über ein Problem für Kreisscheibenfamilien. *Acta Sci. Math. (Szeged)*, 26:273–282, 1965.

[21] L. Stachó. A solution of Gallai's problem on pinning down circles. *Mat. Lapok*, 32(1-3):19–47, 1981/84.

[22] D. Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. In *Proc. 38th Annu. ACM Sympos. Theory Comput. (STOC)*, pages 681–690, 2006.