The Discrete 2-Center Problem*

Pankaj K. Agarwal[†] Micha Sharir[‡] Emo Welzl[§]

March 10, 1998

Abstract

We present an algorithm for computing the *discrete 2-center* of a set P of n points in the plane; that is, computing two congruent disks of smallest possible radius, centered at two points of P, whose union covers P. Our algorithm runs in time $O(n^{4/3}\log^5 n)$.

1 Introduction

Problem statement and previous results. Let P be a set of n points in the plane. The *discrete 2-center* problem for P is to cover P by (the union of) two congruent closed disks whose radius is as small as possible, and whose centers are two points of P. This is a restricted version of the standard 2-center problem, where the centers of the two covering disks can be any pair of points in the plane. This latter problem has been studied extensively, where the best algorithm, due to Sharir [14] and slightly improved by Eppstein [6], runs in randomized expected $O(n\log^2 n)$ time.

The discrete 2-center problem has been studied in [7], where a near-quadratic algorithm is proposed (such an algorithm is briefly described later in this introduction). Before discussing it further, we note that the discrete 1-center problem, seeking the smallest disk

[*] Pankaj Agarwal and Micha Sharir have been supported by a joint grant from the U.S.–Israeli Binational Science Foundation. Work by Pankaj Agarwal has also been supported by a National Science Foundation Grant CCR-93–01259, by an Army Research Office MURI grant DAAH04-96-1-0013, by a Sloan fellowship, and by an NYI award and matching funds from Xerox Corporation. Work by Micha Sharir has also been supported by National Science Foundation Grants CCR-94-24398 and CCR-93-11127, and by a grant from the G.I.F. — the German Israeli Foundation for Scientific Research and Development. Work by Micha Sharir and Emo Welzl has been supported by a Max-Planck Research Award.

[†] Center for Geometric Computing, Department of Computer Science, Box 90129, Duke University, Durham, NC 27708-0129, USA.

[‡] School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA.

[§] Institut für Theoretische Informatik, ETH Zürich, CH-8092 Zürich, Switzerland.

[0] A preliminary version of this paper appeared in *Proceedings of the 13th Anuual Symposium on Computational Geometry, 1997, pp. 147–155.*

centered at a point of P and containing P, is much easier to solve, in time $O(n \log n)$, using the furthest-neighbor Voronoi diagram of P. That is, the diagram allows us to find, in $O(n \log n)$ time, the furthest neighbor $f(p)$ of each point $p \in P$. The point p that minimizes the distance between p and $f(p)$ is the center of the desired smallest enclosing disk.

The discrete 2-center appears to be more difficult than the standard 2-center problem. Both problems involve a 'decision procedure' that, given a fixed radius r, aims to determine whether P can be covered by two disks of radius r. As an informal explanation of the additional difficulty of the discrete 2-center problem, suppose that we have already guessed one center p. The second center must then lie in ('pierce') each of the disks of radius r centered at the points of P and not containing p. In the standard 2-center problem we simply need to determine whether the intersection of all these disks is nonempty, whereas in the discrete 2-center problem we need to determine whether this intersection contains a point of P, which is a harder task.

Main results and the overall approach. In this paper we obtain an efficient solution to the discrete 2-center problem that runs in time $O(n^{4/3} \log^5 n)$. This is the first subquadratic algorithm for solving the problem. We note that a near-quadratic solution is rather easy: It suffices to show a near-quadratic solution to the fixed-size problem, and then follow (a simpler version of) the binary-search technique (on the radius r) that is described below in Section 5, to find the optimal solution. The fixed-size problem, for a given radius r, determines whether there exist $p, q \in P$ so that $P \subset D(p,r) \cup D(q,r)$, where $D(x,r)$ denotes the closed disk of radius r centered at x. We try each point $p \in P$ as the first center and obtain the set $N_p \subset P$ of points not contained in $D(p,r)$. By computing the farthest-point Voronoi diagram of N_p, we can determine in $O(n \log n)$ time whether there exists a point $q \in P$ so that $N_p \subset D(q,r)$. The running time of the fixed-size procedure is therefore $O(n^2 \log n)$.

In order to improve the running time of the fixed-size problem, we proceed as follows: For each $p \in P$, let K_p be the intersection of all the disks $D(q,r)$ centered at the points of P and not containing p. If any set K_p contains a point q of P, then we are done: p and q are centers of two disks of radius r whose union covers P. Conversely, if $p, q \in P$ are centers of two such disks, then $p \in K_q$ and $q \in K_p$. In other words, we need to compute the union U of all the K_p's, and determine whether $U \cap P \neq \emptyset$. The difficult step is to compute U in time close to $n^{4/3}$.

We consider a more general problem: Let P be a set of m points and $\mathcal{D}$ a set of n congruent disks. For each $p \in P$, define $\mathcal{D}_p = \{D \in \mathcal{D} \mid p \notin D\}$, $K_p = \bigcap_{D \in \mathcal{D}_p} D$, $\mathcal{K} = \{K_p \mid p \in P\}$, and $U = \bigcup_{p \in P} K_p$. In Section 2 we present some important properties of $\mathcal{K}$, which we believe to be of independent interest. The main property is that $\mathcal{K}$ is a collection of *convex pseudo-disks*; i.e., these sets are compact and convex, and for any pair K_p, K_q of such sets, both $K_p \setminus K_q$ and $K_q \setminus K_p$ are connected.

In Section 3 we show that the combinatorial complexity of U is $O(m^{2/3} n^{2/3} \log^{1/3} n +$

$n \log n$). While this bound is non-trivial, and 'consistent' with the running time we are aiming at, we have so far been unable to exploit this bound to obtain an alternative simpler solution, of comparable complexity, of the discrete 2-center problem. The reasons for this are technical, and will be noted below.

In Section 4 we present an $O(n^{4/3} \log^4 n)$-time algorithm for computing U (and for testing whether $U \cap P \neq \emptyset$) for the case in which $\mathcal{D} = \{D(p,r) \mid p \in P\}$. The algorithm constructs and searches in U in a semi-implicit manner, using appropriate range-searching data structures and techniques similar to those used in parametric searching, for performing various primitive operations on the semi-implicit representation of U. Finally, we describe the overall algorithm for the discrete 2-center in Section 5. As mentioned, its running time is $O(n^{4/3} \log^5 n)$.

2 Structure of $\mathcal{K}$

In this section we prove some interesting properties of $\mathcal{K}$. These properties, besides being of independent interest, are crucial for making our algorithm efficient.

Theorem 2.1 *Let $\mathcal{D}$ be a finite set of congruent disks in the plane, and let P be a finite set of points. Let $\mathcal{K}$ be the same as defined in the introduction. Then $\mathcal{K}$ is a family of convex pseudo-disks; that is, each K_p is a compact convex set, and for each pair of distinct sets K_p, K_q, both sets $K_p \setminus K_q$ and $K_q \setminus K_p$ are connected.*

We prove the theorem by a sequence of lemmas.

Lemma 2.2 *For a point p, two distinct disks $D_1, D_2 \in \mathcal{D}$ that do not contain p and another disk $D \in \mathcal{D}$ that contains p, the set $D \setminus (D_1 \cap D_2)$ is connected.*

Proof: Suppose to the contrary that $D \setminus (D_1 \cap D_2)$ is disconnected. Since $p \in D$ and $p \notin D_1, D_2$, all three disks D, D_1, D_2 are distinct. Since $D \setminus (D_1 \cap D_2)$ is disconnected, ∂D and $\partial(D_1 \cap D_2)$ must cross at exactly 4 points, all lying on the boundary of $E = D \cap D_1 \cap D_2$. This however is impossible, since the intersection of three congruent disks can have at most three such intersection points on its boundary. □

Corollary 2.3 *For a point $p \in P$ and a disk $D \in \mathcal{D}$ that contains p, the set $K_p \setminus D$ is connected.*

Proof: Suppose to the contrary that $K_p \setminus D$ is disconnected. Since $D \notin \mathcal{D}_p$, it is distinct from any of the disks that form K_p, so any intersection of ∂D with ∂K_p must be a proper crossing. Moreover, since $K_p \setminus D$ is disconnected, the boundaries of D and of K_p must cross at least four times. This, however, implies that $D \setminus K_p$ is also disconnected (this follows from the convexity of K_p). But

$$D \setminus K_p = D \setminus \bigcap_{D' \in \mathcal{D}_p} D' = D \cap \bigcup_{D' \in \mathcal{D}_p} (D')^c = \bigcup_{D' \in \mathcal{D}_p} (D \setminus D').$$

If a union of a collection of sets is disconnected, then either one of the sets is disconnected, or there exist two disjoint sets in the collection. In our setting, $D \setminus D'$ is always connected, and the second case contradicts Lemma 2.2 (because for $D', D'' \in \mathcal{D}_p$, $(D \setminus D') \cup (D \setminus D'') = D \setminus (D' \cap D'')$). Hence, $K_p \setminus D$ is connected. □

Lemma 2.4 *For a point p, two distinct disks $D_1, D_2 \in \mathcal{D}$ that do not contain p and two other distinct disks $D_3, D_4 \in \mathcal{D}$ that contain p, the set $(D_1 \cap D_2) \setminus (D_3 \cap D_4)$ is connected.*

Proof: Suppose again to the contrary that $(D_1 \cap D_2) \setminus (D_3 \cap D_4)$ is disconnected. Lemma 2.2 implies (using the argument in the proof of Corollary 2.3) that $(D_1 \cap D_2) \setminus D_3$ is connected, and so is $(D_1 \cap D_2) \setminus D_4$. It follows that $\partial(D_1 \cap D_2)$ and ∂D_3 intersect at most twice, and the same holds for $\partial(D_1 \cap D_2)$ and ∂D_4. These conditions, along with our assumption, imply that $\partial(D_1 \cap D_2)$ and $\partial(D_3 \cap D_4)$ intersect exactly 4 times. Moreover, put $E = D_1 \cap D_2 \cap D_3 \cap D_4$, fix a point o in the interior of E, and consider the boundaries of $D_1, \ldots, D_4$ and E as graphs of functions $r = D_1(\theta), \ldots, r = E(\theta)$ in polar coordinates about o. Let u, v, w, z be the four points of intersection between $\partial(D_1 \cap D_2)$ and $\partial(D_3 \cap D_4)$, appearing in this circular counter-clockwise order along ∂E. Let $\theta_u < \theta_v < \theta_w < \theta_z$ be the polar orientations of u, v, w, z, respectively. Since $D_1, \ldots, D_4$ are congruent disks, each ∂D_i appears along ∂E in a single connected arc. Hence, with no loss of generality, we may assume that ∂E is attained by ∂D_1 over $[\theta_u, \theta_v]$, by ∂D_3 over $[\theta_v, \theta_w]$, by ∂D_2 over $[\theta_w, \theta_z]$, and by ∂D_4 over $[\theta_z, \theta_u]$. See Figure 1. Let θ_p be the polar orientation of $\vec{op}$. It is impossible that

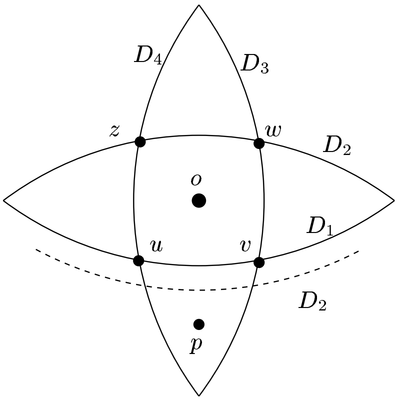

Figure 1: The proof of Lemma 2.4

θ_p lies in $[\theta_v, \theta_w]$, for otherwise, since $p \in D_3$, we have $|op| \leq D_3(\theta_p) \leq D_1(\theta_p)$, implying that $p \in D_1$, contrary to assumption. Similarly, θ_p cannot lie in $[\theta_z, \theta_u]$. (We use the notation $[\theta, \theta']$ to denote the angular interval extending counter-clockwise from θ to θ'.) Suppose then that $\theta_p \in [\theta_u, \theta_v]$. Let $F = D_3 \cap D_4$ and regard it too as a graph $r = F(\theta)$. Since $p \in F$ but $p \notin D_2$, we have $D_2(\theta_p) < F(\theta_p)$, and this inequality is reversed over the intervals $[\theta_v, \theta_w]$ and $[\theta_z, \theta_u]$. It follows that ∂F and ∂D_2 intersect at least twice over the interval $[\theta_u, \theta_v]$, which, together with w and z, yields 4 points of intersection between these boundaries, all lying along $\partial(D_2 \cap D_3 \cap D_4)$. This is impossible for congruent disks (cf. the proof of Lemma 2.2). A similar contradiction occurs when $\theta_p \in [\theta_w, \theta_z]$. All these contradictions establish the lemma. □

Following the same argument as in the proof of Corollary 2.3, we obtain:

Corollary 2.5 *For a point $p \in P$ and two disks $D, D' \in \mathcal{D}$ that contain p, the set $K_p \setminus (D \cap D')$ is connected.*

We are now ready to prove Theorem 2.1.

Proof of Theorem 2.1: Suppose to the contrary that there exist $p, q \in P$ such that $K_p \setminus K_q$ is disconnected. As in the proof of Corollary 2.3, we can express $K_p \setminus K_q$ as

$$\begin{aligned} K_p \setminus K_q &= K_p \setminus \bigcap_{D \in \mathcal{D}_q} D = K_p \cap \bigcup_{D \in \mathcal{D}_q} D^c \\ &= \bigcup_{D \in \mathcal{D}_q} (K_p \setminus D) = \bigcup_{D \in \mathcal{D}_q \setminus \mathcal{D}_p} (K_p \setminus D)\,. \end{aligned}$$

(The last equality follows from the fact that the disks in $\mathcal{D}_q \cap \mathcal{D}_p$ contribute empty sets to this union.) Again, as in the proof of Corollary 2.3, if $K_p \setminus K_q$ is disconnected, then either there exists a disk $D \in \mathcal{D}_q \setminus \mathcal{D}_p$ so that $K_p \setminus D$ is disconnected, or there exists two disks $D_1, D_2 \in \mathcal{D}_q \setminus \mathcal{D}_p$ so that $K_p \setminus (D_1 \cap D_2)$ is disconnected. The first condition contradicts Corollary 2.3 and the second contradicts Corollary 2.5. Hence, $K_p \setminus K_q$ is connected (and so is $K_q \setminus K_p$). □

Theorem 2.1 fails for non-congruent disks, as is illustrated in Figure 2. Nevertheless, the following variant of the theorem holds in even more generality:

Theorem 2.6 *Let $\mathcal{D}$ be a finite set of convex pseudo-disks in the plane; that is, each $D \in \mathcal{D}$ is a compact convex set, and for each pair of distinct sets $D, D' \in \mathcal{D}$, both sets $D \setminus D'$ and $D' \setminus D$ are connected; we also assume that ∂D and $\partial D'$ cross each other transversally at any point of intersection. Let P be a finite set of points. For each $p \in P$, let $\mathcal{D}_p$ denote the set of pseudo-disks in $\mathcal{D}$ that do not contain p, and let K_p denote their intersection. Then, for any $p, q \in P$, ∂K_p and ∂K_q can cross each other at most twice.*

(Note that for the sets K_p and K_q in Figure 2, their boundaries do not cross at all.)

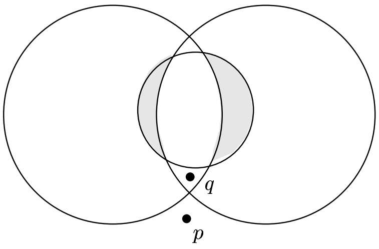

Figure 2: $K_q \setminus K_p$ (the shaded region) consists of two connected components.

Proof: We partition $\mathcal{D}$ into four subsets: the subset $\mathcal{D}_p \cap \mathcal{D}_q$ of pseudo-disks that contain neither p nor q, the subset $\mathcal{D}_p \setminus \mathcal{D}_q$ of pseudo-disks that contain q but not p, the subset $\mathcal{D}_q \setminus \mathcal{D}_p$ of pseudo-disks that contain p but not q, and the subset $\mathcal{D} \setminus (\mathcal{D}_p \cup \mathcal{D}_q)$ of pseudo-disks that contain both p and q. We can ignore the last subset since the pseudo-disks in this set have no effect on K_p or K_q. Let $\mathcal{I} = \bigcap(\mathcal{D}_p \cap \mathcal{D}_q)$. Clearly, both K_p and K_q are contained in $\mathcal{I}$, so any crossing between their boundaries must be interior to $\mathcal{I}$. In particular, if such a crossing occurs between a pseudo-disk $D \in \mathcal{D}_p$ and a pseudo-disk $D' \in \mathcal{D}_q$, then we must have $D \in \mathcal{D}_p \setminus \mathcal{D}_q$ and $D' \in \mathcal{D}_q \setminus \mathcal{D}_p$ (that is, $q \in D$ and $p \in D'$).

Now suppose that ∂K_p and ∂K_q cross each other three times, at points u, v, and w. By the above argument, there exist six (not necessarily all distinct) pseudo-disks, $D_u^{(p)}$, $D_u^{(q)}$, $D_v^{(p)}$, $D_v^{(q)}$, $D_w^{(p)}$, $D_w^{(q)}$, such that $D_u^{(p)}$, $D_v^{(p)}$, and $D_w^{(p)}$ are in $\mathcal{D}_p \setminus \mathcal{D}_q$, $D_u^{(q)}$, $D_v^{(q)}$, and $D_w^{(q)}$ are in $\mathcal{D}_q \setminus \mathcal{D}_p$, $D_u^{(p)}$ and $D_u^{(q)}$ cross at u, $D_v^{(p)}$ and $D_v^{(q)}$ cross at v, and $D_w^{(p)}$ and $D_w^{(q)}$ cross at w.

Let o be a point in the interior of $K_p \cap K_q$. There must exist two of the crossing points, say u and v, such that p and q appear between u and v in counter-clockwise angular order about o. Without loss of generality, assume that u, p, q, and v appear in this counter-clockwise order about o, and let $\theta_u < \theta_p < \theta_q < \theta_v$ be the orientations of the vectors $\vec{ou}$, $\vec{op}$, $\vec{oq}$, and $\vec{ov}$, respectively.

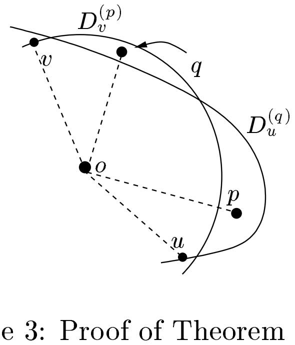

Figure 3: Proof of Theorem 2.6.

Now consider the two pseudo-disks $D_u^{(q)}$ and $D_v^{(p)}$, and regard their boundaries as functions $r = D_u^{(q)}(\theta)$ and $r = D_v^{(p)}(\theta)$ in polar coordinates about o. Then we have (see Figure 3):

$$\begin{aligned} D_u^{(q)}(\theta_u) &\le D_v^{(p)}(\theta_u) \\ D_u^{(q)}(\theta_p) &> D_v^{(p)}(\theta_p) \\ D_u^{(q)}(\theta_q) &< D_v^{(p)}(\theta_q) \\ D_u^{(q)}(\theta_v) &\ge D_v^{(p)}(\theta_v)\,. \end{aligned}$$

These inequalities follow from the convexity of K_p and K_q, from the fact that u and v lie on their boundaries, and from the fact that $D_v^{(p)} \in \mathcal{D}_p \setminus \mathcal{D}_q$ and $D_u^{(q)} \in \mathcal{D}_q \setminus \mathcal{D}_p$. However, this implies that $D_u^{(q)}$ and $D_v^{(p)}$ intersect at least three times, contradicting the assumption that $\mathcal{D}$ is a set of pseudo-disks. This completes the proof. □

The following corollary is an immediate consequence of the results of [10]:

Corollary 2.7 *In the setting of Theorem 2.6, if P has m points, then the boundary of $\bigcup_{p\in P} K_p$ consists of $O(m)$ connected portions of the boundaries of the individual K_p's.*

We now return to the assumption that $\mathcal{D}$ is a set of congruent disks. For a point $p \in P$, we say that p lies *above* (resp. *below*) K_p if the downward-directed (resp. upward-directed) vertical ray from p intersects K_p, and p lies to the *left* (resp. *right*) of K_p if p and K_p are separated by a vertical line, so that p lies to the left (resp. right) of the line. To facilitate our solution to the fixed-size decision problem, presented in Section 4, we need the following stronger property of the K_p's: Define

$$\begin{aligned} P_T &= \{p \in P \mid p \text{ lies above } K_p\}, \\ P_B &= \{p \in P \mid p \text{ lies below } K_p\}, \\ P_L &= \{p \in P \mid p \text{ lies to the left of } K_p\}, \\ P_R &= \{p \in P \mid p \text{ lies to the right of } K_p\}. \end{aligned}$$

Theorem 2.8 *Let p, q be two distinct points in P_T. Then the top boundaries of K_p and K_q cross at most once, and the same holds for the bottom boundaries. The same properties hold for each of the other three sets P_B, P_L, and P_R.*

Proof: Suppose that the top boundaries of K_p and K_q cross at two points u and v. The intersection u must be witnessed by two disks $D_p^{(u)}$ and $D_q^{(u)}$ with u on the top boundaries of these disks, $D_p^{(u)} \cap \{p, q\} = \{q\}$ and $D_q^{(u)} \cap \{p, q\} = \{p\}$. Similarly, there exist witness disks $D_p^{(v)}$ and $D_q^{(v)}$, with similar properties, for the intersection v.

Let us first prove that two such intersections are not possible if $p, q \in P_T$.

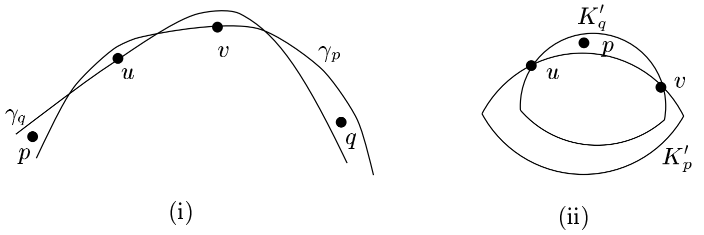

Figure 4: Impossible crossings between the top boundaries: (i) $p, q \in P_T$; (ii) $p, q \in P_B$.

We call the top boundary of a disk D extended by vertical rays downwards at its endpoints the *top curve of* D. Since we are dealing with disks of equal radius, the top curves of $D_p^{(u)}$ and $D_q^{(u)}$ intersect in exactly one point (they have to intersect, since otherwise K_p and K_q are disjoint). Since p is above $D_p^{(u)}$, and not above $D_q^{(u)}$, and vice versa for q, the x-coordinate of this unique intersection has to lie between the x-coordinates of p and q. So we have shown that the x-coordinate of u has to lie between p and q, and the same is true for v. We may assume that the x-coordinates of p, u, v, and q appear in this increasing order.

Now consider the top curves $\gamma_p^{(v)}$ and $\gamma_q^{(u)}$ of disks $D_p^{(v)}$ and $D_q^{(u)}$, respectively, and refer to Figure 4 (i). The curve $\gamma_p^{(v)}$ lies below p (since $p \notin D_p^{(v)}$ and p lies above K_p), lies above u or passes through u (otherwise u cannot lie on the boundary of K_p), passes through v, and lies above q (since $q \in D_p^{(v)}$). The curve $\gamma_q^{(u)}$ lies above p, passes through u, continues above v and lies below q. It follows that the two curves switch sides three times: between p and u, between u and v, and between v and q. (This also covers the case where, say, $\gamma_p^{(v)}$ passes through u, because the curves $\gamma_p^{(v)}$ and $\gamma_q^{(u)}$ must *cross* at this point, as is easily verified.) This gives three intersections of these curves, a contradiction, which concludes the argument for the case in which $p, q \in P_T$.

Suppose next that $p, q \in P_B$. Let $D_p^{(u)}$, $D_q^{(u)}$, $D_p^{(v)}$, and $D_q^{(v)}$ be four respective witness disks, defined as above. We exploit now the previously proved fact that $K'_p := D_p^{(v)} \cap D_p^{(u)}$ and $K'_q := D_q^{(v)} \cap D_q^{(u)}$ behave like pseudo-disks and thus their boundaries do not cross at any point other than u and v. Let us assume that the top boundary of K'_q lies above the top boundary of K'_p in the range between the x-coordinates of u and v. Now recall that p must lie in $K'_q \setminus K'_p$, and thus it lies below the top boundary of K'_q and *above* the top boundary of K'_p, in contradiction to the fact that p lies below K_p which is contained in K'_p; see Figure 4 (ii).

We now switch to the case of P_L, where p and q lie to the left of their regions K_p and K_q. Without loss of generality, suppose that p lies to the left of q. Since K_q is to the right of q, any intersection of the boundaries of K_p and K_q must lie to the right of q. So let us assume that two such intersections u and v exist, both between the top boundaries of K_p and K_q, and that the x-coordinates of p, q, u, and v appear in this increasing order. Let $D_p^{(u)}$, $D_q^{(u)}$, $D_p^{(v)}$, and $D_q^{(v)}$ be four respective witness disks, defined as above.

First consider the top boundary of the disk $D_q^{(v)}$. It must lie above p and u and pass through v. Let us assume that it lies below q (which we will lead to a contradiction) and refer to Figure 5(i). The top boundary of $D_p^{(u)}$ lies above q, goes through u and lies above (or passes through) v, and so it must intersect the top boundary of $D_q^{(v)}$ twice, once between q and u, and once between u and v; a contradiction (as in a preceding argument, this also covers the case where $\partial D_q^{(v)}$ passes through u). So the top boundary of $D_q^{(v)}$ must lie above q, which implies that the whole disk $D_q^{(v)}$ must lie above q, since $q \notin D_q^{(v)}$. This implies that $D_q^{(u)}$ must also lie above q, for otherwise it must lie entirely below q, and so the vertical line through q is disjoint from the intersection of $D_q^{(v)}$ and $D_q^{(u)}$. But this intersection contains p to the left of this line, and the point u to the right of this line, which is a contradiction, since this intersection has to be connected.

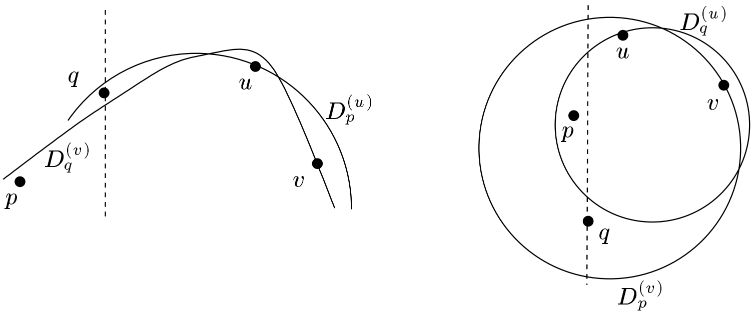

Figure 5: The proof of $p, q \in P_L$: (i) $D_q^{(v)}$ lies below q; (ii) the other situation.

Now we investigate the interplay between $D_q^{(u)}$ and $D_p^{(v)}$. Their top boundaries intersect between u and v. Since the top boundaries of two congruent disks intersect at most once, the top boundary of $D_p^{(v)}$ lies above the top boundary of $D_q^{(u)}$ at the x-coordinate of q, which in turn lies above q, as just proved. Now the bottom boundary of $D_q^{(u)}$ must lie above q, while the bottom boundary of $D_p^{(v)}$ must lie below (one disk must not contain q, the other has to). So either the boundaries of these two disks intersect twice to the left of q (which yields at least three intersections between these boundaries, including the one between u and v), or they do not intersect there, which implies that $D_p^{(v)}$ contains $D_q^{(u)}$ to the left of the vertical line through q; see Figure 5 (ii). But $p \in D_q^{(u)}$, which implies $p \in D_p^{(v)}$; a contradiction.

The cases of bottom boundaries, and of P_R, are symmetric, which concludes the proof of the theorem. □

Remark: Two top boundaries (or two bottom boundaries) of two sets K_p, K_q, for two points p, q in, say P_T, may also interact in somewhat more involved manners. First, we can have a *weak crossing* between two such top boundaries, in which the two boundaries have an overlapping portion, so that the top portion of ∂K_p lies below the top portion of ∂K_q to the left of the overlap, lies above it to the right of the overlap. See Figure 6(i) for an illustration. Another possibility is that these top boundaries meet twice, without crossing, and overlap between these two meeting points, as is illustrated in Figure 6(ii). Situations of the second type will not affect our algorithm, and we will have to exercise some care to accommodate situations of the first type in the algorithm; see Section 4 for details.

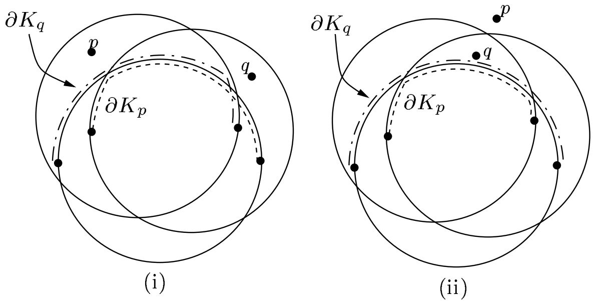

Figure 6: Counter-examples for Theorem 2.8 if we just consider intersection points (including noncrossings).

3 Complexity of $\mathcal{K}$

Let $\mathcal{D}$ be a set of n congruent disks and P a set of m points in the plane. Let $\mathcal{K}$ and U be the same as defined in the introduction. In this section we obtain a bound on the combinatorial complexity of $\mathcal{K}$, which is defined as follows. Let $V(\mathcal{D}, P)$ be the set of intersection points of disks in $\mathcal{D}$ that lie on the boundary of some K_p. Set $\kappa(\mathcal{D}, P) = |V(\mathcal{D}, P)|$ and $\kappa(n, m) = \max\ \kappa(\mathcal{D}, P)$, where the maximum is taken over all sets of n congruent disks and over all sets of m points in the plane. Note that if a vertex appears on the boundaries of several sets in $\mathcal{K}$, we count it only once. If we count the vertices with multiplicity, then $\kappa(n, m) \geq mn$ — take n congruent disks all of whose boundaries appear on their common intersection and choose m points in their common exterior. The main result of this section is the following theorem.

Theorem 3.1 $\kappa(n, m) = O(m^{2/3} n^{2/3} \log^{1/3} n + n \log n)$.

The proof of the theorem is based on the random-sampling technique, and proceeds along the same lines as the proof by Clarkson et al. [4] for the bound on the complexity of many faces in an arrangement of lines in the plane. We first prove a technical lemma and a weaker bound on $\kappa(n,m)$, and then prove the theorem.

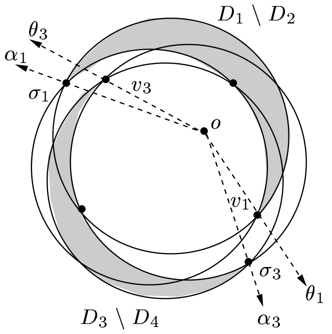

Figure 7: Proof of Lemma 3.2.

Lemma 3.2 *Let $D_1, D_2, \ldots, D_k$ be a set of congruent disks, all of whose boundaries appear on their common intersection $\mathcal{I}$. Assume that $\partial D_1, \partial D_2, \ldots, \partial D_k$ appear in this clockwise order along $\partial\mathcal{I}$. Then the sets $D_i \setminus D_{i+1}$, for $1 \le i \le k$ (where we put $D_{k+1} = D_1$), are pairwise disjoint, and the same holds for the sets $D_{i+1} \setminus D_i$.*

Proof: Suppose that there exist a pair of indices $1 \le i < j \le k$ so that $D_i \setminus D_{i+1}$ and $D_j \setminus D_{j+1}$ intersect. Note that j must be at least $i+2$, and i must be at least $j+2-k$; without loss of generality, we can assume that $i=1$, $j=3$, and $k \ge 4$. Consider the arrangement $\mathcal{A}(\{D_1, D_2, D_3, D_4\})$, and let $\mathcal{I}' = \bigcap_{i=1}^{4} D_i$. Let us assume that the origin, o, lies in the interior of $\mathcal{I}'$. Let v_i be the (unique) intersection point of ∂D_i and ∂D_{i+1} that appears on $\mathcal{I}'$, and let σ_i be the other intersection point of these circles. Let θ_i (resp. α_i) denote the orientation of v_i (resp. σ_i). We will regard ∂D_i as the graph of a univariate function $D_i(\theta)$ in polar coordinates. We will denote by (θ_1, θ_2) the (open) counter-clockwise circular interval from θ_1 to θ_2.

By construction, $D_1(\theta) > D_2(\theta)$ for $\theta \in (\theta_1, \alpha_1)$, and therefore $D_1 \setminus D_2$ is nonempty only for $\theta \in (\theta_1, \alpha_1)$. Similarly, $D_3(\theta) > D_4(\theta)$ for $\theta \in (\theta_3, \alpha_3)$ and $D_3 \setminus D_4$ is nonempty only for $\theta \in (\theta_3, \alpha_3)$. Since $\partial D_1, \partial D_2, \partial D_3$, and ∂D_4 appear in this counter-clockwise order along $\partial\mathcal{I}'$, it follows that $\theta_2 \in (\theta_1, \alpha_1)$ and $\theta_4 \in (\theta_3, \alpha_3)$, and that $\alpha_1 \in (\theta_2, \theta_4)$ and $\alpha_3 \in (\theta_4, \theta_2)$.

Let ξ be a point in $(D_1 \setminus D_2) \cap (D_3 \setminus D_4)$, and let η be the orientation of ξ. Then we have $\min\,\{D_3(\eta), D_1(\eta)\} > \max\,\{D_4(\eta), D_2(\eta)\}$. Moreover, $\eta \in (\theta_1, \alpha_1) \cap (\theta_3, \alpha_3)$. The order relationships noted at the preceding paragraph are easily seen to imply that only the following two cases can arise:

CASE (I) $\eta \in (\theta_3, \alpha_1) \subseteq (\theta_3, \theta_4)$:

CASE (II) $\eta \in (\theta_1, \alpha_3) \subseteq (\theta_1, \theta_2)$:

In case (i), ∂D_2 appears along $\partial(D_1 \cap D_2 \cap D_3)$ in at least two disjoint arcs—the arc with angular range (θ_1, θ_2) and another arc containing a point at orientation η (observe that ∂D_2 cannot appear on the boundary of this intersection in the angular range (θ_2, θ_3)). This however is impossible for congruent disks. Symmetrically, in case (ii), ∂D_4 appears along $\partial(D_1 \cap D_3 \cap D_4)$ in at least two disjoint arcs—the arc with angular range (θ_3, θ_4) and another arc containing a point at orientation η. These two contradictions complete the proof that the regions $D_i \setminus D_{i+1}$, for $1 \le i \le k$, are pairwise disjoint. A symmetric argument shows that the regions $D_i \setminus D_{i-1}$ are also pairwise disjoint. This completes the proof of the lemma. □

Lemma 3.3 *For $m, n \ge 1$, $\kappa(n,m) = O(m\sqrt{n} + n)$.*

Proof: Let $\mathcal{D}$ be a set of n congruent disks and P a set of m points in the plane. It suffices to prove that $\kappa(n,m) \le 2m^2 + n$. By partitioning P into $t = \lceil m/\sqrt{n} \rceil$ subsets, each of size at most $\sqrt{n}$, and observing that $\kappa(\mathcal{D}, P_i) = O(n)$ for each $i \le t$, the bound on $\kappa(n,m)$ can be improved to $O(m\sqrt{n} + n)$; see e.g. [4].

We partition $\mathcal{D}$ into maximal subsets $\mathcal{D}_1, \ldots, \mathcal{D}_k$, so that all disks within each $\mathcal{D}_i$ contain the same subset of P. Let V_i be the set of vertices on the boundary of $\bigcap \mathcal{D}_i$. Obviously $\sum_{i=1}^k |V_i| \le \sum_i |\mathcal{D}_i| = n$. We partition the vertices in $V(\mathcal{D}, P)$ into two subsets A and B, where A consists of those vertices v for which the two disks on whose boundaries v lies belong to the same $\mathcal{D}_i$, and B consists of those vertices whose two associated disks belong to different $\mathcal{D}_i$'s. In the first case, v is a vertex of $\bigcap \mathcal{D}_i$, therefore $|A| \le \sum_{i=1}^k |V_i| \le n$. We next bound $|B|$.

For each point $p \in P$, let $B_p \subseteq B$ be the set of vertices in B that appear on the boundary of K_p, and let $(D_1, D_2, \ldots)$ be the circular sequence of disks whose boundaries appear in this counter-clockwise order along ∂K_p. Suppose $v \in B_p$ is an intersection point of the circles bounding two consecutive disks in this sequence, say D_1 and D_2. Since D_1 and D_2 belong to two different subsets, the symmetric difference $D_1 \oplus D_2$ contains at least one point q of P. We charge v to q. If v' is another vertex of B_p, which is an intersection point of the circles bounding two other consecutive disks D_3 and D_4, then, by Lemma 3.2, $D_1 \setminus D_2$ and $D_3 \setminus D_4$ are disjoint, and the same holds for $D_2 \setminus D_1$ and $D_4 \setminus D_3$. Hence, each point $q \in P$ can be charged at most twice (once for lying in some $D_i \setminus D_{i+1}$ and once for lying in some $D_j \setminus D_{j-1}$), thereby implying that $|B_p| \le 2m$. Summing over all points $p \in P$, we obtain that $|B| \le 2m^2$, and therefore

$$\kappa(\mathcal{D}, P) \le 2m^2 + n,$$

as asserted. □

Proof of Theorem 3.1. Let $r \ge 1$ be a fixed parameter, to be specified later. We choose a random subset $R \subseteq \mathcal{D}$ of size r, where each subset of size r is chosen with equal probability, and consider the *vertical decomposition* $\mathcal{A}^*(R)$ of the arrangement $\mathcal{A}(R)$ [2, 4]. For each cell $\Delta \in \mathcal{A}^*(R)$, let $\mathcal{D}_\Delta \subseteq \mathcal{D}$ be the set of disks whose boundaries intersect Δ (including the edges of Δ), let $E_\Delta \subseteq \mathcal{D}$ be the set of disks that are disjoint from Δ, and let $P_\Delta \subseteq P$ be the set of points that lie in Δ (a point lying on an edge (or a vertex) of $\mathcal{A}^*(R)$ is assigned to one of the cells adjacent to it. Put $m_\Delta = |P_\Delta|$ and $n_\Delta = |\mathcal{D}_\Delta|$. We denote by $\mathcal{I}_\Delta$ the common intersection of the disks in E_Δ.

Let v be a vertex of K_p, for some $p \in P_\Delta$, not lying on a vertex of $\mathcal{A}^*(R)$. Suppose that v is an intersection point of the boundaries of two disks D and D'. Since, by definition, none of these disks can fully contain Δ, we can classify v into three categories:

(i) Both D and D' belong to $\mathcal{D}_\Delta$,

(ii) $D \in \mathcal{D}_\Delta$ and $D' \in E_\Delta$ (or vice versa), or

(iii) both $D, D' \in E_\Delta$.

A vertex of type (i) is also a vertex of $V(\mathcal{D}_\Delta, P_\Delta)$, so the number of such vertices is at most $\kappa(\mathcal{D}_\Delta, P_\Delta) \leq \kappa(n_\Delta, m_\Delta)$. Since $E_\Delta \subseteq \mathcal{D}_p$ for every $p \in P_\Delta$, every vertex of type (ii) lies on the boundary of $\mathcal{I}_\Delta$. The boundary of each disk in $\mathcal{D}_\Delta$ intersects $\mathcal{I}_\Delta$ in at most two points, so the number of type (ii) vertices is at most $2n_\Delta$. Summing over all cells, the number of type (i) and type (ii) vertices is $\sum_{\Delta \in \mathcal{A}^*(R)} O(n_\Delta + \kappa(n_\Delta, m_\Delta))$.

Finally, each vertex of type (iii) is a vertex of $\mathcal{I}_\Delta$. Hence, in order to bound the number of (distinct) vertices of type (iii), we need an upper bound on the total number of distinct vertices of all the $\mathcal{I}_\Delta$'s, over all cells $\Delta \in \mathcal{A}^*(R)$. Let G be the graph dual to $\mathcal{A}^*(R)$, that is, each node of G corresponds to a cell of $\mathcal{A}^*(R)$, and two nodes corresponding to cells Δ, Δ' are connected by an edge if the boundaries of Δ and Δ' overlap along (a portion of) an edge. We compute a path Π in G that visits each node of G at least once and at most 4 times. The existence of such a path was proved in [2]. We traverse Π, and at each node corresponding to a cell Δ, we maintain $\mathcal{I}_\Delta$, as follows. When we move from a node corresponding to Δ to the next node in Π, corresponding to a cell $\Delta' \in \mathcal{A}^*(R)$, we delete all the disks of $\mathcal{D}_\Delta \setminus \mathcal{D}_{\Delta'}$ from the intersection, and insert the disks of $\mathcal{D}_{\Delta'} \setminus \mathcal{D}_\Delta$ into the intersection. Since $E_\Delta \cup E_{\Delta'} \subseteq \mathcal{D}_\Delta \cup \mathcal{D}_{\Delta'}$, we now have the set $E_{\Delta'}$. We thus perform at most $n_\Delta + n_{\Delta'}$ insertions and deletions as we move from one node of Π to the next. Summing over all nodes of Π, we perform $O(\sum_\Delta n_\Delta)$ insertions and deletions. We wish to bound the number of distinct vertices that ever appear on the intersection, we traverse Π. Tamir [15] (see also Agarwal [1]) has shown that the number of distinct vertices that ever appear on the intersection of halfplanes, as we perform a mixed sequence of k insertions and deletions (starting at the empty set), is $O(k \log k)$. Using the same argument, one can show that the total number distinct vertices that ever appear on the intersection of a set of congruent disks, as we perform a sequence of k insertions and deletions (again, starting at the empty set), is also $O(k \log k)$. Hence, the number of distinct type (iii) vertices is $O(\sum_\Delta n_\Delta \log n)$.

Finally, each vertex of $\mathcal{A}^*(R)$ may be a vertex of $V(\mathcal{D}, P)$. Putting everything together, we obtain

$$\begin{aligned}\kappa(\mathcal{D}, P) &\leq \sum_{\Delta \in \mathcal{A}^*(R)} \kappa(\mathcal{D}_\Delta, P_\Delta) + O\Big(\sum_{\Delta \in \mathcal{A}^*(R)} n_\Delta \log n \Big) + O(r^2) \\ &= O\Big(\sum_{\Delta \in \mathcal{A}^*(R)} (m_\Delta \sqrt{n_\Delta} + n_\Delta \log n) \Big) + O(r^2).\end{aligned}$$

Since R is a random subset of $\mathcal{D}$, the random-sampling technique of Clarkson and Shor [5] implies that there exists R for which

$$\begin{aligned}\sum_{\Delta \in \mathcal{A}^*(R)} n_\Delta &= O(nr), \\ \sum_{\Delta \in \mathcal{A}^*(R)} m_\Delta \sqrt{n_\Delta} &= O\left(m \sqrt{\frac{n}{r}} \right).\end{aligned}$$

Substituting these values and choosing $r = \left\lceil \dfrac{m^{2/3}}{n^{1/3} \log^{2/3} n} \right\rceil$, we obtain

$$K(n, m) = O(m^{2/3} n^{2/3} \log^{1/3} n + n \log n).$$

This completes the proof of the theorem. □

An immediate consequence of Theorem 3.1 and Corollary 2.7 is the following.

Corollary 3.4 *The complexity of U is $O(n^{4/3} \log^{1/3} n)$.*

Unfortunately, we have not been able to exploit this bound to obtain an efficient algorithm, of comparable complexity, that computes U explicitly. The results of this section, although of interest in their own right, are not needed for the analysis of the algorithm that we present in the next two sections.

4 The Decision Algorithm

Let P be a set of n points in the plane, and let $\mathcal{D} = \{D(p, r) \mid p \in P\}$. Let $\mathcal{K}$ and U be the same as defined in the introduction. We describe an $O(n^{4/3} \log^4 n)$-time algorithm to determine whether $U \cap P \neq \emptyset$. Our strategy is to construct separately each of the four sub-unions $U_T = \bigcup_{p \in P_T} K_p$, $U_B = \bigcup_{p \in P_B} K_p$, $U_L = \bigcup_{p \in P_L} K_p$, and $U_R = \bigcup_{p \in P_R} K_p$, and to test whether any of them contains any point of P. We next describe in detail the construction of one such sub-union, say U_T. As already mentioned, we do not know how to compute these unions efficiently in an explicit manner (for instance, it may be too expensive to construct each K_p explicitly, for all $p \in P$). We will therefore represent the K_p's and their unions implicitly; this implicit representation will be sufficient to determine whether $U_T \cap P \neq \emptyset$.

4.1 Representation of $\mathcal{K}$ and of its union

For a subset $A \subseteq P$, let U^A denote the union $\bigcup_{p\in A} K_p$. For each connected component of ∂U^A, we store its concave vertices (points of crossing of the boundaries of two distinct K_p's) and the points that are locally x-extremal along ∂U^A. If two top (or two bottom) boundaries have a weak crossing along ∂U^A, as in the remark following Theorem 2.8, we store the left endpoint of the common overlap between these boundaries, and think of it as a 'weakly-concave' vertex of ∂U^A. A maximal portion γ of ∂U^A that does not contain any of these points is x-monotone and lies on the boundary of a single K_p (such a portion, γ, may overlap with the boundaries of many K_p's, but there is (at least) one point p such that γ is fully contained in ∂K_p). We refer to γ as a *boundary arc* of U^A. We maintain γ implicitly, by recording the point p for which $\gamma \subset \partial K_p$ and a bit that indicates whether γ is a portion of the top or bottom part of ∂K_p.

Next, to represent each K_p implicitly, we compute a family $\{\mathcal{D}^{(1)}, \ldots, \mathcal{D}^{(s)}\}$ of "canonical" subsets of $\mathcal{D}$ such that $\sum_{i=1}^{s} |\mathcal{D}^{(i)}| = O(n^{4/3}\log n)$, and such that for any $p \in P$, $\mathcal{D}_p$ can be represented as the union of $O(n^{1/3}\log n)$ canonical subsets. Let J_p be the set of indices of these canonical subsets (i.e., $\mathcal{D}_p = \bigcup_{i\in J_p} \mathcal{D}^{(i)}$). Katz and Sharir [9] have shown that the construction of such a family of canonical sets, and of the corresponding sets of indices $\{J_p\}_{p\in P}$, can be accomplished in time $O(n^{4/3}\log n)$. For each canonical subset $\mathcal{D}^{(j)}$, we compute the intersection $\mathcal{I}^{(j)} = \bigcap \mathcal{D}^{(j)}$ in $O(|\mathcal{D}^{(j)}|\log|\mathcal{D}^{(j)}|)$ time. We store the vertices of the top and bottom parts of $\mathcal{I}^{(j)}$ in separate lists, each sorted in increasing order of their x-coordinates. For each vertex $v \in \mathcal{I}^{(j)}$, we also store the disk whose boundary appears on $\partial\mathcal{I}^{(j)}$ immediately to its right. Finally, we store the vertices of all the $\mathcal{I}^{(j)}$'s in a single master list Λ, sorted in increasing order of their x-coordinates. The total time spent in computing this implicit representation of the K_p's is $\sum_{i=1}^{s} O(|\mathcal{D}^{(j)}|\log n) = O(n^{4/3}\log^2 n)$.

4.2 Basic operations on $\mathcal{K}$

In order to compute the implicit representation of U_T, we will need subroutines for the following basic operations on the boundaries of the sets in $\mathcal{K}$.

(S1) Leftmost and rightmost points. *Given a point p, compute the leftmost and the rightmost points of K_p.*

This requires computing the leftmost and the rightmost points of $\bigcap_{j\in J_p} \mathcal{I}^{(j)}$. Reichling [13] has shown that the leftmost (or rightmost) point of the intersection of k convex polygons with a total of n vertices can be computed in time $O(k\log^2 n)$. In fact, his algorithm can also be applied to a family of intersections of congruent disks. Applying Reichling's algorithm to the set $\{\mathcal{I}^{(j)} \mid j \in J_p\}$, we can compute the leftmost (or rightmost) point of K_p in time $O(|J_p|\log^2 n) = O(n^{1/3}\log^3 n)$.

(S2) Intersection points with a vertical line. *Given a vertical line ℓ and a point $p \in P$, determine the intersection points of ℓ with ∂K_p.*

For each $j \in J_p$, we can compute $\ell \cap \mathcal{I}^{(j)}$ in $O(\log n)$ time. Repeating this step for all $j \in J_p$, we obtain a collection of $O(n^{1/3} \log n)$ intervals along ℓ. We can compute the endpoints of the intersection of these intervals (or detect that the intersection is empty) in an additional $O(n^{1/3} \log n)$ time. Hence, the total running time of this procedure is $O(n^{1/3} \log^2 n)$. This procedure can also be used to determine whether a query point in the plane lies above, below, or on a boundary arc γ.

(S3) Crossing points of two top (or two bottom) boundary arcs. *Given two points $p, q \in P_T$ and an x-interval $[a, b]$ contained in the x-span of both K_p and K_q, determine whether the top boundaries of K_p and K_q cross in $[a, b]$. If so, return their crossing point. If they weakly cross in $[a, b]$, then return the leftmost endpoint of their common overlap in $[a, b]$. A similar operation is prescribed for the bottom boundaries of K_p and K_q.*

Let γ_p (resp. γ_q) be the portion of the top boundary of K_p (resp. K_q) in the x-interval $[a, b]$, and let $\ell_a : x = a$ and $\ell_b : x = b$. By computing the intersection points of ∂K_p and ∂K_q with ℓ_a, we can determine, in $O(n^{1/3} \log n)$ time, whether γ_p lies above or below γ_q at ℓ_a. Suppose γ_p lies above γ_q at ℓ_a. We repeat the same procedure at ℓ_b. Note that, by Theorem 2.8, γ_p lies below γ_q at ℓ_b too if and only if γ_p and γ_q cross (or weakly cross). If they do cross, then, by performing a binary search over the points stored in the master list Λ, we obtain two consecutive vertices $\alpha, \beta \in \Lambda$ so that the crossing point (or, in case of weak crossing, the leftmost point of the common overlap) of γ_p and γ_q lies in the x-interval I between α and β. Each step of the binary search involves determining whether γ_p lies above γ_q at a vertical line $\ell : x = x_0$, for some $x_0 \in \Lambda$, and is performed using subroutine (S2). Hence the total cost of the binary search is $O(n^{1/3} \log^3 n)$. The top boundary of each $\mathcal{I}^{(j)}$, for $j \in J_p$, is composed of a single circular arc in the x-interval I. We therefore collect the $O(n^{1/3} \log n)$ corresponding disks, and compute, in $O(n^{1/3} \log^2 n)$ time, the top boundary $\hat{\gamma}_p$ of their intersection within I. Similarly, we compute $\hat{\gamma}_q$, the top boundary of K_q over I. We can now compute the crossing point (or the leftmost point of the common overlap) of $\hat{\gamma}_p$ and $\hat{\gamma}_q$ in an additional $O(n^{1/3} \log n)$ time, by merging the lists of vertices of $\hat{\gamma}_p$ and $\hat{\gamma}_q$, and by inspecting each 'atomic' interval formed by this merge. The total time spent by this procedure is thus $O(n^{1/3} \log^3 n)$. A symmetric procedure can compute the unique crossing point (or the leftmost point of the common overlap of a weak crossing) of the bottom boundaries of K_p and K_q, within the same time bound.

(S4) Crossing points of a top boundary arc and a bottom boundary arc. *Given two points $p, q \in P_T$ and an x-interval $[a, b]$ contained in the x-span of both K_p and K_q, determine whether the top boundary of K_p crosses the bottom boundary of K_q in the interval*

$[a, b]$. If so, return their crossing point(s).

Let γ_p be the portion of the top boundary of K_p lying in the interval $[a, b]$, and let γ_q be the portion of the bottom boundary of K_q lying in the interval $[a, b]$. Note that, by convexity, γ_p and γ_q can cross in at most two points. By comparing the y-coordinates of the endpoints of γ_p, γ_q, using an appropriate variant of subroutine (S2), we can determine whether they cross exactly once. In this case, we can determine their unique crossing point (or the leftmost point of overlap of a weak crossing) using an appropriate variant of subroutine (S3). Suppose that we determine that γ_p and γ_q cross at zero or two points, and that γ_q lies above γ_p at the vertical line $x = a$ (and also at $x = b$); if γ_q lies below γ_p at these points, the arcs do not intersect. If we regard γ_p and γ_q as graphs of univariate, partially defined functions $\gamma_p(x), \gamma_q(x)$, respectively, then $\Delta\gamma(x) = \gamma_q(x) - \gamma_p(x)$ is a convex function. Therefore, by a binary search through Λ, each step of which requires determining the intersection points of a vertical line with γ_p and γ_q, we can determine, in overall $O(n^{1/3} \log^3 n)$ time, the unique x-value x_0 at which $\Delta\gamma$ attains its minimum. If $\Delta\gamma(x_0) > 0$, then γ_p and γ_q do not intersect. If $\Delta\gamma(x_0) = 0$, then the minimum of $\Delta\gamma$ is the unique point of intersection (actually, of tangency) of γ_p and γ_q. If $\Delta\gamma(x_0) < 0$, then γ_p and γ_q have two crossings, one of which lies in the interval $[a, x_0]$ and the other lies in the interval $[x_0, b]$. Now we can compute both crossing points in $O(n^{1/3} \log^3 n)$ time, using an appropriate variant of subroutine (S3).

4.3 Computing U_T

We now describe an algorithm for computing the implicit representation of U_T described above, and for determining whether $U_T \cap P \neq \emptyset$. We first compute, using subroutine (S1), the leftmost and rightmost points, l_p, r_p, of each K_p, for $p \in P$. This, combined with calls to subroutine (S2), allows us to compute the sets P_T, P_B, P_L, and P_R, in overall $O(n^{4/3} \log^3 n)$ time. Next, we compute U_T, using a divide-and-conquer algorithm. If $|P_T| = 1$, then $U_T = K_p$, where p is the only point in P_T. In this case, we output ∂U_T as consisting of two boundary arcs, both connecting l_p and r_p, where the top (resp. bottom) arc is the top (resp. bottom) boundary of K_p. If $|P_T| > 1$, we partition P_T into two subsets P_T^1 and P_T^2, each of size at most $\lceil |P_T|/2 \rceil$. We recursively compute $U_T^1 = \bigcup_{p \in P_T^1} K_p$ and $U_T^2 = \bigcup_{p \in P_T^2} K_p$, and then compute $U_T = U_T^1 \cup U_T^2$, using a sweep-line algorithm. This 'merge' step computes the implicit representation of U_T from those of U_T^1, U_T^2, which are output by the respective recursive calls.

The sweep line scans the plane from left to right, stopping at the concave vertices and the locally x-extremal points of U_T^1, U_T^2, and U_T. By Corollary 2.7, the number of such "event points" is only $O(n)$. The algorithm maintains those arcs of U_T^1, U_T^2 that currently intersect the sweep line in a height-balanced tree T, sorted in the increasing order of the y-coordinates of their intersection points with the line. At each event point, the algorithm inserts a new arc, deletes an arc, or swaps two adjacent arcs in the tree T. In order to insert a new arc into T, the algorithm has to perform $O(\log n)$ comparisons of the following form: given a point

q and a boundary arc γ, determine whether q lies above, below, or on γ. Using subroutine (S2), such a comparison can be performed in $O(n^{1/3}\log^2 n)$ time. The time spent in inserting an arc is thus $O(n^{1/3}\log^3 n)$. The deletion of an arc follows a standard deletion procedure of a height-balanced tree. After having inserted or deleted an arc, we obtain the new $O(1)$ adjacent pairs of arcs in T, compute their (leftmost) intersection points to the right of the current sweep line, and insert them into the event queue. We thus need to perform $O(1)$ calls to the subroutines (S3) and (S4), each of which takes $O(n^{1/3}\log^3 n)$ time. Omitting all the other straightforward and standard details of the sweep-line algorithm, we conclude that the algorithm spends $O(n^{1/3}\log^3 n)$ time at each event point, therefore the total time spent by the sweep-line algorithm is $O(n^{4/3}\log^3 n)$. The overall time spent in computing the implicit representation of U_T is thus $O(n^{4/3}\log^4 n)$. (Note that the computation of the sets $\mathcal{D}^{(j)}$, $\mathcal{I}^{(j)}$, and J_p is performed only once, before starting the recursive construction of U_T.)

We next have to determine whether $U_T \cap P \neq \emptyset$. This can easily be done, at no increase in the asymptotic running time, during the topmost sweep of the recursion, in which the entire U_T is constructed. We include the points of P as additional event points of the line sweep. Whenever we encounter a point $p \in P$, we find the arc γ of U_T lying immediately above p. If γ is a portion of the top boundary of some K_q, then $p \in U_T$. Moreover, $D(p,r) \cup D(q,r)$ covers all points of P, so we can return p, q as the solution to the fixed-size problem; if γ is a portion of the bottom boundary of some K_q, then $p \notin U_T$. The arc γ can be determined by searching the tree with p, where each step of the search determines whether p lies above, below, or on an arc γ'. Since each such step can be performed in $O(n^{1/3}\log^2 n)$ time, using subroutine (S2), we can determine in $O(n^{1/3}\log^3 n)$ time whether $p \in U_T$. Summing this cost over all points $p \in P$, the total time spent by this stage is $O(n^{4/3}\log^3 n)$.

We now construct and search in U_B, U_L and U_R, using the algorithm just described. If at least one of these unions contains a point of P then we have found, and can output, two points $p, q \in P$ such that $D(p,r) \cup D(q,r)$ covers P; otherwise, no two such points exist. The overall running time is $O(n^{4/3}\log^4 n)$. Hence, we obtain the following result.

Theorem 4.1 *Given a set P of n points in the plane and a real value $r > 0$, we can determine, in $O(n^{4/3}\log^4 n)$ time, whether there are two points $p, q \in P$ so that $P \subset D(p,r) \cup D(q,r)$. If so, we can also find such a pair within the same time bound.*

Remark. The running time can be improved by a logarithmic factor by exploiting the special structure of canonical subsets and using fractional cascading. We omit the details of this slight improvement from here.

5 The Overall Algorithm

The overall algorithm for the discrete 2-center problem proceeds as follows. We note that the optimum radius r^* is a distance between two points of P, so we will run a binary search over these $\binom{n}{2}$ distances, using the fixed-size decision procedure given in the preceding section to determine whether the optimum r^* is larger than, smaller than, or equal to a distance r. Note that $r < r*$ if $U = \bigcup_{p \in P} K_p$, as defined in the introduction, contains a point of P, that $r > r^*$ if the interior of U contains a point of P, and $r = r^*$ if the interior of U does not contain a point but its boundary contains a point of P.

Of course, running this binary search explicitly will require quadratic time, so we use instead the distance-selection algorithm of [9] (see also [2]), which computes the k-th smallest distance in a set of n points in the plane in time $O(n^{4/3} \log^2 n)$. Since we need to invoke this procedure, and also the fixed-size decision procedure, only $O(\log n)$ times, the overall running time of the algorithm is $O(n^{4/3} \log^5 n)$.

Theorem 5.1 *The discrete 2-center problem for a set of n points in the plane can be solved in time $O(n^{4/3} \log^5 n)$.*

6 Conclusion

We have presented an $O(n^{4/3} \log^5 n)$-time algorithm for the planar discrete 2-center problem. By remark following Theorem 4.1, the running time can be improved by a logarithmic factor.

Our decision algorithm relies heavily on the properties of the combinatorial structure of $\mathcal{K}$ that we have proved in Section 2. Although we have shown in Section 3 that the complexity of U is roughly $n^{4/3}$, we do not have an algorithm with comparable running time that computes U explicitly. There are, in fact, a number of substructures in an arrangement of a set of congruent disks, whose worst-case complexity has the same asymptotic upper bound as that of the corresponding structure in an arrangement of lines. For example, the number of incidences between points and congruent disks, the complexity of many faces, and the complexity of U (one can define a structure analogous to U for a set of halfplanes in the plane). However, unlike the case of lines, no efficient algorithm is known for computing most of these substructures. It is an open problem whether m distinct faces in an arrangement of n congruent disks can be computed in time close to $m^{2/3}n^{2/3} + n$. A solution to any of these problems will most likely offer insights for developing a simpler algorithm (still with running time close to $n^{4/3}$) for the discrete 2-center problem. A more challenging open problem is whether a near-linear-time algorithm can be developed for the discrete 2-center problem.

Acknowledgments

The authors thank Arie Tamir for helpful discussions and the referee for several useful comments.

References

[1] P.K. Agarwal, On stabbing lines for convex polyhedra in 3D, *Comput. Geom: Theory Appls.* 4 (1994), 177–189.

[2] P.K. Agarwal, B. Aronov, M. Sharir and S. Suri, Selecting distances in the plane, *Algorithmica* 9 (1993), 495–514.

[3] P.K. Agarwal and M. Sharir, Planar geometric location problems, *Algorithmica* 11 (1994), 185–195.

[4] K. Clarkson, H. Edelsbrunner, L. Guibas, M. Sharir, and E. Welzl, Combinatorial complexity bounds for arrangements of curves and spheres, *Discrete Comput. Geom.* 5 (1990), 99–160.

[5] K. Clarkson and P. W. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.* 4 (1989), 387–421.

[6] D. Eppstein, Faster construction of planar 2-centers, *Proc. 8th ACM-SIAM Symp. on Discrete Algorithms*, 1997, 131–138.

[7] J. Hershberger and S. Suri, Finding tailored partitions, *J. Algorithms* 12 (1991), 431–463.

[8] J. Jaromczyk and M. Kowaluk, An efficient algorithm for the Euclidean two-center problem, *Proc. 10th ACM Symp. on Computational Geometry*, 1994, 303–311.

[9] M. Katz and M. Sharir, An expander-based approach to geometric optimization, *SIAM J. Comput.* 26 (1997), 1384–1408.

[10] K. Kedem, R. Livne, J. Pach and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete Comput. Geom.* 1 (1986), 59–71.

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

[12] N. Megiddo and K. Supowit, On the complexity of some common geometric location problems, *SIAM J. Comput.* 13 (1984), 1182–1196.

[13] M. Reichling, On the detection of a common intersection of k convex objects in the plane, *Inform. Process. Lett.* 29 (1988), 25–29.

[14] M. Sharir, A near-linear algorithm for the planar two-center problem, *Discrete Comput. Geom.* (in press).

[15] A. Tamir, Improved complexity bounds for center location problem on networks by using dynamic data structures, *SIAM J. Discrete Math.* 1 (1988), 377-396.