

Optimal Algorithm for the Planar Two-Center Problem

Received May 1, 2024

Revised Sep 30, 2024

Accepted Oct 5, 2024

Published Nov 5, 2024

Key words and phrases

Two centers, disk coverage, circular hulls, r -coverage

Kyungjin Cho^a ✉

Eunjin Oh^a ✉

Haitao Wang^b ✉

Jie Xue^c ✉

^a Pohang University of Science and Technology, Korea

^b University of Utah, USA

^c New York University Shanghai, China

ABSTRACT. We study a fundamental problem in computational geometry, the *planar two-center* problem. In this problem, the input is a set S of n points in the plane and the goal is to find two smallest congruent disks whose union contains all points of S . A longstanding open problem has been to obtain an $O(n \log n)$ -time algorithm for the planar two-center problem, matching the $\Omega(n \log n)$ lower bound given by Eppstein [SODA'97]. Towards this, researchers have made many efforts over decades. The previous best algorithm, given by Wang [SoCG'20], solves the problem in $O(n \log^2 n)$ time. In this paper, we present an $O(n \log n)$ -time (deterministic) algorithm for the planar two-center problem, which completely resolves this open problem.

1. Introduction

Given a set S of n points in the plane, the *planar k -center* problem asks for k smallest congruent disks in the plane whose union contains S . Planar k -center is NP-hard if k is given as part of the input. The best known algorithm for the problem takes $n^{O(\sqrt{k})}$ time [19].

In the special case where $k = 1$, the problem can be solved in linear time [7, 12, 24, 23], and this result not only holds in the plane but also in any fixed dimension. When $k = 2$, the problem becomes substantially more challenging. First of all, unlike the one-center case, the two-center problem requires $\Omega(n \log n)$ time to solve in the algebraic decision tree model, as shown by Eppstein [15] by a reduction from the max-gap problem. A longstanding open problem has been

Invited paper from the 40th International Symposium on Computational Geometry (SoCG 2024). Kyungjin Cho and Eunjin Oh were supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No.RS-2023-00209069). Haitao Wang was supported in part by NSF under Grant CCF-2300356.

to prove a matching upper bound for the problem, i.e., obtain an $O(n \log n)$ -time algorithm. Towards this, researchers have made many efforts over decades. The first published result on the planar two-center problem appeared in 1991 by Hershberger and Suri [18], who considered the decision version of the problem only and presented an $O(n^2 \log n)$ -time algorithm. This algorithm was subsequently improved to $O(n^2)$ time by Hershberger [17]. Using the decision algorithm in [17] and the parametric search technique [22], Agarwal and Sharir [2] presented an $O(n^2 \log^3 n)$ -time algorithm for the (original) two-center problem. The running time was slightly improved in a sequence of works [14, 20, 21]. A major breakthrough was achieved by Sharir [27], who proposed an $O(n \log^9 n)$ -time algorithm. This is the first algorithm that solves the problem in near-linear time (in fact, even the first algorithm with subquadratic runtime). Eppstein [15] and Chan [6] further improved this result, obtaining an $O(n \log^2 n)$ -time randomized algorithm and an $O(n \log^2 n \log \log n)$ -time deterministic algorithm, respectively. Since then, no progress had been made for over two decades, until recently in SoCG 2020 Wang [29] gave an $O(n \log^2 n)$ -time deterministic algorithm¹.

Choi and Ahn [8] considered a special case of the two-center problem where the two disks in an optimal solution overlap and a point in their intersection is given in the input (we refer to it as the *anchored case*); they gave an $O(n \log n)$ -time algorithm by improving the parametric search scheme in [29], where an algorithm of $O(n \log n \log \log n)$ time is provided for this case. In addition, Choi and Ahn's technique also solves in $O(n \log n)$ time the convex case in which the points of S are in a convex position; this again improves the $O(n \log n \log \log n)$ -time result in [29]. In spite of this progress on special cases, prior to this work, Wang's $O(n \log^2 n)$ -time (deterministic) algorithm [29] remains the best known result for the general planar two-center problem.

Our result. In this paper, we resolve this longstanding open problem completely, by presenting an $O(n \log n)$ -time deterministic algorithm for the planar two-center problem. Remarkably, in terms of running time, this is the first improvement since the $O(n \log^2 n)$ -time (randomized) algorithm of Eppstein [15] in 1997, although Wang [29] later presented a deterministic version with the same time complexity. Our result follows from a new decision algorithm for the problem which runs in $O(n)$ time after $O(n \log n)$ -time preprocessing, summarized in Theorem 1.1 (unless otherwise stated, all time complexities in the following discussion are deterministic.) Along with known techniques in the previous works, such a decision algorithm immediately leads to an $O(n \log n)$ -time algorithm for the two-center problem, as elucidated in the proof of Corollary 1.2.

¹ Independently, Tan and Jiang [28] also claimed a simple $O(n \log^2 n)$ -time algorithm, but unfortunately that algorithm was later found to be incorrect [29].

THEOREM 1.1. *Let S be a set of n points in the plane. After a preprocessing step on S in $O(n \log n)$ time, for any given $r \geq 0$, one can compute in $O(n)$ time two congruent disks of radius r in the plane that together covers S , or decide the nonexistence of two such disks.*

COROLLARY 1.2. *The planar two-center problem can be solved deterministically in $O(n \log n)$ time.*

PROOF. Consider an instance S of the two-center problem, which is a set of n points in the plane. Let opt be the radius of the two disks in an optimal solution for S , and δ be the smallest distance between the two disk centers in any optimal solution for S .

If $\delta \leq \frac{3}{2} \cdot \text{opt}$, then the algorithm of Choi and Ahn [8] can compute an optimal solution for S in $O(n \log n)$ time (in this case it is possible to sample $O(1)$ points in the plane (not necessarily from S) so that at least one point is in the intersection of the two optimal disks; thus the problem reduces to the anchored case).

For the case $\delta > \frac{3}{2} \cdot \text{opt}$, Eppstein [15] showed that if one can decide whether $r \geq \text{opt}$ for any given r in $T(n)$ time after $T_0(n)$ -time preprocessing on S , then one can compute an optimal solution for S in $O(T(n) \cdot \log n + T_0(n))$ time by Cole's parametric search [10]. The algorithm of Theorem 1.1 solves this decision problem with $T(n) = O(n)$ and $T_0(n) = O(n \log n)$. Therefore, we can compute an optimal solution for S in $O(n \log n)$ time.

Combining the two cases finally gives us an $O(n \log n)$ -time algorithm for the planar two-center problem. Note that since we do not know which case holds for S , we run the algorithms for both cases and return the better solution. ■

To prove Theorem 1.1, which is the main goal of this paper, we design an algorithm that is relatively simple. However, the correctness analysis of the algorithm is technical, entailing a nontrivial combination of several novel and intriguing insights into the problem, along with known observations in the literature.

Other related work. A number of variants of the two-center problem have also attracted much attention in the literature, e.g., [1, 3, 4, 11, 16, 25, 26, 31]. For example, if the centers of the two disks are required to be in S , the problem is known as the *discrete* two-center problem. Agarwal, Sharir, and Welzl [3] solved the problem in $O(n^{4/3} \log^5 n)$ time; the logarithmic factor in the runtime was slightly improved recently by Wang [30]. An outlier version of the problem was studied in [1]. Other variants include those involving obstacles [16, 25, 26], the bichromatic version [4, 31], the kinetic version [11], etc. It should be noted that although the general planar k -center problem is NP-hard [19], Choi, Lee, and Ahn recently gave a polynomial-time algorithm for the case where the points are given in convex position [9].

Outline. The rest of the paper is organized as follows. In Section 2, we define some basic notions that will be used throughout the paper. Section 3 introduces a new concept, called

r -coverage, which serves an important role in our algorithm. To the best of our knowledge, we are not aware of any previous work that used this concept to tackle the two-center problem before. We prove several properties about the r -coverage, which may be interesting in their own right. Our algorithm for Theorem 1.1 is presented in Section 4, while its correctness is proved in Section 5. Finally, in Section 6, we conclude the paper and pose some open questions.

2. Preliminaries

Basic notions. For two points p and q in the plane, we denote by \overline{pq} the line segment connecting p and q , and denote by $|pq|$ the distance between p and q (throughout the paper, “distance” always refers to the Euclidean distance). For a compact region R in the plane, we use ∂R to denote its boundary. For a disk D in the plane, let $\text{ctr}(D)$ and $\text{rad}(D)$ denote the center and the radius of D , respectively. For a point $p \in \partial D$ where D is a disk, the *antipodal point* of p on ∂D refers to the (unique) point of ∂D that is farthest from p ; we often use \hat{p} to denote it. For two disks D and D' in the plane, denote by $\text{dist}(D, D')$ the distance between the two centers $\text{ctr}(D)$ and $\text{ctr}(D')$. For a number $r \geq 0$, we say that a set Q of points in the plane is r -coverable if there exists a radius- r disk that contains all points of Q .

Circular hulls. The r -circular hull of a point set Q in the plane with respect to a value $r > 0$ is defined as the common intersection of all radius- r disks that contain Q if such disks exist [13, 18]. We use $\alpha_r(Q)$ to denote the r -circular hull of Q . See Figure 1(a). Note that Q is r -coverable if and only if $\alpha_r(Q)$ exists and is not empty.

Wang [29] considered the problem of dynamically tracking the circular hull of a set Q of points in the plane under *monotone* insertions, that is, each inserted point is always to the right of all points in the current Q . The following result is obtained in [29].

LEMMA 2.1 (Theorem 6.7 in [29]). *Let Q be a (dynamic) set of points in the plane which is initially empty. For any fixed number $r > 0$, there exists a data structure that maintains $\alpha_r(Q)$ in $O(1)$ amortized update time under monotone insertions to Q .*

More specifically, $\alpha_r(Q)$ is maintained in a linked list so that it can be output in time linear in its size [29].

Solutions for the two-center problem. For any $r > 0$, an r -solution for a planar two-center instance S is a set of two radius- r disks in the plane whose union covers S ; a *solution* for S refers to an r -solution with some $r \geq 0$. We denote by $\text{opt}(S)$ the minimum r such that S has an r -solution. An r -solution $\{D, D'\}$ for S is *tight* if $\text{dist}(D, D')$ is minimized (among all r -solutions). An r -solution $\{D, D'\}$ for S is p -anchored for a point $p \in \mathbb{R}^2$ if $p \in D \cap D'$. The *anchored two-center problem* takes as input a set S of points in the plane and a point $p \in \mathbb{R}^2$, and asks for a




Figure 1. (a) Illustrating the r -circular hull $\alpha_r(Q)$, which is bounded by the solid arcs. The radius of the dashed circle is r . (b) Illustrating the r -coverage $CR_r(Q)$ (bounded by the outer cycle of solid arcs) for the set of five black points, while the inner cycle of solid arcs bounds $\alpha_r(Q)$. The larger dashed arc is of radius $2r$ while the smaller dashed circle is of radius r .

p -anchored r -solution for S with the minimum r . As discussed in Section 1, Choi and Ahn [8] presented an $O(n \log n)$ -time algorithm for the anchored problem.

LEMMA 2.2 (Choi and Ahn [8]). *The anchored planar two-center problem can be solved by a deterministic algorithm which runs in $O(n \log n)$ time.*

3. r -Coverage

We introduce a new concept, called r -coverage, that is critical to our algorithm. To the best of our knowledge, we are not aware of any previous work that used it to tackle the two-center problem before. We prove several properties about it, which are needed for our algorithm. We believe these properties are interesting in their own right and will find applications elsewhere.

DEFINITION 3.1. For a point set Q in the plane and a value $r > 0$, we define the r -coverage of Q as the union of all radius- r closed disks containing Q . We use $CR_r(Q)$ to denote the r -coverage of Q .

In the following, we prove several properties about $CR_r(Q)$ based on an interesting relationship with $\alpha_r(Q)$. More specifically, we will show that $CR_r(Q)$ is convex and $\partial CR_r(Q)$ consists of circular arcs of radii r and $2r$ alternately. Roughly speaking, these arcs correspond to the vertices and arcs of $\alpha_r(Q)$ in the following way: Each arc of radius r of $\partial CR_r(Q)$ is an “antipodal arc” of an arc of $\alpha_r(Q)$ and each arc of radius $2r$ of $\partial CR_r(Q)$ has a vertex of $\alpha_r(Q)$ as its center (see Figure 1(b)); further, the cyclic order of the vertices and arcs of $\alpha_r(Q)$ on $\partial \alpha_r(Q)$ is consistent with the order of their corresponding arcs of $CR_r(Q)$ on $\partial CR_r(Q)$. As such, given $\alpha_r(Q)$, $CR_r(Q)$ can be constructed in time linear in the number of vertices of $\alpha_r(Q)$. We formally prove these properties below.

We assume that Q can be covered by a radius- r disk since otherwise $C\mathcal{R}_r(Q)$ would not exist. For ease of exposition, we assume that the distance of any two points of Q is strictly less than $2r$ (otherwise, there is a unique radius- r disk covering Q , which itself forms $C\mathcal{R}_r(Q)$). We start with the following lemma.

LEMMA 3.2. *For a point p on $\partial\alpha_r(Q)$, let D be a radius- r disk containing Q with $p \in \partial D$. Then, the antipodal point of p on ∂D must lie on $\partial C\mathcal{R}_r(Q)$.*

PROOF. First of all, by the definition of $\alpha_r(Q)$, such a disk D must exist for $p \in \partial\alpha_r(Q)$. In this proof, we let q be the antipodal point of p on ∂D , and show that q lies on $\partial C\mathcal{R}_r(Q)$.

For contradiction, we assume that q is not on $\partial C\mathcal{R}_r(Q)$. Then, q must be in the interior of $C\mathcal{R}_r(Q)$, since q lies in $C\mathcal{R}_r(Q)$ by the definition of $C\mathcal{R}_r(Q)$. Consider the ray ρ from p to q . Since q is in the interior of $C\mathcal{R}_r(Q)$, ρ contains a point $q' \in C\mathcal{R}_r(Q)$ such that q is in the interior of $\overline{pq'}$. Since q is the antipodal point of p on ∂D and the radius of D is r , we have $|pq| = 2r$. Hence, $|pq'| > 2r$. Since $q' \in C\mathcal{R}_r(Q)$, there exists a radius- r disk D' containing Q and q' . Since D' contains Q , D' contains $\alpha_r(Q)$ as well. As $p \in \alpha_r(Q)$, we obtain that D' contains both p and q' . But this is impossible as the radius of D' is r while $|pq'| > 2r$. ■

For any point $p \in \partial\alpha_r(Q)$, define $\tau(p)$ as the set of points q satisfying the conditions of Lemma 3.2, i.e., there is a radius- r disk D containing Q and having both p and q on its boundary as antipodal points. By Lemma 3.2, $\tau(p) \subseteq \partial C\mathcal{R}_r(Q)$.

Let e be an arc of $\alpha_r(Q)$ with two endpoints u and v . By slightly abusing the notation, define $\tau(e) = \cup_{p \in e} \tau(p)$. Let D be the unique radius- r disk whose boundary contains e . Note that D contains $\alpha_r(Q)$ and thus contains Q . Therefore, $\tau(e)$ consists of the antipodal points of all points of e on ∂D . We call $\tau(e)$ the *antipodal arc* of e . Note that $\tau(e)$ and e are congruent to each other (and thus they have the same length). Since we have assumed that the distance of any two points of Q is less than $2r$, $\tau(e)$ does not intersect e . For later reference purposes, we have the following corollary, following directly from Lemma 3.2.

COROLLARY 3.3. *For an open arc e of $\alpha_r(Q)$, $\tau(e)$ is an arc lying on $\partial C\mathcal{R}_r(Q)$.*

Let u be a vertex of $\alpha_r(Q)$. We now give a characterization of $\tau(u)$. Let v and w be the counterclockwise and clockwise neighboring vertices of u on $\alpha_r(Q)$, respectively. See Figure 2. Let $e(u, v)$ be the open arc of $\alpha_r(Q)$ connecting u and v . Let \hat{u}_v be the antipodal point of u on the circle containing $e(u, v)$. Define $e(u, w)$ and \hat{u}_w similarly. Observe that $\tau(u)$ is the arc from \hat{u}_v clockwise to \hat{u}_w on the radius- $(2r)$ circle centered at u . Also observe that the clockwise endpoint of $\tau(u)$, which is \hat{u}_w , is actually the counterclockwise endpoint of $\tau(e(u, w))$; similarly, the counterclockwise endpoint of $\tau(u)$, which is \hat{u}_v , is the clockwise endpoint of $\tau(e(u, v))$. Hence, if we consider the arcs e and vertices u of $\alpha_r(Q)$ in cyclic order on $\partial\alpha_r(Q)$, their corresponding arcs $\tau(\cdot)$ form a closed cycle, denoted by $\Lambda_r(Q)$. By Lemma 3.2, $\Lambda_r(Q) \subseteq \partial C\mathcal{R}_r(Q)$. In the following,




Figure 2. The two solid arcs connecting v, u, w are $e(u, v)$ and $e(u, w)$, respectively. The red dashed arc is $\tau(u)$. The two dotted circles contain $e(u, w)$ and $e(u, v)$, respectively. The two blue solid arcs on the two circles are $\tau(e(u, v))$ and $\tau(e(u, w))$, respectively, and they share endpoints \hat{u}_v and \hat{u}_w with $\tau(u)$.

we show that $\partial C\mathcal{R}_r(Q) \subseteq \Lambda_r(Q)$ (and thus $\partial C\mathcal{R}_r(Q) = \Lambda_r(Q)$). This is an almost immediate consequence of the following lemma.

LEMMA 3.4. *Let q be a point on $\partial C\mathcal{R}_r(Q)$. Then there is a unique disk D of radius r covering Q with $q \in \partial D$. Moreover, the antipodal point p of q on ∂D lies on $\partial\alpha_r(Q)$.*

PROOF. Consider a disk D of radius r covering Q and $q \in \partial D$. Such a disk always exists by the definition of $C\mathcal{R}_r(Q)$. Since D covers Q , we have $\alpha_r(Q) \subseteq D$.

We first show that the antipodal point p of q on ∂D lies on $\partial\alpha_r(Q)$. Assume to the contrary that this is not the case. Then, p must lie outside of $\alpha_r(Q)$ since $\alpha_r(Q) \subseteq D$. Thus there is another disk D' of radius r containing Q but not containing p . As such, $Q \subseteq D \cap D'$ and $p \notin D \cap D'$. See Figure 3. Since $Q \subseteq D'$, we have $D' \subseteq C\mathcal{R}_r(Q)$. As $q \in \partial C\mathcal{R}_r(Q)$, q cannot be in the interior of D' . Further, since p and q are antipodal points on ∂D , it is possible to slightly change D to a new position to obtain another radius- r disk D'' such that D'' contains $D \cap D'$ (and thus contains Q) and also contains q in its interior (e.g., first slightly rotate D around q towards the interior of D' , and then slightly translate the disk along the direction from the center toward q). Since $Q \subseteq D''$, it holds that $D'' \subseteq C\mathcal{R}_r(Q)$. As q is in the interior of D'' , q is also in the interior of $C\mathcal{R}_r(Q)$. But this leads to a contradiction since q lies on $\partial C\mathcal{R}_r(Q)$.

We proceed to show that such a disk D is unique. Assume there is another radius- r disk D' with $Q \subseteq D'$ and $q \in \partial D'$. Then the antipodal point p of q on ∂D lies outside of D' . On the other hand, since p lies on $\partial\alpha_r(Q)$ (which has been proved above), p must be contained in D' because D' contains Q (and thus contains $\alpha_r(Q)$). This leads to a contradiction. ■

By Lemma 3.4, for any point $q \in \partial C\mathcal{R}_r(Q)$, q must be in $\tau(p)$ for a point $p \in \partial\alpha_r(Q)$. As such, $q \in \Lambda_r(Q)$. This also proves that $\partial C\mathcal{R}_r(Q) \subseteq \Lambda_r(Q)$. We thus obtain $\partial C\mathcal{R}_r(Q) = \Lambda_r(Q)$. We




Figure 3. Illustrating D , D' , and D'' (the red dashed circle). Q is contained in the gray region, which is $D \cap D'$.

summarize these in the following corollary (note that the arc $\tau(e)$ for each arc e and $\tau(u)$ for each vertex u of $\alpha_r(Q)$ can be computed in $O(1)$ time).

COROLLARY 3.5. *We have the following relations between $C\mathcal{R}_r(Q)$ and $\alpha_r(Q)$.*

- (1) $\partial C\mathcal{R}_r(Q)$ is the union of the arcs $\tau(u)$ and $\tau(e)$ for all vertices u and arcs e of $\alpha_r(Q)$ following their cyclical order on $\partial\alpha_r(Q)$.
- (2) Given $\alpha_r(Q)$, $C\mathcal{R}_r(Q)$ can be computed in time linear in the number of vertices of $\alpha_r(Q)$.

LEMMA 3.6. $C\mathcal{R}_r(Q)$ is convex.

PROOF. One way to prove this is to examine the definition of $\Lambda_r(Q)$ by observing that the three arcs $\tau(u)$, $\tau(e(u, v))$, and $\tau(e(u, w))$ always form a convex chain, for any consecutive vertices v, u, w of $\alpha_r(Q)$; see Figure 2. Here we give a different proof, which provides an alternative view of $C\mathcal{R}_r(Q)$.

For each $q \in Q$, let D_q be the radius- r disk centered at q . Note that a radius- r disk contains Q if and only if its center lies in the common intersection $\bigcap_{q \in Q} D_q$. Therefore, $C\mathcal{R}_r(Q)$ is the union of all radius- r disks with centers in $\bigcap_{q \in Q} D_q$. Equivalently, $C\mathcal{R}_r(Q)$ is the Minkowski sum of $\bigcap_{q \in Q} D_q$ and the radius- r disk centered at the origin. Note that $\bigcap_{q \in Q} D_q$ is convex. As the Minkowski sum of two convex bodies is also convex, we obtain that $C\mathcal{R}_r(Q)$ must be convex. ■

4. Algorithm of Theorem 1.1

In this section, we prove Theorem 1.1. As discussed in Section 1, compared to most of the previous work on the planar two-center problem, our algorithm is relatively simple, at least conceptually. However, proving its correctness is highly intricate and we devote the entire Section 5 to it.

Let S be a set of n points in the plane. Without loss of generality, we assume that the minimum enclosing disk of S is the unit disk with equation $x^2 + y^2 \leq 1$. Choose a sufficiently

large constant integer c (say $c = 100$). Set $\mathbb{Z}/c = \{z/c : z \in \mathbb{Z}\}$ and $\theta = \frac{2\pi}{c}$. Define $\Gamma = \{(\cos i\theta, \sin i\theta) : i \in [c]\}$, which is a set of c unit vectors that evenly partition any disk centered at the origin.

Preprocessing. At the outset, we compute a farthest pair (a, b) of points in S . Let o be the midpoint of the segment \overline{ab} . Define

$$A = \{o\} \cup \{(x, y) \in \mathbb{Z}/c : x^2 + y^2 \leq 4\}.$$

Clearly, $|A| = O(1)$. In other words, A consists of o and all the possible $O(1)$ points of \mathbb{Z}/c in the radius-2 disk $x^2 + y^2 \leq 4$.

For each point $p \in A$, we compute an optimal p -anchored solution for S by Lemma 2.2. Among all these solutions, we take the one with smallest radius, and denote it by $\{D_0, D_0'\}$. Next, for each unit vector $\vec{y} \in \Gamma$, we sort the points in S along the direction \vec{y} ; let $S_{\vec{y}}$ denote the corresponding sorted sequence. This completes the preprocessing of our algorithm.

Decision procedure. Given a number $r > 0$, our goal is to determine whether there exist two radius- r disks whose union covers S , and if yes, return two such disks. For a sequence L of points in the plane, we define the *maximal r -coverable prefix* of L , denoted by $\Phi_r(L)$, as the longest prefix of L that is r -coverable. For a region R in the plane, we let $L \cap R$ denote the subsequence (not necessarily contiguous) of L consisting of the points inside R .

Algorithm 1 gives the entire algorithm. We begin with checking whether $r \geq \text{rad}(D_0)$ (note that $\text{rad}(D_0) = \text{rad}(D_0')$). If yes, we simply return $\{D_0, D_0'\}$ (Line 2). Otherwise, we consider the vectors in Γ one by one (the for-loop in Lines 3-8). For each $\vec{y} \in \Gamma$, we do the following. First, we compute X , the maximal r -coverable prefix of $S_{\vec{y}}$, and construct its r -coverage $\mathcal{CR}_r(X)$. Then, we compute Y , the maximal r -coverable prefix of $S_{\vec{y}} \cap \mathcal{CR}_r(X)$, and construct $\mathcal{CR}_r(Y)$. Finally, we compute $Z = S_{\vec{y}} \cap \mathcal{CR}_r(Y)$. Observe that $X \subseteq Y \subseteq Z$. Now we partition S into two subsets, Z and $S \setminus Z$ (here we view Z as a set instead of a sequence). We simply check whether Z and $S \setminus Z$ are both r -coverable. If so, we find a radius- r disk E (resp., E') that covers Z (resp., $S \setminus Z$) and return $\{E, E'\}$ as the solution. If no solution is found after all vectors in Γ are considered, then we return NO (Line 9). Figure 4 illustrates the algorithm with a simple example.

Time analysis. We first show that the preprocessing can be implemented in $O(n \log n)$ time. Finding a farthest pair (a, b) of S takes $O(n \log n)$ time. For each $p \in A$, computing the optimal p -anchored solution can be done in $O(n \log n)$ time by Lemma 2.2. For each $\vec{y} \in \Gamma$, computing the sorted sequence $S_{\vec{y}}$ takes $O(n \log n)$ time. Since A and Γ are both of constant size, the overall preprocessing time is $O(n \log n)$. Then we consider the time cost of the decision procedure.

```

1:  if  $r \geq \text{rad}(D_0)$  then
2:      return  $\{D_0, D_0'\}$ 
3:  for every  $\vec{\gamma} \in \Gamma$  do
4:       $X \leftarrow \Phi_r(S_{\vec{\gamma}})$ 
5:       $Y \leftarrow \Phi_r(S_{\vec{\gamma}} \cap \mathcal{CR}_r(X))$ 
6:       $Z \leftarrow S_{\vec{\gamma}} \cap \mathcal{CR}_r(Y)$ 
7:      if there exist two radius- $r$  disks  $E \supseteq Z$  and  $E' \supseteq S \setminus Z$  then
8:          return  $\{E, E'\}$ 
9:  return NO

```

Algorithm 1. $\text{Decide}(S, r, \{D_0, D_0'\}, \{S_{\vec{\gamma}} : \vec{\gamma} \in \Gamma\})$




Figure 4. Illustration of Algorithm 1. (a) The red points are not in $\mathcal{CR}_r(X)$: X consists of the three points inside the dashed circle. (b) The blue point is not in $\mathcal{CR}_r(Y)$: Y consists of the four points inside the dashed circle. (c) The points of Z and $S \setminus Z$ are covered by two radius- r disks, respectively. Observe that $X \subseteq Y \subseteq Z$.

LEMMA 4.1. *Lines 4-6 of Algorithm 1 can be implemented in $O(n)$ time.*

PROOF. For Line 4, computing $\Phi_r(S_{\vec{\gamma}})$ can be done in linear time by Lemma 2.1 since points of $S_{\vec{\gamma}}$ are already sorted along the direction $\vec{\gamma}$. Specifically, we begin with $Q = \emptyset$ and insert points of $S_{\vec{\gamma}}$ to Q one by one. We can maintain the circular hull $\alpha_r(Q)$ in $O(1)$ amortized time using the data structure of Lemma 2.1, as the insertions to Q are monotone. Note that Q is r -coverable if and only if $\alpha_r(Q)$ exists. We keep inserting points to Q until $\alpha_r(Q)$ does not exist. At this point, we obtain the maximal r -coverable prefix of $S_{\vec{\gamma}}$.




Figure 5. Illustration of the proof of Lemma 5.1.

For Line 5, we first compute the r -coverage $C\mathcal{R}_r(X)$. Recall that we can compute $\alpha_r(X)$ in linear time using Lemma 2.1 as discussed above. After that, $C\mathcal{R}_r(X)$ can be obtained from $\alpha_r(X)$ in linear time by Corollary 3.5. Next, we need to compute $S_{\vec{y}} \cap C\mathcal{R}_r(X)$. To this end, since the r -coverage $C\mathcal{R}_r(X)$ is convex by Lemma 3.6 and $S_{\vec{y}}$ is sorted along \vec{y} , computing $S_{\vec{y}} \cap C\mathcal{R}_r(X)$ can be easily done in linear time, e.g., by sweeping a line perpendicular to \vec{y} . With $S_{\vec{y}} \cap C\mathcal{R}_r(X)$ available, computing $\Phi_r(S_{\vec{y}} \cap C\mathcal{R}_r(X))$ can again be done in linear time by Lemma 2.1.

Similarly to the above for Line 5, Line 6 can also be implemented in linear time. ■

The above observation implies that each iteration of the for-loop in Algorithm 1 takes $O(n)$ time, since Lines 7-8 can be done in $O(n)$ time using the linear-time algorithm for the planar one-center problem [7, 12, 24, 23], or alternatively, using Lemma 2.1. As $|I| = O(1)$, the for-loop has $O(1)$ iterations. As such, the total runtime of Algorithm 1 is $O(n)$.

5. Correctness of the algorithm

In this section, we show the correctness of Algorithm 1, which is the last missing piece for proving Theorem 1.1. Before giving the analysis in Section 5.2, we first prove two geometric lemmas in Section 5.1 that will be needed in our analysis.

5.1 Geometric lemmas

Let S be a set of points in the plane. Recall the definition of a tight r -solution for S in Section 2. The following lemma describes a property of a tight r -solution (similar observations were made previously, e.g., [27]).

LEMMA 5.1. *Let $\{D, D'\}$ be a tight r -solution for S , and ℓ (resp., σ) be the line (resp., open segment) connecting the two disk centers $\text{ctr}(D)$ and $\text{ctr}(D')$. Then, there exist two (possibly the same) points*

$u, v \in (S \cap \partial D) \setminus D'$ such that $\overline{uv} \cap (\ell \setminus \sigma) \neq \emptyset$; see Figure 5. Similarly, there exist two (possibly the same) points $u', v' \in (S \cap \partial D') \setminus D$ such that $\overline{u'v'} \cap (\ell \setminus \sigma) \neq \emptyset$.

PROOF. It suffices to find $u, v \in (S \cap \partial D) \setminus D'$ satisfying the desired condition — the existence of $u', v' \in (S \cap \partial D') \setminus D$ can be shown likewise. Without loss of generality, we assume that ℓ is horizontal and $\text{ctr}(D)$ is left of $\text{ctr}(D')$.

Let ℓ' be the line perpendicular to ℓ and passing through $\text{ctr}(D)$. See Figure 5. We define u as a point in $(S \cap \partial D) \setminus D'$ on the side of ℓ' not containing σ , i.e., the left side of ℓ' by our assumption. Such a point always exists since $\{D, D'\}$ is a tight r -solution of S ; otherwise, we could move D towards D' so that $\{D, D'\}$ is still an r -solution while the distance between $\text{ctr}(D)$ and $\text{ctr}(D')$ decreases, a contradiction. If there are multiple such points, we choose a leftmost one. Next, we choose v as follows. If u lies on ℓ , we set $v = u$. In this case, \overline{uv} , which refers to a single point, intersects $\ell \setminus \sigma$ at the point v , and thus the lemma trivially holds. Below we assume $u \notin \ell$.

Without loss of generality, we assume that u is above ℓ . See Figure 5. Let \hat{u} be the antipodal point of u on ∂D . We set v as a point in $(S \cap \partial D) \setminus D'$ on the side of $\overline{u\hat{u}}$ not containing σ and $v \notin \overline{u\hat{u}}$. Such a point always exists; otherwise, we could slightly rotate D counterclockwise around u so that $\{D, D'\}$ is still an r -solution and the distance between $\text{ctr}(D)$ and $\text{ctr}(D')$ decreases, which contradicts the tightness of the original $\{D, D'\}$. Note that $v \neq u$ since $v \notin \overline{u\hat{u}}$.

The above finds two different points u and v in $(S \cap \partial D) \setminus D'$. Now we prove $\overline{uv} \cap (\ell \setminus \sigma) \neq \emptyset$. Note that $\overline{uv} \cap \sigma = \emptyset$, due to our choice of v (recall that v is on ∂D that lies on the side of $\overline{u\hat{u}}$ not containing σ and $v \notin \overline{u\hat{u}}$). Hence, it suffices to prove $\overline{uv} \cap \ell \neq \emptyset$. Assume to the contrary that $\overline{uv} \cap \ell = \emptyset$. Then, u and v lie on the same side of ℓ . As v is a point on ∂D that lies on the side of $\overline{u\hat{u}}$ not containing σ , v must be strictly to the left of u . But this contradicts our choice of u , i.e., u is a leftmost point among all points of $(S \cap \partial D) \setminus D'$ to the left of ℓ' . This proves $\overline{uv} \cap \ell \neq \emptyset$. ■

The following lemma proves a property regarding two radius- r disks.

LEMMA 5.2. *Let $\{D, D'\}$ be a set of two radius- r disks, $u, v \in \partial D \setminus D'$ be two points whose antipodal points on ∂D are not in D' , and ℓ (resp., σ) be the line (resp., open segment) connecting $\text{ctr}(D)$ and $\text{ctr}(D')$. If $\overline{uv} \cap (\ell \setminus \sigma) \neq \emptyset$, then $C\mathcal{R}_r(\{u, v\}) \cap D' = D \cap D'$.*

PROOF. Note first that $D \subseteq C\mathcal{R}_r(\{u, v\})$ since $\{u, v\} \subseteq D$ and the radius of D is r . Hence, $D \cap D' \subseteq C\mathcal{R}_r(\{u, v\}) \cap D'$. To prove the lemma, it suffices to show $C\mathcal{R}_r(\{u, v\}) \cap D' \subseteq D \cap D'$.

Let \hat{u} and \hat{v} be the antipodal points of u and v on ∂D , respectively. By Corollary 3.5, $C\mathcal{R}_r(\{u, v\}) \setminus D$ is a subset of the union of the following parts: (1) a disk D'' , not D , of radius r with u and v on its boundary, and (2) two circular sectors centered at u and v of radius $2r$ and tangent to D and D'' at their antipodal points of u and v , respectively (in Figure 6, the red regions are the two circular sectors excluding the points in D). Let ℓ' be the bisector of the two disk centers $\text{ctr}(D)$ and $\text{ctr}(D')$. Note that u, v, \hat{u} , and \hat{v} lie on ∂D but not in D' by the assumption




Figure 6. Illustration of the proof of Lemma 5.2. The gray region is $D'' \setminus D$. The two red regions are $CR_r(\{u, v\})$ excluding $D \cup D''$. $CR_r(\{u, v\}) \setminus D$ consists of the gray region and the two red regions.

in the lemma statement. By definition, the union of the above parts (1) and (2) completely lies in the side of ℓ' containing $\text{ctr}(D)$. Therefore, $CR_r(\{u, v\}) \setminus D$ also completely lies in the side of ℓ' containing $\text{ctr}(D)$. Hence, no point of $CR_r(\{u, v\}) \setminus D$ lies in the side of ℓ' containing $\text{ctr}(D')$. On the other hand, every point of D' in the side of ℓ' containing $\text{ctr}(D)$ must be in D . It thus follows that $CR_r(\{u, v\}) \cap D' \subseteq D \cap D'$. ■

5.2 Proving the correctness

We are now in a position to prove the correctness of Algorithm 1. Recall our assumption that the unit disk $\{(x, y) : x^2 + y^2 \leq 1\}$ is the minimum enclosing disk of S . In what follows, let $r > 0$ be the input value of the decision algorithm. Note that it suffices to consider the case $r < 1$ since otherwise S is contained in a single disk of radius r .

First of all, it is clear that whenever Algorithm 1 returns a pair of disks (instead of NO), it is always an r -solution for S . Therefore, if $r < \text{opt}(S)$ (recall the definition of $\text{opt}(S)$ in Section 2), Algorithm 1 definitely returns NO. As such, it suffices to consider the case $r \geq \text{opt}(S)$. We show below that the algorithm must return an r -solution for S .

Consider a tight r -solution $\{D, D'\}$ for S . Recall that the preprocessing step of the algorithm computes $\{D_0, D_0'\}$, which is the best solution among the p -anchored optimal solutions for all $p \in A$. If $\{D, D'\}$ is p -anchored for some $p \in A$, then $r \geq \text{rad}(D_0) = \text{rad}(D_0')$ and thus Algorithm 1 terminates in Line 2 and returns $\{D_0, D_0'\}$. In this case, the algorithm is correct. The following lemma provides two sufficient conditions for this to happen.

LEMMA 5.3. *If $\{D, D'\}$ satisfies at least one of the following two conditions, then it is p -anchored for some $p \in A$.*

- (1) $\text{dist}(D, D') \leq (2 - \frac{5}{c}) \cdot r$.
- (2) *There exist a point $u \in S \cap \partial D$ whose antipodal point on ∂D is contained in D' and a point $u' \in S \cap \partial D'$ whose antipodal point on $\partial D'$ is contained in D .*

PROOF. The proof for Condition (1) is somewhat similar to the previous work, e.g., [27]; we provide some details here for completeness.

By assumption, S is contained in the unit disk $\{(x, y) : x^2 + y^2 \leq 1\}$ and is thus 1-coverable. Recall that $r < 1$. Since $\text{dist}(D, D') \leq (2 - \frac{5}{c}) \cdot r$, r is at least $1/2$. Otherwise, a single disk of radius strictly less than 1 can cover S , which contradicts our assumption that the radius of the minimum enclosing disk of S is 1. Thus, there exists a disk of radius $1/c$ contained in $D \cap D'$, and the disk in turn contains an axis-parallel square \square of side-length $1/c$. Note that one of D and D' must intersect the unit disk $\{(x, y) : x^2 + y^2 \leq 1\}$ because $S \subseteq D \cup D'$. As $r < 1$, we know that $\square \subseteq \{(x, y) : x^2 + y^2 \leq 4\}$, and thus, $\square \cap \mathbb{Z}_{1/c}^2 \neq \emptyset$. Therefore, \square contains some point $p \in \{(x, y) \in \mathbb{Z}_{1/c}^2 : x^2 + y^2 \leq 4\} \subseteq A$, and thus $\{D, D'\}$ is p -anchored.

The proof for Condition (2) is more interesting and is new (we are not aware of any previous work that uses this condition). We assume that Condition (1) is not satisfied, since otherwise the above already completes the proof. Thus, we have $\text{dist}(D, D') > (2 - \frac{5}{c}) \cdot r$. Recall that (a, b) is a farthest pair in S , and we included in A the midpoint o of the segment \overline{ab} . Assuming that Condition (2) holds, in the following we argue that $\{D, D'\}$ must be o -anchored, which will prove the lemma since $o \in A$.

Suppose there exist a point $u \in S \cap \partial D$ whose antipodal point on ∂D is contained in D' and a point $u' \in S \cap \partial D'$ whose antipodal point on $\partial D'$ is contained in D . Let \hat{u} and \hat{u}' be the antipodal points of u on ∂D and u' on $\partial D'$, respectively.

For any point p in the plane, we use \vec{p} to denote the vector whose head is p and whose tail is $\text{ctr}(D)$, i.e., the center of D . Also, let $\|\vec{p}\|$ denote the magnitude of the vector \vec{p} , or equivalently, the distance between p and $\text{ctr}(D)$. Note that $\|\vec{p} - \vec{q}\| = |pq|$ for any two points p and q . We have

$$\begin{aligned} \|\vec{p} + \vec{q}\| &= \sqrt{\|\vec{p}\|^2 + \|\vec{q}\|^2 + 2\langle \vec{p}, \vec{q} \rangle}, & \|\vec{p} - \vec{q}\| &= \sqrt{\|\vec{p}\|^2 + \|\vec{q}\|^2 - 2\langle \vec{p}, \vec{q} \rangle}, \text{ and} \\ \|\vec{p} + \vec{q}\|^2 + \|\vec{p} - \vec{q}\|^2 &= 2\|\vec{p}\|^2 + 2\|\vec{q}\|^2, \end{aligned} \quad (1)$$

where $\langle \cdot, \cdot \rangle$ denotes the dot product.

We claim that neither D nor D' contains both a and b . Indeed, assume to the contrary this is not the case. Then, $|ab| \leq 2r$. This implies $|uu'| \leq 2r$ since (a, b) a farthest pair of S and $u, u' \in S$. Consequently, the distance between any two points among u, u', \hat{u} , and \hat{u}' is at most $2r$ since the antipodal points \hat{u} and \hat{u}' are contained in $D \cap D'$. Note that $\text{ctr}(D)$ is the midpoint of $\overline{u\hat{u}}$. Analogously, $\text{ctr}(D')$ is the midpoint of $\overline{u'\hat{u}'}$. As such, we have $\text{dist}(D, D') \leq 3r/2$ (refer to the caption of Figure 7 for a detailed explanation). But this leads to a contradiction since $\text{dist}(D, D') > (2 - \frac{5}{c}) \cdot r > 3r/2$, with $c = 100$. Hence, neither D nor D' contains both a and b .

Without loss of generality, we assume $b \in D$ and $a \in D'$. Recall that our goal is to show that $o \in D \cap D'$. In the following, we only prove $o \in D$ since $o \in D'$ can be proved analogously. Let $\delta = |ao|$, i.e., $\delta = |ab|/2$.

We first show that the distance between a and $\text{ctr}(D)$ is at most $\sqrt{r^2 + 2\delta^2}$, i.e., $\|\vec{a}\| \leq \sqrt{r^2 + 2\delta^2}$. Indeed, since (a, b) is a farthest pair of S and $u, a \in S$, we have $|au| \leq |ab| = 2\delta$.




Figure 7. The length of the red line segment connecting the two disk centers, which is $\text{dist}(D, D')$, is at most the average of the two distances from $\text{ctr}(D')$ to u and \hat{u} , respectively. The distance between u and $\text{ctr}(D')$, which is the length of the red dashed segment, is at most $2r$ since $|u\hat{u}'| \leq 2r$ and $|uu'| \leq 2r$. Furthermore, since $\hat{u} \in D'$, the distance between \hat{u} and $\text{ctr}(D')$ is at most r . Therefore, $\text{dist}(D, D') \leq 3r/2$.

Because both a and \hat{u} are in D' , $|a\hat{u}| \leq 2r$. In addition, $\|\vec{u}\| = r$ holds since u lies on ∂D . By Equality (1), we have the following (note that $\|\vec{a} + \vec{u}\| = \|\vec{a} - \vec{\hat{u}}\| = |a\hat{u}| \leq 2r$):

$$\begin{aligned} 2\|\vec{a}\|^2 &= \|\vec{a} + \vec{u}\|^2 + \|\vec{a} - \vec{u}\|^2 - 2\|\vec{u}\|^2 \\ &= |a\hat{u}|^2 + |au|^2 - 2r^2 \\ &\leq 4r^2 + 4\delta^2 - 2r^2 = 2r^2 + 4\delta^2. \end{aligned}$$

We are now in a position to prove $o \in D$. It suffices to show $\|\vec{o}\| \leq r$. Recall that $\vec{a} - \vec{o} = \vec{o} - \vec{b}$ since o is the midpoint of \overline{ab} . Notice that $\|\vec{b}\| \leq r$ since $b \in D$. From $\|\vec{a}\| \leq \sqrt{r^2 + 2\delta^2}$ and Equation (1), we have the following (which leads to $\|\vec{o}\| \leq r$):

$$\begin{aligned} (r^2 + 2\delta^2) + r^2 &\geq \|\vec{a}\|^2 + \|\vec{b}\|^2 = \|\vec{o} + (\vec{a} - \vec{o})\|^2 + \|\vec{o} + (\vec{b} - \vec{o})\|^2 \\ &= \|\vec{o} + (\vec{a} - \vec{o})\|^2 + \|\vec{o} - (\vec{a} - \vec{o})\|^2 \\ &= 2\|\vec{o}\|^2 + 2\|\vec{a} - \vec{o}\|^2 \\ &= 2\|\vec{o}\|^2 + 2\delta^2 \end{aligned}$$

As such, we obtain $o \in D$.

Similarly, we can prove $o \in D'$. Therefore, $\{D, D'\}$ is o -anchored. ■

In what follows, it suffices to focus on the case where neither condition in Lemma 5.3 holds. The failure of condition 1 implies $\text{dist}(D, D') > (2 - \frac{5}{c}) \cdot r$. For the failure of condition 2, every point $u \in S \cap \partial D$ has its antipodal point on ∂D outside D' , or this holds for D' . Without loss of generality, we assume that this is true for D . We argue below that there exists at least




Figure 8. Illustration of $u, v, s, \hat{s}, \vec{\gamma}, \vec{\eta}$. For convenience, only the directions of $\vec{\gamma}, \vec{\eta}$ are presented.

one vector $\vec{\gamma} \in \Gamma$ such that when Algorithm 1 considers $\vec{\gamma}$ in the for-loop, it returns a solution in Line 8.

Choosing $\vec{\gamma}$. The choice of $\vec{\gamma}$ is as follows. Let ℓ (resp., σ) be the line (resp., open segment) connecting the two centers $\text{ctr}(D)$ and $\text{ctr}(D')$. For any point $p \in \partial D$, we denote by \hat{p} the antipodal point of p on ∂D . The line ℓ intersects ∂D at two antipodal points s and \hat{s} , where s (resp., \hat{s}) denotes the one farther (resp., closer) to $\text{ctr}(D')$. By Lemma 5.1, there exist $u, v \in S \cap (\partial D \setminus D')$ such that $\overline{uv} \cap (\ell \setminus \sigma) \neq \emptyset$. Note that this implies $\overline{uv} \cap \ell \neq \emptyset$ and $\overline{uv} \cap \sigma = \emptyset$. By our assumption, $u, v, \hat{u}, \hat{v} \notin D'$. Without loss of generality, we can assume that $|su| \leq |sv|$. See Figure 8. The points s and \hat{s} partition ∂D into two arcs; the u -arc refers to the one where u lies on. Since $\overline{uv} \cap \ell \neq \emptyset$, v must lie on the arc other than the u -arc, which we call the v -arc. (If $u = v = s$, we pick an arbitrary arc as the u -arc and let the other one be the v -arc.) For each vector $\vec{\gamma} \in \Gamma$, we shoot a ray from $\text{ctr}(D)$ with direction $\vec{\gamma}$. These rays intersect ∂D at c points (which evenly partition ∂D into c arcs). Among these c intersection points, we pick the one on the u -arc that is closest to \hat{s} . Then we define $\vec{\gamma} \in \Gamma$ as the vector corresponding to this intersection point.

Justifying $\vec{\gamma}$. We show that when considering $\vec{\gamma}$, Algorithm 1 returns a solution. Let $\vec{\eta} \in \mathbb{S}^1$ be the unit vector with the same direction as the one originated from $\text{ctr}(D)$ to \hat{s} . Note that the angle between $\vec{\eta}$ and $\vec{\gamma}$ is at most $\frac{2\pi}{c}$ by the construction of Γ and the choice of $\vec{\gamma}$. By rotating and reflecting the coordinate system, we can assume that $\vec{\gamma} = (1, 0)$ and the clockwise angle from $\vec{\gamma}$ to $\vec{\eta}$ is at most $\frac{2\pi}{c}$; see Figure 8.

For two points $p, q \in \mathbb{R}^2$, we write $p < q$ if the x -coordinate of p is smaller than that of q . The *angle* of an arc of ∂D is defined as the central angle formed by the arc at $\text{ctr}(D)$. The following observation pertains to the relative positions of certain points with respect to u and v .

LEMMA 5.4. *The points u and v satisfy the following (see Figure 8).*

- (1) $v < \hat{u}$.
- (2) $u < p$ holds for any point $p \in D'$.

(3) $|pu| > 2r$ for any point $p \in D' \setminus D$ with $p < \hat{u}$. Note that $|pu| > 2r$ implies $p \notin CR_r(\{u\})$.

PROOF. We prove these three conditions in order.

Proving (1). To see Condition (1), observe that the angle of the subarc within the u -arc between s and u is at most $\frac{\pi}{2}$, since otherwise v would be closer to s due to $\overline{uv} \cap (\ell \setminus \sigma) \neq \emptyset$, contradicting our assumption $|su| \leq |sv|$. As the smaller angle between $\vec{\eta}$ and $\vec{\gamma}$ is at most $\frac{2\pi}{c} < \frac{\pi}{4}$, $s < \hat{u}$ must hold. So we have $v < \hat{u}$, since v lies on the subarc of the v -arc between s and \hat{u} .

Proving (2). For Condition (2), we define a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ as $f(p) = \langle p, \vec{\gamma} \rangle - \langle \text{ctr}(D'), \vec{\gamma} \rangle$, where $\langle \cdot, \cdot \rangle$ denotes the dot product. Note that for any points p, q in the plane, $p < q$ if and only if $f(p) < f(q)$. It suffices to show that $f(u) < -r$, since $f(p) \geq -r$ for any $p \in D'$. Recall that r is the radius of D (and of D'). We then have

$$\begin{aligned} f(u) &\leq f(s) + \frac{\pi r}{2} = f(\text{ctr}(D)) - r \cdot \langle \vec{\eta}, \vec{\gamma} \rangle + \frac{\pi r}{2} \\ &\leq -\left(3 - \frac{5}{c}\right) \cdot r \cdot \langle \vec{\eta}, \vec{\gamma} \rangle + \frac{\pi r}{2} \\ &\leq -\left(3 - \frac{5}{c}\right) \cdot \left(1 - \frac{2\pi}{c}\right) \cdot r + \frac{\pi r}{2} < -r. \end{aligned}$$

The first inequality holds because the distance between s and u is at most $\frac{\pi r}{2}$ (this is in turn because the angle of the subarc within the u -arc between s and u is at most $\frac{\pi}{2}$ as already discussed above). The equality in the first line holds because $\vec{\eta}$ is the unit vector directed from s to $\text{ctr}(D)$ and the distance between s and $\text{ctr}(D)$ is r . The second inequality holds since $f(\text{ctr}(D)) = -\text{dist}(D, D') \cdot \langle \vec{\eta}, \vec{\gamma} \rangle$ and we have assumed $\text{dist}(D, D') > (2 - \frac{5}{c}) \cdot r$. The last one follows from $\langle \vec{\eta}, \vec{\gamma} \rangle \geq 1 - \frac{2\pi}{c}$, as the angle between $\vec{\gamma}$ and $\vec{\eta}$ is at most $\frac{2\pi}{c}$. This proves the second condition.

Proving (3). We finally prove Condition (3). Let p be any point in $D' \setminus D$ with $p < \hat{u}$. Our goal is to prove $|pu| > 2r$. Define ℓ' to be the perpendicular bisector line of σ . Let ℓ_p be the vertical line through p . By the way, we chose $\vec{\gamma}$, u and the vector $\vec{\gamma}$ originating from $\text{ctr}(D)$ lie on the same side of ℓ . Recall that we have rotated the coordinate system so that $\vec{\gamma} = (1, 0)$. Hence, if u is above ℓ , then the direction of $\vec{\eta}$ is downward, i.e., $\text{ctr}(D)$ is higher than $\text{ctr}(D')$ (see Figure 9(a)). If u is below ℓ , then $\vec{\eta}$ is upward and $\text{ctr}(D')$ is higher than $\text{ctr}(D)$ (see Figure 9(b)). In the following, we only consider the case where u is above ℓ , since the other case can be proved analogously. Since \hat{u} is the antipodal point of u on ∂D and ℓ contains $\text{ctr}(D)$, \hat{u} is below ℓ .

We claim that \hat{u} must lie in the triangle bounded by the lines ℓ , ℓ' , and ℓ_p ; see Figure 9. To see this, first of all, since $p < \hat{u}$, \hat{u} is to the right of ℓ_p . Also, notice that $D \setminus D'$ is to the left of ℓ' . Since $\hat{u} \in D$ and $\hat{u} \notin D'$, we obtain that \hat{u} is to the left of ℓ' . Finally, since $\text{ctr}(D)$ is higher than $\text{ctr}(D')$, \hat{u} is below ℓ , $\hat{u} \in \partial D$, and $\hat{u} \notin D'$, we obtain that the intersection $q = \ell \cap \ell'$ must be to the




Figure 9. Illustration of the proof of Condition (3) of Lemma 5.4.

right of the vertical line through \hat{u} . Since \hat{u} is right of ℓ_p , we obtain that q is to the right of ℓ_p . The above discussions together lead to that \hat{u} must be in the triangle bounded by ℓ , ℓ' , and ℓ_p .

By the definition of ℓ' , the portion of D' to the left of ℓ' must be inside D . Since $p \in D' \setminus D$, p must be to the right of ℓ' . Recall that \hat{u} is to the left of ℓ' . Since u is above ℓ while \hat{u} is below ℓ , we obtain that the triangle $\Delta p\hat{u}u$ has an obtuse angle at \hat{u} ; see Figure 9. As a consequence, $|up| > |u\hat{u}| = 2r$. ■

Consider the sets X, Y, Z computed in Lines 4-6 of Algorithm 1 for the vector \vec{y} . Our goal is to prove that Z is exactly equal to $S \cap D$ (when viewed as a set), which implies $S \setminus Z \subseteq D'$ because $S \subseteq D \cup D'$. Note that as long as this is true, both Z and $S \setminus Z$ are r -coverable, and thus Algorithm 1 will return a solution in Line 8 when considering \vec{y} .

LEMMA 5.5. *Let u and v be as above. We have $u \in X \subseteq D$, $\{u, v\} \subseteq Y \subseteq D$, and $Z = S \cap D$.*

PROOF. Let (p_1, \dots, p_n) be the sorted sequence $S_{\vec{y}}$; let j and k be the indices with $u = p_j$ and $v = p_k$. Recall that Algorithm 1 computes X, Y and Z in this order.

Proving $u \in X \subseteq D$. We first prove $u \in X \subseteq D$. Recall that X is the maximal r -coverable prefix of $S_{\vec{y}}$. According to Lemma 5.4(2), $p_i \notin D'$ for all indices $i \leq j$, and therefore $\{p_1, \dots, p_j\} \subseteq D$. Hence, $\{p_1, \dots, p_j\}$ is an r -coverable, and thus $\{p_1, \dots, p_j\} \subseteq X$. In particular, $u = p_j \in X$. To show $X \subseteq D$, we distinguish two cases: $v \in X$ and $v \notin X$.

- If $v \in X$, then we have $X \subseteq \mathcal{CR}_r(\{u, v\})$ since $u, v \in X$ and X is r -coverable. By Lemma 5.2, $\mathcal{CR}_r(\{u, v\}) \cap D' = D \cap D'$. As such, we obtain $X \cap D' \subseteq D \cap D'$, and consequently, $X \cap (D' \setminus D) = \emptyset$. This further implies that $X \subseteq D$ since $X \subseteq D \cup D'$.
- If $v \notin X$, we assume $X \not\subseteq D$ for the sake of contradiction. Then, there exists p_i such that $\{p_1, \dots, p_i\} \subseteq X$ is r -coverable but p_i lies outside D (and thus $p_i \in D' \setminus D$). Since $v \notin X$, we have $p_i < v$. As $v < \hat{u}$ by Lemma 5.4(1), $p_i < \hat{u}$ holds. Because $p_i \in D' \setminus D$ and $p_i < \hat{u}$, we further have $p_i \notin \mathcal{CR}_r(\{u\})$ by Lemma 5.4(3).

On the other hand, since X is r -coverable and $u \in X$, $X \subseteq C\mathcal{R}_r(\{u\})$ must hold. As $p_i \in X$, we have $p_i \in C\mathcal{R}_r(\{u\})$. We thus obtain contradiction.

Proving $\{u, v\} \subseteq Y \subseteq D$. We now prove $\{u, v\} \subseteq Y \subseteq D$. Recall that Y is the maximal r -coverable prefix of $S_{\bar{y}} \cap C\mathcal{R}_r(X)$. As $X \subseteq D$, $D \subseteq C\mathcal{R}_r(X)$ holds. Since $v \in D$, we have $v \in S \cap C\mathcal{R}_r(X)$. Furthermore, since X is r -coverable, it follows that $X \subseteq C\mathcal{R}_r(X)$, which further implies $X \subseteq Y$. As $u \in X$, we have $u \in Y$.

To show $v \in Y$, we claim that all points of $\{p_1, \dots, p_k = v\} \cap C\mathcal{R}_r(X)$ lie in D . Assume to the contrary this is not true. Then, there is a point p_i in $\{p_1, \dots, p_k = v\} \cap C\mathcal{R}_r(X)$ such that $p_i \notin D$ (and thus $p_i \in D' \setminus D$). Thus, either $p_i = v$ or $p_i < v$. Since $v < \hat{u}$ by Lemma 5.4(1), we have $p_i < \hat{u}$. Since $p_i \in D' \setminus D$ and $p_i < \hat{u}$, we further have $p_i \notin C\mathcal{R}_r(\{u\})$ by Lemma 5.4(3). On the other hand, since $u \in X$ and X is r -coverable, $C\mathcal{R}_r(X) \subseteq C\mathcal{R}_r(\{u\})$ holds. As $p_i \in C\mathcal{R}_r(X)$, we obtain $p_i \in C\mathcal{R}_r(\{u\})$, a contradiction.

The above claim implies that $\{p_1, \dots, p_k = v\} \cap C\mathcal{R}_r(X)$ is r -coverable, and thus v must be in Y . As $\{u, v\} \subseteq Y$ and Y is r -coverable, it follows that $Y \subseteq C\mathcal{R}_r(\{u, v\})$. According to Lemma 5.2, $C\mathcal{R}_r(\{u, v\}) \cap D' = D \cap D'$. Therefore, $Y \cap D' \subseteq D \cap D'$, which leads to $Y \subseteq D$ since $Y \subseteq D \cup D'$.

Proving $Z = S \cap D$. We finally prove $Z = S \cap D$, using the fact $\{u, v\} \subseteq Y \subseteq D$. Recall that $Z = S \cap C\mathcal{R}_r(Y)$. As $Y \subseteq D$ and Y is r -coverable, we have $D \subseteq C\mathcal{R}_r(Y)$, which implies $S \cap D \subseteq Z$. It suffices to show $Z \subseteq S \cap D$, or equivalently, $Z \subseteq D$. Observe that $Z = (Z \cap D) \cup (Z \cap D')$, since $D \cup D'$ contains all points in Z . Therefore, we only need to show $Z \cap D' \subseteq D$. Since $\{u, v\} \subseteq Y$ and Y is r -coverable, we have $C\mathcal{R}_r(Y) \subseteq C\mathcal{R}_r(\{u, v\})$. By Lemma 5.2, $C\mathcal{R}_r(\{u, v\}) \cap D' = D \cap D'$. It follows that

$$Z \cap D' \subseteq C\mathcal{R}_r(Y) \cap D' \subseteq C\mathcal{R}_r(\{u, v\}) \cap D' = D \cap D' \subseteq D.$$

We then have $Z \cap D' \subseteq D$ and thus $Z \subseteq D$.

This completes the proof. ■

We have thus completed the full proof of the correctness of Algorithm 1, from which we derive Theorem 1 and Corollary 2.

6. Conclusion and future work

This paper presents an $O(n \log n)$ -time deterministic algorithm for the planar two-center problem, which improves the previous $O(n \log^2 n)$ -time algorithm by Wang [29] and matches the $\Omega(n \log n)$ lower bound for the problem given by Eppstein [15]. Our algorithm has a straightforward design while its analysis is technically involved and requires new approaches to the problem. We believe that some of our techniques can find applications in other variants of the planar two-center problem (or more generally, other geometric coverage problems).

Next, we pose some interesting open questions for future study. First, given that the planar two-center problem can now be solved optimally, it is natural to ask for the largest k such that planar k -center can be solved in $O(n \log n)$ time (or near-linear time). This question is completely open: even for $k = 3$ with the assumption that all input points are in convex positions, the best-known algorithm runs in near-quadratic time [5]. Furthermore, one can also study the two-center problem in higher dimensions and ask whether the problem can be solved in near-linear time in any fixed dimension (which is true for the one-center problem). Finally, we notice that, while near-linear time algorithms for the planar two-center problem have been known for more than two decades since the work of Sharir [27], the best known algorithms for the *discrete* planar two-center problem run in $O(n^{4/3} \cdot \text{poly}(\log n))$ time [3, 30], as mentioned in the introduction. Therefore, it is interesting to ask whether the techniques in this paper can also be applied to design near-linear time algorithms for the discrete problem.

References

- [1] Pankaj K. Agarwal and Jeff M. Phillips. An efficient algorithm for 2D Euclidean 2-center with outliers. *Proceedings of the 16th Annual European Symposium on Algorithms (ESA)*, pages 64–75, 2008. [DOI](#) (3)
- [2] Pankaj K. Agarwal and Micha Sharir. Planar geometric location problems. *Algorithmica*, 11(2):185–195, 1994. [DOI](#) (2)
- [3] Pankaj K. Agarwal, Micha Sharir, and Emo Welzl. The discrete 2-center problem. *Discrete and Computational Geometry*. 20(3):287–305, 1998. [DOI](#) (3, 20)
- [4] Esther M. Arkin, José Miguel Díaz-Báñez, Ferran Hurtado, Piyush Kumar, Joseph S.B. Mitchell, Belén Palop, Pablo Pérez-Lantero, Maria Saumell, and Rodrigo I. Silveira. Bichromatic 2-center of pairs of points. *Computational Geometry: Theory and Applications*, 48(2):94–107, 2015. [DOI](#) (3)
- [5] Donglai Bian, Bo Jiang, and Zhiying Cao. An efficient algorithm of the planar 3-center problem for a set of the convex position points. *MATEC Web of Conferences*, volume 232, page 03022. EDP Sciences, 2018. [DOI](#) (20)
- [6] Timothy M. Chan. More planar two-center algorithms. *Computational Geometry: Theory and Applications*, 13(3):189–198, 1999. [DOI](#) (2)
- [7] Bernard Chazelle and Jiří Matoušek. On linear-time deterministic algorithms for optimization problems in fixed dimension. *Journal of Algorithms*, 21:579–597, 1996. [DOI](#) (1, 11)
- [8] Jongmin Choi and Hee-Kap Ahn. Efficient planar two-center algorithms. *Computational Geometry: Theory and Applications*, 97(101768), 2021. [DOI](#) (2, 3, 5)
- [9] Jongmin Choi, Jaegun Lee, and Hee-Kap Ahn. Efficient k -center algorithms for planar points in convex position. *Proceedings of the 18th Algorithms and Data Structures Symposium (WADS)*, pages 262–274, 2023. [DOI](#) (3)
- [10] Richard Cole. Slowing down sorting networks to obtain faster sorting algorithms. *Journal of the ACM*, 34(1):200–208, 1987. [DOI](#) (3)
- [11] Mark de Berg, Marcel Roeloffzen, and Bettina Speckmann. Kinetic 2-centers in the black-box model. *Proceedings of the Twenty-Ninth Annual Symposium on Computational Geometry (SoCG)*, pages 145–154, 2013. [DOI](#) (3)
- [12] M.E. Dyer. On a multidimensional search technique and its application to the Euclidean one centre problem. *SIAM Journal on Computing*, 15(3):725–738, 1986. [DOI](#) (1, 11)
- [13] Herbert Edelsbrunner, David G. Kirkpatrick, and Raimund Seidel. On the shape of a set of points in the plane. *IEEE Transactions on Information Theory*, 29:551–559, 1983. [DOI](#) (4)
- [14] David Eppstein. Dynamic three-dimensional linear programming. *INFORMS Journal on Computing*, 4(4):360–368, 1992. [DOI](#) (2)
- [15] David Eppstein. Faster construction of planar two-centers. *Proceedings of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 131–138, 1997. (1–3, 19)
- [16] Dan Halperin, Micha Sharir, and Ken Goldberg. The 2-center problem with obstacles. *Journal of Algorithms*, 42(1):109–134, 2002. [DOI](#) (3)
- [17] John Hershberger. A faster algorithm for the two-center decision problem. *Information Processing Letters*, 1:23–29, 1993. [DOI](#) (2)

- [18] John Hershberger and Subhash Suri. Finding tailored partitions. *Journal of Algorithms*, 12(3):431–463, 1991. DOI (2, 4)
- [19] R. Z. Hwang, R. C. T. Lee, and R. C. Chang. The slab dividing approach to solve the Euclidean p -center problem. *Algorithmica*, 9(1):1–22, 1993. DOI (1, 3)
- [20] Jerzy W. Jaromczyk and Mirosław Kowaluk. An efficient algorithm for the Euclidean two-center problem. *Proceedings of the Tenth Annual Symposium on Computational Geometry (SoCG)*, pages 303–311, 1994. DOI (2)
- [21] Matthew J. Katz and Micha Sharir. An expander-based approach to geometric optimization. *SIAM Journal on Computing*, 26(5):1384–1408, 1997. DOI (2)
- [22] Nimrod Megiddo. Applying parallel computation algorithms in the design of serial algorithms. *Journal of the ACM*, 30(4):852–865, 1983. DOI (2)
- [23] Nimrod Megiddo. Linear programming in linear time when the dimension is fixed. *Journal of the ACM*, 31(1):114–127, 1984. DOI (1, 11)
- [24] Nimrod Megiddo. Linear-time algorithms for linear programming in R^3 and related problems. *SIAM Journal on Computing*, 12(4):759–776, 1983. DOI (1, 11)
- [25] Eunjin Oh, Sang Won Bae, and Hee-Kap Ahn. Computing a geodesic two-center of points in a simple polygon. *Computational Geometry: Theory and Computation*, 82:45–59, 2019. DOI (3)
- [26] Eunjin Oh, Jean-Lou De Carufel, and Hee-Kap Ahn. The geodesic 2-center problem in a simple polygon. *Computational Geometry: Theory and Computation*, 74:21–37, 2018. DOI (3)
- [27] Micha Sharir. A near-linear algorithm for the planar 2-center problem. *Discrete and Computational Geometry*, 2(18):125–134, 1997. DOI (2, 11, 14, 20)
- [28] Xuehou Tan and Bo Jiang. Simple $O(n \log^2 n)$ algorithms for the planar 2-center problem. *Proceedings of the 23rd International Computing and Combinatorics Conference (COCOON)*, pages 481–491, 2017. DOI (2)
- [29] Haitao Wang. On the planar two-center problem and circular hulls. *Discrete and Computational Geometry*, 68(4):1175–1226, 2022. DOI (2, 4, 19)
- [30] Haitao Wang. Unit-disk range searching and applications. *Journal of Computational Geometry*, 14:343–394, 2023. DOI (3, 20)
- [31] Haitao Wang and Jie Xue. Improved algorithms for the bichromatic two-center problem for pairs of points. *Computational Geometry: Theory and Applications*, 100(101806):1–12, 2022. DOI (3)