

An Efficient Algorithm for Euclidian 2-Center with Outliers*

Pankaj K. Agarwal[†]

Jeff M. Phillips[‡]

Abstract

For a set P of n points in \mathbb{R}^2 , the Euclidean 2-center problem computes a pair of congruent disks of the minimal radius that cover P . We extend this to the $(2, k)$ -center problem where we compute the minimal radius pair of congruent disks to cover $n - k$ points of P . We present a randomized algorithm with $O(nk^7 \log^3 n)$ expected running time for the $(2, k)$ -center problem. We also study the (p, k) -center problem in \mathbb{R}^2 under the ℓ_∞ -metric. We give solutions for $p = 4$ in $O(k^{O(1)}n \log n)$ time and for $p = 5$ in $O(k^{O(1)}n \log^5 n)$ time.

1 Introduction

Let P be a set of n points in \mathbb{R}^2 . For a pair of integers $0 \leq k \leq n$ and $p \geq 1$, a family of p congruent disks is called a (p, k) -center if the disks cover at least $n - k$ points of P ; and $(p, 0)$ -center is the standard p -center. The Euclidean (p, k) -center problem asks for computing a (p, k) -center of P of the smallest radius. In this paper we study the $(2, k)$ -center problem. We also study the (p, k) -center problem under the ℓ_∞ -metric for small values of p and k . Here we wish to cover all but k points of P by p congruent axis-aligned squares of the smallest side length. Our goal is to develop algorithms whose running time is $n(k \log n)^{O(1)}$.

Related work. There has been extensive work on the p -center problem in algorithms and operations research communities [3, 10, 13]. If p is part of the input, the problem is NP-hard [17] even for the Euclidean case in \mathbb{R}^2 . The Euclidean 1-center problem is known to be LP-type [15], and therefore can be solved in linear time for any fixed dimensions. The Euclidean 2-center problem is not LP-type. Agarwal and Sharir [2] proposed an $O(n^2 \log^3 n)$ time algorithm for the 2-center problem. The running time was improved to $O(n \log^{O(1)} n)$ by Sharir [19]. The exponent of the $\log n$ factor was subsequently improved in [11, 4]. The best known deterministic algorithm takes $O(n \log^2 \log^2 n)$ time in the worst case, and the best known randomized algorithm takes $O(n \log^2 n)$ expected time.

There is little work on the (p, k) -center problem. Using a framework described by Matoušek [14] the $(1, k)$ -center problem can be solved in $O(n \log k + k^3 n^\varepsilon)$ time for any $\varepsilon > 0$. In general, Matoušek shows how to solve this problem with k outliers in $O(nk^d)$ time where d is the inherent number of constraints in the solution. The bound for the $(1, k)$ -center problem (and any LP-type problem with $d = 3$) is improved by Chan [5] to $O(n \log n + k^2 \log^2 n)$.

*This work is supported by NSF under grants CNS-05-40347, CFF-06-35000, and DEB-04-25465, by ARO grants W911NF-04-1-0278 and W911NF-07-1-0376, by an NIH grant 1P50-GM-08183-01, by a DOE grant OEGP200A070505, and by a grant from the U.S. Israel Binational Science Foundation.

[†]Department of Computer Science, Duke University, Durham, NC 27708: pankaj@cs.duke.edu

[‡]Department of Computer Science, Duke University, Durham, NC 27708: jeffp@cs.duke.edu

The p -center problem under ℓ_∞ -metric is dramatically simpler. Sharir and Welzl [20] show how to compute the ℓ_∞ p -center in near-linear time for $p \leq 5$. In fact, they show that the rectilinear 2- and 3-center problems are LP-type problems and can be solved in $O(n)$ time. Also, they show the 1-dimensional version of the problem is an LP-type problem for any p , with combinatorial dimension $O(p)$. Thus applying Matoušek’s framework [14], the ℓ_∞ (p, k) -center in \mathbb{R}^2 for $p \leq 3$, can be found in $O(k^{O(1)}n)$ time and in $O(k^{O(p)}n)$, for any p , if the points lie in \mathbb{R}^1 .

Our results. Our main result is a randomized algorithm for the Euclidean $(2, k)$ -center problem in \mathbb{R}^2 whose expected running time is $O(nk^7 \log^3 n)$. We follow the general framework of Sharir and subsequent improvements by Eppstein. But in order to handle outliers we first prove, in Section 2, a few structural properties of levels in an arrangement of unit disks, which are of independent interest.

As in [19, 11], our solution breaks the $(2, k)$ -center problem into two subproblems: when the optimal disk’s centers are further apart than the optimal radius, and when the optimal disk’s centers are closer than their radius. The first subproblem, which we refer to as the *well-separated case* and describe in Section 3, takes $O(k^6 n \log^3 n)$ time in the worst case and uses parametric search [16]. The second subproblem, which we refer to as the *nearly concentric case* and describe in Section 4, takes $O(k^7 n \log^3 n)$ expected time. Thus we solve the $(2, k)$ -center problem in $O(k^7 n \log^3 n)$ expected time. We can solve the nearly concentric case and hence the $(2, k)$ -center problem in $O(k^7 n^{1+\delta})$ deterministic time, for any $\delta > 0$. We present near-linear algorithms for the ℓ_∞ (p, k) -center in \mathbb{R}^2 for $p = 4, 5$. The ℓ_∞ $(4, k)$ -center problem takes $O(k^{O(1)}n \log n)$ time, and the ℓ_∞ $(5, k)$ -center problem takes $O(k^{O(1)}n \log^5 n)$ time. Because of lack of space, these results are not described in this abstract. We have not made an attempt to minimize the exponent of k . We believe that it can be improved by a more careful analysis.

2 Arrangement of Unit Disks

Let $\mathcal{D} = \{D_1, \dots, D_n\}$ be a set of n unit disks in \mathbb{R}^2 . Let $\mathcal{A}(\mathcal{D})$ be the arrangement of \mathcal{D} .¹ $\mathcal{A}(\mathcal{D})$ consists of $O(n^2)$ vertices, edges, and faces. For a subset $\mathcal{R} \subseteq \mathcal{D}$, let $\mathcal{J}(\mathcal{R}) = \bigcap_{D \in \mathcal{R}} D$ denote the intersection of disks in \mathcal{R} . Each disk in \mathcal{R} contributes at most one edge in $\mathcal{J}(\mathcal{R})$. We refer to $\mathcal{J}(\mathcal{R})$ as a *unit-disk polygon* and a connected portion of $\partial\mathcal{J}(\mathcal{R})$ as a *unit-disk curve*. We introduce the notion of level in $\mathcal{A}(\mathcal{D})$, prove a few structural properties of levels, and describe an algorithm that will be useful for our overall algorithm.

Levels and their structural properties. For a point $x \in \mathbb{R}^2$, the *level* of x with respect to \mathcal{D} , denoted by $\lambda(x, \mathcal{D})$, is the number of disks in \mathcal{D} that *do not* contain x . (Our definition of level is different from the more common definition in which it is defined as the number of disks whose interiors contain x .) All points lying on an edge or face ϕ of $\mathcal{A}(\mathcal{D})$ have the same level,

¹The *arrangement* of \mathcal{D} is the planar decomposition induced by \mathcal{D} ; its vertices are the intersection points of boundaries of two disks, its edges are the maximal portions of disk boundaries that do not contain a vertex, and its faces are the maximal connected regions of the plane that do not intersect the boundary of any disk.

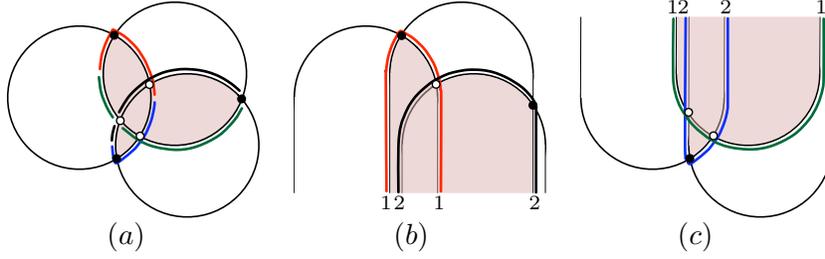


Figure 1: (a) $\mathcal{A}(\mathcal{D})$, shaded region is $\mathcal{A}_{\leq 1}(\mathcal{D})$, filled (resp. hollow) vertices are convex (resp. concave) vertices of $\mathcal{A}_{\leq 1}(\mathcal{D})$, covering of $\mathcal{A}_{\leq 1}(\mathcal{D})$ edges by six unit-disk curves. (b) $\mathcal{A}(\Gamma^+)$, shaded region is $\mathcal{A}_{\leq 1}(\Gamma^+)$, and the covering of $\mathcal{A}_{\leq 1}(\Gamma^+)$ edges by two concave chains. (c) $\mathcal{A}(\Gamma^-)$, shaded region is $\mathcal{A}_{\leq 1}(\Gamma^-)$, and the covering of $\mathcal{A}_{\leq 1}(\Gamma^-)$ edges by two convex chains.

which we denote by $\lambda(\phi)$. For $k \leq n$, let $\mathcal{A}_k(\mathcal{D})$ (resp. $\mathcal{A}_{\leq k}(\mathcal{D})$) denote the set of points in \mathbb{R}^2 whose level is k (resp. at most k); see Figure 1. By definition, $\mathcal{A}_0(\mathcal{D}) = \mathcal{A}_{\leq 0}(\mathcal{D}) = \mathcal{J}(\mathcal{D})$.

The boundary of $\mathcal{A}_{\leq k}(\mathcal{D})$ is composed of the edges of $\mathcal{A}(\mathcal{D})$. Let $v \in \partial D_1 \cap \partial D_2$, for $D_1, D_2 \in \mathcal{D}$, be a vertex of $\partial \mathcal{A}_{\leq k}(\mathcal{D})$. We call v *convex* (resp. *concave*) if $\mathcal{A}_{\leq k}(\mathcal{D})$ lies in $D_1 \cap D_2$ (resp. $D_1 \cup D_2$) in a sufficiently small neighborhood of v . $\partial \mathcal{A}_{\leq 0}(\mathcal{D})$ is composed of convex vertices; see Figure 1(a). We define the complexity of $\mathcal{A}_{\leq k}(\mathcal{D})$ to be the number of edges of $\mathcal{A}(\mathcal{D})$ whose levels are at most k . Since the complexity of $\mathcal{A}_{\leq 0}(\mathcal{D})$ is n , the following lemma follows from the result by Clarkson and Shor [8] (see also Sharir [18] and Chan [6]).

Lemma 2.1. [8] *For $k \geq 0$, the complexity of $\mathcal{A}_{\leq k}(\mathcal{D})$ is $O(nk)$.*

Remark. The argument by Clarkson and Shor can also be used to prove that $\mathcal{A}_{\leq k}(\mathcal{D})$ has $O(k^2)$ connected components and that it has $O(k^2)$ local minimums in $(+y)$ -direction. See also [14, 7]. These bounds are tight in the worst case; see Figure 2.

It is well known that the edges in $\leq k$ -level of a line arrangement can be covered by $k + 1$ concave chains. We prove a similar result for $\mathcal{A}_{\leq k}(\mathcal{D})$; it can be covered by $O(k)$ unit-disk curves.

For a disk D_i , let γ_i^+ (resp. γ_i^-) denote the set of points that lie on or below (resp. above) D_i ; $\partial \gamma_i^+$ consists of the upper semicircle of ∂D_i plus two vertical downward rays emanating from the left and right endpoints of the semicircle — we refer to these rays as left and right rays. The curve $\partial \gamma_i^-$ has a similar structure. See Figure 1(b). Set $\Gamma^+ = \{\gamma_i^+ \mid 1 \leq i \leq n\}$ and $\Gamma^- = \{\gamma_i^- \mid 1 \leq i \leq n\}$. Each pair of curves $\partial \gamma_i^+, \partial \gamma_j^+$ intersect in at most one point. (If we assume that the left and right rays are not vertical but have very large positive and negative slopes, respectively, then each pair of boundary curves intersects in exactly one point.) We define the level of a point with respect to Γ^+ , Γ^- , or $\Gamma^+ \cup \Gamma^-$ in the same way as with respect to \mathcal{D} . A point lies in a disk D_i if and only if it lies in both γ_i^+ and γ_i^- , so we obtain the following inequalities:

$$\max\{\lambda(x, \Gamma^+), \lambda(x, \Gamma^-)\} \leq \lambda(x, \mathcal{D}). \quad (2.1)$$

$$\lambda(x, \mathcal{D}) \leq \lambda(x, \Gamma^+ \cup \Gamma^-) \leq 2\lambda(x, \mathcal{D}). \quad (2.2)$$

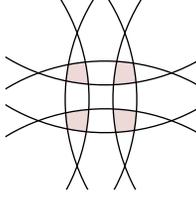


Figure 2: Constant size version of lower bound. $\mathcal{A}_{\leq 2}(\mathcal{D})$ has 4 connected components, shaded.

We cover the edges of $\mathcal{A}_{\leq k}(\Gamma^+)$ by $k + 1$ concave chains as follows. The level of the $(k + 1)$ st rightmost left ray is at most k at $y = -\infty$. Let ρ_i be such a ray, belonging to γ_i^+ . We begin tracing $\partial\gamma_i^+$, beginning from the point at $y = -\infty$ on ρ_i , as long as $\partial\gamma_i^+$ remains in $\mathcal{A}_{\leq k}(\Gamma^+)$. We stop when we have reached a vertex $v \in \mathcal{A}_{\leq k}(\Gamma^+)$ at which it leaves $\mathcal{A}_{\leq k}(\Gamma^+)$; v is a convex vertex on $\mathcal{A}_{\leq k}(\Gamma^+)$. Suppose $v = \partial\gamma_i^+ \cap \partial\gamma_j^+$. Then $\partial\mathcal{A}_{\leq k}(\Gamma^+)$ follows $\partial\gamma_j^+$ immediately to the right of v , so we switch to $\partial\gamma_j^+$ and repeat the same process. It can be checked that we finally reach $y = -\infty$ on a right ray. Since we switch the curve on a convex vertex, the chain Λ_i^+ we trace is a concave chain, which is composed of a left ray, followed by a unit-disk curve ξ_i^+ , and then followed by a right ray. Let $\Lambda_0^+, \Lambda_1^+, \dots, \Lambda_k^+$ be the $k + 1$ chains traversed by this procedure. These chains cover all edges of $\mathcal{A}_{\leq k}(\Gamma^+)$, and each edge lies exactly on one chain. Similarly we cover the edges of $\mathcal{A}_{\leq k}(\Gamma^-)$ by $k + 1$ convex curves $\Lambda_0^-, \Lambda_1^-, \dots, \Lambda_k^-$. Let $\Xi = \{\xi_0^+, \dots, \xi_k^+, \xi_0^-, \dots, \xi_k^-\}$ be the family of unit-disk curves induced by these convex and concave chains. By (2.1), Ξ covers all edges of $\mathcal{A}_{\leq k}(\mathcal{D})$. Since a unit circle intersects a unit-disk curve in at most two points, we conclude the following.

Lemma 2.2. *The edges of $\mathcal{A}_{\leq k}(\mathcal{D})$ can be covered by at most $2k + 2$ unit-disk curves, and a unit circle intersects $O(k)$ edges of $\mathcal{A}_{\leq k}(\mathcal{D})$.*

The curves in Ξ may contain edges of $\mathcal{A}(\mathcal{D})$ whose levels are greater than k . If we wish to find a family of unit-disk curves whose union is the set of edges in $\mathcal{A}_{\leq k}(\mathcal{D})$, we proceed as follows. We add the x -extremal points of each disk as vertices of $\mathcal{A}(\mathcal{D})$, so each edge is now x -monotone and lies in a lower or an upper semicircle. By (2.1), only $O(k)$ such vertices lie in $\mathcal{A}_{\leq k}(\mathcal{D})$. We call a vertex of $\mathcal{A}_{\leq k}(\mathcal{D})$ *extremal* if it is an x -extremal point on a disk or an intersection point of a lower and an upper semicircle. Lemma 2.2 implies that there are $O(k^2)$ extremal vertices. For each extremal vertex v we do the following. If there is an edge e of $\mathcal{A}_{\leq k}(\mathcal{D})$ lying to the right of v , we follow the arc containing e until we reach an extremal vertex or we leave $\mathcal{A}_{\leq k}(\mathcal{D})$. In the former case we stop. In the latter case we are at a convex vertex v' of $\partial\mathcal{A}_{\leq k}(\mathcal{D})$, and we switch to the other arc incident on v' and continue. These curves have been drawn in Figure 1(a). This procedure returns an x -monotone unit-disk curve that lies in $\mathcal{A}_{\leq k}(\mathcal{D})$. It can be shown that this procedure covers all edges of $\mathcal{A}_{\leq k}(\mathcal{D})$. We thus obtain the following:

Lemma 2.3. *Let \mathcal{D} be a set of n unit disks in \mathbb{R}^2 . Given $\mathcal{A}_{\leq k}(\mathcal{D})$, we can compute, in time $O(nk)$, a family of $O(k^2)$ x -monotone unit-disk curves whose union is the set of edges of $\mathcal{A}_{\leq k}(\mathcal{D})$.*

Remark. Since $\mathcal{A}_{\leq k}(\mathcal{D})$ can consist of $\Omega(k^2)$ connected components, the $O(k^2)$ bound is tight in the worst case; see Figure 2.

Emptiness detection of $\mathcal{A}_{\leq k}(\mathcal{D})$. We need a dynamic data structure for storing a set \mathcal{D} of unit disks that supports the following two operations:

(O1) Insert a disk into \mathcal{D} or delete a disk from \mathcal{D} ;

(O2) For a given k , determine whether $\mathcal{A}_{\leq k}(\mathcal{D}) \neq \emptyset$.

As described by Sharir [19], $\mathcal{J}(\mathcal{D})$ can be maintained under insertion/deletion in $O(\log^2 n)$ time per update. Matoušek [14] has described a data structure for solving LP-type problems with violations. Finding the lowest point in $\mathcal{J}(\mathcal{D})$ can be formulated as an LP-type problem. Therefore using the dynamic data structure with Matoušek's algorithm, we can obtain the following result. The details are omitted from this abstract.

Lemma 2.4. *There exists a dynamic data structure for storing a set of n unit disks so that (O1) can be performed in $O(\log^2 n)$ time, and (O2) takes $O(k^3 \log^2 n)$ time.*

3 Well-Separated Disks

In this section we describe an algorithm for the case in which the two disks D_1, D_2 of the optimal solution are well separated. That is, let c_1 and c_2 be the centers of D_1 and D_2 , and let r^* be their radius. Then $\|c_1 c_2\| \geq r^*$; see Figure 3. Without loss of generality, let us assume that c_1 lies to the left of c_2 . Let D_1^- be the semidisk lying to the left of the line passing through c_1 in direction normal to $c_1 c_2$. If $D_1 \cap D_2 = \emptyset$, then a line ℓ is called a *separator line* if ℓ separates D_1^- from D_2 . If $D_1 \cap D_2 \neq \emptyset$, then ℓ is called a separator line if it separates D_1^- from the intersection points of $\partial D_1 \cap \partial D_2$. We first show that we can quickly compute a set of $O(k^2)$ lines that contains a separator line. Next, we describe a decision algorithm, and then we describe the algorithm for computing D_1 and D_2 provided they are well separated.

Computing separator lines. We fix a sufficiently large constant h and we choose a set $U = \{u_1, \dots, u_h\} \subseteq \mathbb{S}^1$ of uniformly distributed directions, i.e., $u_i = (\cos(2\pi i/h), \sin(2\pi i/h))$.

For a point $p \in \mathbb{R}^2$ and a direction u_i , let $p^{[i]}$ be the projection of p in the direction normal to u_i . Let $P^{[i]} = \langle p_1^{[i]}, \dots, p_n^{[i]} \rangle$ be the sorted sequence of projections of points in the direction normal to u_i . For each pair a, b such that $a + b \leq k$, we choose the interval $\delta_{a,b}^{[i]} = [p_a^{[i]}, p_{n-b}^{[i]}]$ and we place $O(1)$ equidistant points in this interval. See Figure 3. Let $L_{a,b}^{[i]}$ be the set of (oriented) lines in the direction normal to u_i and passing through these points. Set

$$L = \bigcup_{\substack{1 \leq i \leq h \\ a+b \leq k}} L_{a,b}^{[i]}.$$

We claim that L contains at least one separator line. Intuitively, let $u_i \in U$ be the direction closest to $\overrightarrow{c_1 c_2}$. Suppose p_a and p_{n-b} are the first and the last points of P in the direction u_i that lie inside $D_1 \cup D_2$. Since $|P \setminus (D_1 \cup D_2)| \leq k$, $a + b \leq k$. If $D_1 \cap D_2 = \emptyset$, then let q be the extreme points of D_1 in direction $\overrightarrow{c_1 c_2}$. Otherwise, let q be the first intersection point of $\partial D_1 \cap \partial D_2$ in direction u_i . Following the same argument as Sharir [19], one can argue that

$$\langle c_1 - q, u_i \rangle \geq \alpha \langle p_{n-b} - p_a, u_i \rangle,$$

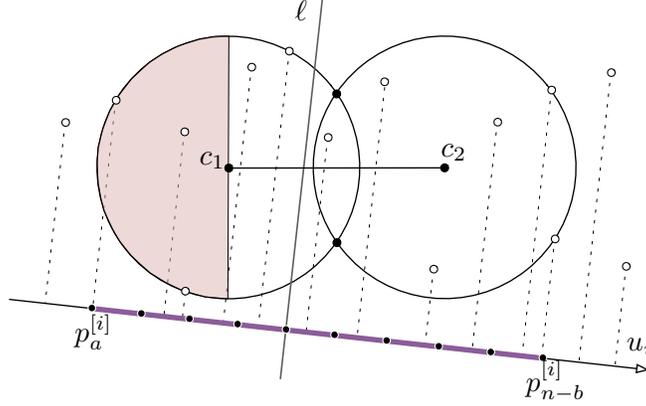


Figure 3: Let ℓ is a separator line for disks D_1 and D_2 .

where $\alpha \leq 1$ is a constant. Hence if at least 2α points are chosen in the interval $\delta_{a,b}^{[i]}$, then one of the lines in $L_{a,b}^{[i]}$ is a separator line. Omitting all the details, which are similar to the one in [19], we conclude the following.

Lemma 3.1. *We can compute in $O(k^2 n \log n)$ time a set L of $O(k^2)$ lines that contains a separator line.*

Let D_1, D_2 be a $(2, k)$ -center of P , let $\ell \in L$ be a line, and let $P^- \subseteq P$ be the set of points that lie in the left halfspace bounded by ℓ . We call D_1, D_2 a $(2, k)$ -center consistent with ℓ if $P^- \cap (D_1 \cup D_2) \subseteq D_1$, the center of D_1 lies to the left of ℓ , and ∂D_1 contains at least one point of P^- . We describe a decision algorithm that determines whether there is a $(2, k)$ -center of unit radius that is consistent with ℓ . Next, we describe an algorithm for computing a $(2, k)$ -center consistent with ℓ , which will lead to computing an optimal $(2, k)$ -center of P , provided there is a well-separated optimal $(2, k)$ -center of P .

Decision algorithm. Let $\ell \in L$ be a line. We describe an algorithm for determining whether there is a unit radius $(2, k)$ -center of P that is consistent with ℓ . Let P^- (resp. P^+) be the subset of points in P that lie in the left (resp. right) halfspace bounded by ℓ ; set $n^- = |P^-|$, $n^+ = |P^+|$. Suppose D_1, D_2 is a unit-radius $(2, k)$ -center of P consistent with ℓ , and let c_1, c_2 be their centers. Then $P^- \cap (D_1 \cup D_2) \subseteq D_1$ and $|P^- \cap D_1| \geq n^- - k$. For a subset $Q \subset P$, let $\mathcal{D}(Q) = \{D(q) \mid q \in Q\}$ where $D(q)$ is the unit disk centered at q . Let $\mathcal{D}^- = \mathcal{D}(P^-)$ and $\mathcal{D}^+ = \mathcal{D}(P^+)$. For a point $x \in \mathbb{R}^2$, let $\mathcal{D}_x^+ = \{D \in \mathcal{D}^+ \mid x \in D\}$. Since ∂D_1 contains a point of P^- and at most k points of P^- do not lie in D_1 , c_1 lies on an edge of $\mathcal{A}_{\leq k}(\mathcal{D}^-)$.

We first compute $\mathcal{A}_{\leq k}(\mathcal{D}^-)$ in $O(nk \log n)$ time. For each disk $D \in \mathcal{D}^+$, we compute the intersection points of ∂D with the edges of $\mathcal{A}_{\leq k}(\mathcal{D}^-)$. By Lemma 2.2, there are $O(nk)$ such intersection points, and these intersection points split each edge into *edgelets*. The total number of edgelets is also $O(nk)$. Using Lemma 2.2, we can compute all edgelets in time $O(nk \log n)$. All points on an edgelet γ lie in the same subset of disks of \mathcal{D}^+ , which we denote by \mathcal{D}_γ^+ . Let $P_\gamma^+ \subseteq P^+$ be the set of centers of disks in \mathcal{D}_γ^+ , and let $\kappa_\gamma = \lambda(\gamma, \mathcal{D}^-)$. A unit disk centered at

a point on γ contains P_γ^+ and all but κ_γ points of P^- . If at least $k' = k - \kappa_\gamma$ points of $P^+ \setminus P_\gamma^+$ can be covered by a unit disk, which is equivalent to $\mathcal{A}_{\leq k'}(\mathcal{D}^+ \setminus \mathcal{D}_\gamma)$ being nonempty, then all but k points of P can be covered by two disks of unit radius.

When we move from one edgelet γ of $\mathcal{A}_{\leq k}(\mathcal{D}^-)$ to an adjacent one γ' with σ as their common endpoint, then $\mathcal{D}_\gamma^+ = \mathcal{D}_{\gamma'}^+$ (if σ is a vertex of $\mathcal{A}_{\leq k}(\mathcal{D}^-)$), $\mathcal{D}_{\gamma'}^+ = \mathcal{D}_\gamma^+ \cup \{D\}$ (if $\sigma \in \partial D$ and $\gamma' \subset \{D\}$), or $\mathcal{D}_{\gamma'}^+ = \mathcal{D}_\gamma^+ \setminus \{D\}$ (if $\sigma \in \partial D$ and $\gamma \subset \mathcal{D}$). We therefore traverse the graph induced by the edgelets of $\mathcal{A}_{\leq k}(\mathcal{D})$ and maintain \mathcal{D}_γ^+ in the dynamic data structure described in Section 2 as we visit the edgelets γ of $\mathcal{A}_{\leq k}(\mathcal{D}^-)$. At each step we process an edgelet γ , insert or delete a disk into \mathcal{D}_γ^+ , and test whether $\mathcal{A}_{\leq j}(\mathcal{D}_\gamma^+) = \emptyset$ where $j = k - \lambda(\gamma, \mathcal{D}^-)$. If the answer is yes at any step, we stop. We spend $O(k^3 \log^2 n)$ time at each step, by Lemma 2.4. Since the number of edgelets is $O(nk)$, we obtain the following.

Lemma 3.2. *Let P be a set of n points in \mathbb{R}^2 , ℓ a line in L , and $0 \leq k \leq n$ an integer. We can determine in $O(nk^4 \log^2 n)$ time whether the radius of a $(2, k)$ -center of P consistent with ℓ is at most 1.*

Optimization algorithm. Let ℓ be a line in L . Let r^* be the smallest radius of a $(2, k)$ -center of P that is consistent with ℓ . Our goal is to compute a $(2, k)$ -center of P of radius r^* that is consistent with ℓ . We use the parametric search technique [16] — we simulate the decision algorithm generically at r^* and use the decision algorithm to resolve each comparison, which will be of the form: given $r_0 \in \mathbb{R}^+$, is $r_0 \leq r^*$? We simulate a parallel version of the decision procedure to reduce the number of times the decision algorithm is invoked to resolve a comparison. Note that we need to parallelize only those steps of the simulation that depend on r^* , i.e., that require comparing a value with r^* . Instead of simulating the entire decision algorithm, as in [11], we stop the simulation after computing the edgelets and return the smallest $(2, k)$ -center found so far, i.e., the smallest radius for which the decision algorithm returned “yes.” Since we stop the simulation earlier, we do not guarantee that we find the a $(2, k)$ -center of P of radius r^* that is consistent with ℓ . However, as argued by Eppstein [11], this is sufficient for our purpose.

Let P^-, P^+ be the same as in the decision algorithm. Let $\mathcal{D}^-, \mathcal{D}^+$ etc. be the same as above except that each disk is of radius r^* (recall that we do not know the value of r^*). We simulate the algorithm to compute the edgelets of $\mathcal{A}_{\leq k}(\mathcal{D}^-)$ as follows. First, we compute the $\leq k^{\text{th}}$ order farthest point Voronoi diagram of P^- in time $O(n \log n + nk^2)$. Let e be an edge of the diagram with points p and q of P^- as its neighbors, i.e., e is a portion of the bisector of p and q . Then for each point $x \in e$, the disk of radius $\|xp\|$ centered at x contains at least $n^- - k$ points of P^- . We associate an interval $\delta_e = \{\|xp\| \mid x \in e\}$. By definition, e corresponds to a vertex of $\mathcal{A}_{\leq k}(\mathcal{D}^-)$ if and only if $r^* \in \delta_e$; namely, if $\|xp\| = r^*$, for some $x \in e$, then x is a vertex of $\mathcal{A}_{\leq k}(\mathcal{D}^-)$, incident upon the edges that are portions of $\partial D(p)$ and $\partial D(q)$. Let X be the sorted sequence of the endpoints of the intervals. By doing a binary search on X and using the decision procedure at each step, we can find two consecutive endpoints between which r^* lies. We can now compute all edges e of the Voronoi diagram such that $r^* \in \delta_e$. We thus compute all vertices of $\mathcal{A}_{\leq k}(\mathcal{D}^-)$. Since we do not know r^* , we do not have actual coordinates of the vertices. We represent each vertex as a pair of points. Similarly, each edge is represented as a point $p \in P^-$, indicating that e lies in $\partial D(p)$. Once we have all the edges of $\mathcal{A}_{\leq k}(P^-)$, we can construct the graph induced by them and compute $O(k^2)$ x -monotone

unit-disk curves whose union is the set of edges in $\mathcal{A}_{\leq k}(P^-)$, using Lemma 2.3. Since this step does not depend on the value of r^* , we need not parallelize it. Let $\Xi = \{\xi_i, \dots, \xi_u\}$, $u = O(k^2)$, be the set of these curves.

Next, for each disk $D \in \mathcal{D}^+$ and for each $\xi_i \in \Xi$, we compute the edges of ξ_i that ∂D intersects, using a binary search. We perform these $O(nk^2)$ binary searches in parallel and use the decision algorithm at each step. Incorporating Cole's technique [9] in the binary search we need to invoke the decision procedure only $O(\log n)$ times. For an edge $e \in \mathcal{A}_{\leq k}(\mathcal{D})$, let $\mathcal{D}_e^+ \in \mathcal{D}$ be the set of disks whose boundaries intersect e . We sort the disks in \mathcal{D}_e^+ by the order in which their boundaries intersect e . By doing this in parallel for all edges and using a parallel sorting algorithm for each edge, we can perform this step by invoking the decision algorithm $O(\log n)$ times.

Putting pieces together. We repeat the optimization algorithm for all lines in L and return the smallest $(2, k)$ -center that is consistent with a line in L . The argument of Eppstein [11] implies that if there is a well-separated $(2, k)$ -center of P then the above algorithms return a $(2, k)$ -center of P . Hence, we conclude the following:

Lemma 3.3. *Let P be a set of n points in \mathbb{R}^2 and $0 \leq k \leq n$ an integer. If there is a well-separated $(2, k)$ -center of P , then the $(2, k)$ -center problem for P can be solved in $O(nk^6 \log^3 n)$ time.*

4 Nearly Concentric Disks

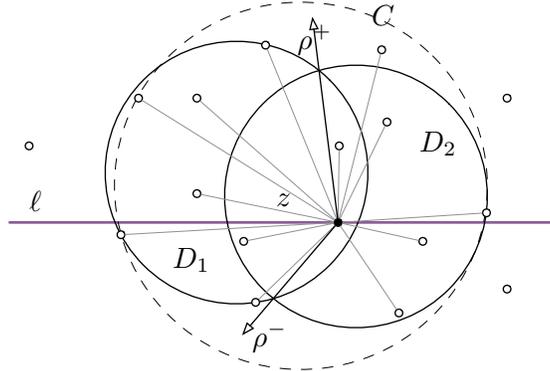


Figure 4: Two unit disks D_1 and D_2 of radius r^* with centers closer than a distance r^* . For the set of points P , shown as hollow points, all but 4 lie in $D_1 \cap D_2$. C is the circumcircle of $P \cap (D_1 \cup D_2)$ and $z \in D_1 \cap D_2$ is selected from C . The points $P \cap (D_1 \cup D_2)$ are sorted radially around z .

In this section we describe an algorithm for when the two disks D_1 and D_2 of the optimal solution are not well separated and are thus nearly concentric. More specifically, let c_1 and c_2 be the centers of D_1 and D_2 and let r^* be their radius. Then this section handles the case where $\|c_1 c_2\| \leq r^*$.

First, we find an *intersector point* z of D_1 and D_2 — a point that lies in $D_1 \cap D_2$. We show how z defines a set \mathcal{P} of $O(n^2)$ possible partitions of P into two sets, such that for one partition $P_{i,j}, P \setminus P_{i,j}$ the following holds: $(D_1 \cup D_2) \cap P = (D_1 \cap P_{i,j}) \cup (D_2 \cap (P \setminus P_{i,j}))$. Finally, we show how to search through the set \mathcal{P} in $O(k^7 n^{1+\delta})$ time, deterministically, for any $\delta > 0$, or in $O(k^7 n \log^3 n)$ expected time.

Finding an intersector point. Let C be the circumcircle of $P \cap (D_1 \cup D_2)$. Eppstein [11] shows that we can select $O(1)$ points inside C such that at least one, z , lies in $D_1 \cap D_2$. Consider two rays ρ^+ and ρ^- from z through the two points at $\partial D_1 \cap \partial D_2$. We can prove the following.

Lemma 4.1. *Let P be a set of n points in \mathbb{R}^2 . We can generate in $O(nk^3)$ time a family \mathcal{C} of $O(k^3)$ circles such that for any nearly concentric $(2, k)$ -center disks D_1, D_2 one of the circles in \mathcal{C} is the circumcircle of $P \cap (D_1 \cup D_2)$.*

Proof. Using an algorithm of Matoušek [14] in $O(k^3 n)$ time we can find the smallest circle that contains $n - k$ points of P . Briefly, the algorithm runs by finding the three points defining the circumcircle, removing each one in turn, and recursing until k points have been removed. Matoušek shows that if we keep track of which nodes in the recursion we reach and halt the recursion if we have seen that node before, then the size of the recursion tree is only $O(k^3)$. In the running of this algorithm we generate all circles which include exactly $n - j$ points of P for $0 \leq j \leq k$. We claim that one of these circles must be C .

If the initial circle is not C , then it must have at least one point on its boundary which is not in $P \cap (D_1 \cup D_2)$. At least one path of the recursion removes this point. Since we can reach the point set $P \cap (D_1 \cup D_2)$ in at most k steps, some step in this recursion must return C . \square

Corollary 4.1. *Let P be a set of n points in \mathbb{R}^2 , and let D_1 and D_2 describe the nearly concentric $(2, k)$ -center of P . In $O(k^3 n)$ time, we can generate $O(k^3)$ points such that at least one of them is an intersector point of D_1 and D_2 .*

Set of solutions, given an intersector point. Let z be an intersector point of D_1 and D_2 , and let ρ^+ and ρ^- be the two rays from z through $\partial D_1 \cap \partial D_2$. As in the case of well-separated disks, we choose a set $U \subseteq \mathbb{S}^1$ of h uniformly distributed directions. Because the angle between ρ^+ and ρ^- is at least some constant θ , we can set $h = 2\pi/\theta$ and this guarantees that, for at least one direction u from this set, a line ℓ_u drawn through z in the direction u separates ρ^- and ρ^+ . Let ρ^+ be above ℓ_u and ρ^- below; see Figure 4.

For each $u \in U$, we divide P into two disjoint sets. Let P^+ be the subset of P above ℓ_u and let P^- be the subset of P below ℓ_u . In what follows, we assume we are using the $u \in U$ which separates ρ^+ and ρ^- , in reality we repeat the following for all $u \in U$ and use the one with the best answer. Let $n^+ = |P^+|$ and $n^- = |P^-|$. We sort each set radially around z , clockwise. Label the points of P^+ as $\langle p_1^+, \dots, p_{n^+}^+ \rangle$ and the points of P^- as $\langle p_1^-, \dots, p_{n^-}^- \rangle$ to reflect this ordering. In the total circular ordering of P , $p_{n^+}^+$ is adjacent to p_1^- and $p_{n^-}^-$ is adjacent to p_1^+ . Note that the rays ρ^+ and ρ^- create a partition P_{i^*,j^*} of P . Let P_{i^*,j^*} be the union of $\{p_1^+, \dots, p_{i^*}^+\}$ and $\{p_{j^*}^-, \dots, p_{n^-}^-\}$, describing the points left of ρ^+ in P^+ and left of ρ^- in P^- , respectively. Thus, the (unknown) location of ρ^+ and ρ^- define one partition, P_{i^*,j^*} , from a set \mathcal{P} of $O(n^2)$ possible partitions.

For a point set X , define $r_t(X)$ as the radius of the smallest disk that can cover all but t points of X . For a partition $P_{i,j} \in \mathcal{P}$ and an integer $0 \leq t \leq k$, let $m_{i,j}^t = \max\{r_t(P_{i,j}), r_t(P \setminus P_{i,j})\}$. These values define $k+1$ matrices M^1, \dots, M^k , where the value at the cell at the intersection of the i th column and j th row, $M^t(i, j) = m_{i,j}^t$. The partition P_{i^*,j^*} that defines the $(2, k)$ -center D_1, D_2 , such that $P \cap (D_1 \cup D_2) = (P_{i^*,j^*} \cap D_1) \cup ((P \setminus P_{i^*,j^*}) \cap D_2)$, is $\arg \min_{P_{i,j}} \min_t m_{i,j}^t$. We note a couple of properties of each matrix M^t that will help search for the minimum $m_{i,j}^t$.

(P1) If $r^t(P_{i,j}) > r^t(P \setminus P_{i,j})$ then $m_{i,j}^t \leq m_{i',j'}^t$ for $i' \geq i$ and $j' \leq j$. These partitions only add points to $P_{i,j}$ and thus cannot decrease $r^t(P_{i,j})$. Similarly, if $r^t(P \setminus P_{i,j}) > r^t(P_{i,j})$, then $m_{i,j}^t < m_{i',j'}^t$ for $i' \leq i$ and $j' \geq j$.

(P2) Given a value r , if $r^t(P_{i,j}) > r$, then $m_{i',j'}^t > r$ for $i' \geq i$ and $j' \leq j$, and if $r^t(P \setminus P_{i,j}) > r$, then $m_{i',j'}^t > r$ for $i' \leq i$ and $j' \geq j$.

Deterministic solution. We now have the machinery to use a technique of Fredrickson and Johnson [12]. For each t , assume that the side length of M^t is $2^\tau + 1$, where $\tau = \log n + O(1)$. We divide M_t into a set $\mathcal{M}_{\tau-1}^t$ of four submatrices each of side length $2^{\tau-1} + 1$ that overlap along one row and one column. We call the cell at the intersection of this row and column the *center cell*. We calculate $r^t(P_{i,j})$ and $r^t(P \setminus P_{i,j})$ for the center cell $M^t(i, j)$. (This can be done in $O(k^3 n)$ time.) We can use **P1** to eliminate one of the four submatrices and we remove it from $\mathcal{M}_{\tau-1}^t$. We repeat this process. For each matrix in $\mathcal{M}_{\tau-1}^t$ we create four submatrices of side length $2^{\tau-2} + 1$ forming a set $\mathcal{M}_{\tau-2}^t$. For each submatrix in $\mathcal{M}_{\tau-1}^t$ we evaluate the center cell and prune one submatrix from $\mathcal{M}_{\tau-2}^t$. After $\tau' = O(\log n)$ iterations, $O(n)$ cells remain in the union of the submatrices in $\mathcal{M}_{\tau-\tau'}^t$ and we evaluate all of them. This algorithm takes $O(k^3 n^2 \log n)$ time, since each stage evaluates at most $O(n)$ centers cells because they can be connected along a monotone path through M_t .

However, this path through M_t can be interpreted as a sequence of insertions and deletions from $P_{i,j}$ and $P \setminus P_{i,j}$. So we can use an algorithm of Agarwal and Matoušek [1] to, in each stage, evaluate all cells along the path in $O(n \cdot k^3 n^\delta)$ time, for any $\delta > 0$. Agarwal and Matoušek's data structure can maintain the value of the radius of the smallest inclosing disk under insertions and deletions in $O(n^\delta)$ time per update. Each step in the path is one update, and then searching through the $O(k^3)$ nodes of the recursion tree of all possible outliers — each requires $O(1)$ updates — takes $O(k^3 n^\delta)$ time per cell. This reduces the running time to $O(k^3 n^{1+\delta})$, for any $\delta > 0$.

Running this algorithm for $O(k^3)$ possible intersector points and on $k + 1$ matrices takes $O(n \cdot k^7 n^\delta)$ time, for any $\delta > 0$.

Lemma 4.2. *Let P be a set of n points in \mathbb{R}^2 and $0 \leq k \leq n$ an integer. If the minimal radius $(2, k)$ -center of P is nearly concentric, then we can calculate it in $O(nk^7 n^\delta)$ time, for any $\delta > 0$.*

Randomized solution. We can improve the dependence on n by using the dynamic data structure in Section 2 and **P2**. As before, in the x th phase, we maintain sets $\mathcal{M}_{\tau-x}^t$ of submatrices of M^t , each of side length $2^{\tau-x} + 1$. Each submatrix is divided into four submatrices of side length $2^{\tau-x-1} + 1$, forming a set $\mathcal{M}_{\tau-x-1}^t$. To reduce the size of $\mathcal{M}_{\tau-x-1}^t$, we choose a random

center cell $M^t(i, j)$ from the submatrices in $\mathcal{M}_{\tau-x}^t$ and evaluate $r = m_{i,j}^t$ in $O(k^3 n)$ time. For each other center cell $M^t(i', j')$ of a submatrix in $\mathcal{M}_{\tau-x}^t$, with probability $1/2$, $m_{i',j'}^t > r$. When this is true, using **P2**, we can remove a submatrix from $\mathcal{M}_{\tau-x-1}^t$. In expectation, this reduces the size of the union of the cells in the submatrices in $\mathcal{M}_{\tau-x-1}^t$ by $1/8$. Eppstein [11] shows this process expects to find $\arg \min_{P_{i,j}} m_{i,j}^t$ in $O(\log n)$ iterations.

On each iteration we use the dynamic data structure described in Section 2. For $O(n)$ insertions and deletions, it can label each remaining center cell in $O(k^3 n \log^2 n)$ time. Finding $\arg \min_{P_{i,j}} m_{i,j}^t$ for $O(k^3)$ possible intersector points and for $k + 1$ values of t , this algorithm runs in expected $O(nk^7 \log^3 n)$ time.

Lemma 4.3. *Let P be a set of n points in \mathbb{R}^2 and $0 \leq k \leq n$ an integer. If the minimal radius $(2, k)$ -center of P is nearly concentric, then we can calculate it in $O(nk^7 \log^3 n)$ expected time.*

And combining this with Lemma 3.3 we can state the following.

Theorem 4.1. *Let P be a set of n points in \mathbb{R}^2 and $0 \leq k \leq n$ an integer. In $O(nk^7 \log^3 n)$ expected time, we can solve the $(2, k)$ -center problem for P .*

5 Acknowledgements

We thank Sarel Har-Peled for posing the problem and for several helpful discussions.

References

- [1] P. K. Agarwal and J. Matoušek. Dynamic half-space range reporting and its applications. *Algorithmica*, 13(4):325–345, 1995.
- [2] P. K. Agarwal and M. Sharir. Planar geometric locations problems. *Algorithmica*, 11:185–195, 1994.
- [3] P. K. Agarwal and M. Sharir. Efficient algorithms for geometric optimization. *ACM Computing Surveys (CSUR)*, 30(4):412–458, 1998.
- [4] T. Chan. More planar two-center algorithms. *Computational Geometry: Theory and Applications*, 13:189–198, 1999.
- [5] T. Chan. Low-dimensional linear programming with violations. *SIAM Journal of Computing*, 34:879–893, 2005.
- [6] T. Chan. On the bichromatic k -set problem. In *19th Annual Symposium on Discrete Algorithms*, 2007.
- [7] K. L. Clarkson. A bound on local minima of arrangements that implies the upper bound theorem. *Discrete and Computational Geometry*, 10:427–433, 1993.
- [8] K. L. Clarkson and P. W. Shor. Applications of random sampling in geometry, II. *Discrete and Computational Geometry*, 4(1):387–421, 1989.

- [9] R. Cole. Slowing down sorting networks to obtain faster sorting algorithms. *Journal of ACM*, 34:200–208, 1987.
- [10] Z. Drezner and H. Hamacher. *Facility Location: Applications and Theory*. Springer, 2002.
- [11] D. Eppstein. Faster construction of planar two-centers. In *8th ACM-SIAM Symposium on Discrete Algorithms*, pages 131–138, 1997.
- [12] G. N. Frederickson and D. B. Johnson. The complexity of selection and ranking in $x + y$ and matrices with sorted columns. *Journal of Computers and Systems Science*, 24(2):197–208, 1982.
- [13] D. S. Hochbaum, editor. *Approximation Algorithms for NP-hard Problems*. PWS Publishing Company, 1995.
- [14] J. Matoušek. On geometric optimization with few violated constraints. *Discrete and Computational Geometry*, 1995.
- [15] J. Matoušek, E. Welzl, and M. Sharir. A subexponential bound for linear programming and related problems. *Algorithmica*, 16:498–516, 1996.
- [16] N. Megiddo. Linear-time algorithms for linear programming in \mathbb{R}^3 and related problems. *SIAM Journal of Computing*, 12:759–776, 1983.
- [17] N. Megiddo and K. J. Supowit. On the complexity of some common geometric location problems. *SIAM Journal of Computing*, 12:759–776, 1983.
- [18] M. Sharir. On k -sets in arrangement of curves and surfaces. *Discrete and Computational Geometry*, 6:593–613, 1991.
- [19] M. Sharir. A near-linear time algorithm for the planar 2-center problem. *Discrete and Computational Geometry*, 18(2):125–134, 1997.
- [20] M. Sharir and E. Welzl. Rectilinear and polygonal p -piercing and p -center problems. In *12th Annual Symposium on Computational Geometry*, pages 122–132, 1996.