

Enclosing Points with Geometric Objects

Timothy M. Chan*

Qizheng He†

Jie Xue‡

Abstract

Let X be a set of points in \mathbb{R}^2 and \mathcal{O} be a set of geometric objects in \mathbb{R}^2 , where $|X| + |\mathcal{O}| = n$. We study the problem of computing a minimum subset $\mathcal{O}^* \subseteq \mathcal{O}$ that encloses all points in X . Here a point $x \in X$ is enclosed by \mathcal{O}^* if it lies in a bounded connected component of $\mathbb{R}^2 \setminus (\bigcup_{O \in \mathcal{O}^*} O)$. We propose two algorithmic frameworks to design polynomial-time approximation algorithms for the problem. The first framework is based on sparsification and min-cut, which results in $O(1)$ -approximation algorithms for unit disks, unit squares, etc. The second framework is based on LP rounding, which results in an $O(\alpha(n) \log n)$ -approximation algorithm for segments, where $\alpha(n)$ is the inverse Ackermann function, and an $O(\log n)$ -approximation algorithm for disks.

1 Introduction

Studying problems related to plane obstacles is a popular topic in computational geometry. In the common setting of such problems, we are given a set of geometric objects in the plane as *obstacles*. An obstacle blocks all paths in the plane it intersects. Various optimization problems have been investigated in this setting. For example, a line of research [3, 10, 13, 14, 20, 28] focused on finding shortest paths between two points amidst the given obstacles. The obstacle-removal problem [7, 17, 18, 19] asks for a minimum subset of obstacles such that after removing these obstacles, two specified points in the plane have a path between them. The point-separation problem [6, 11, 18] aims to select a minimum subset of obstacles that separate a given set X of points, i.e., block all paths between two points in X .

In this paper, we study a natural problem related to plane obstacles, which we call **ENCLOSING-ALL-POINTS**. We say a point x in the plane is *enclosed* by a set \mathcal{O} of (compact) geometric objects if x lies in a bounded connected component of $\mathbb{R}^2 \setminus (\bigcup_{O \in \mathcal{O}} O)$; in other words, any curve connecting x with the point $(+\infty, 0)$ at infinity must intersect with at least one object in \mathcal{O} . An example is shown in Figure 1. The problem simply aims to compute a minimum subset of obstacles to enclose all input points.

ENCLOSING-ALL-POINTS

Input: A set X of points and a set \mathcal{O} of geometric objects in \mathbb{R}^2 .

Output: A minimum subset $\mathcal{O}^* \subseteq \mathcal{O}$ that enclose all points in X .

ENCLOSING-ALL-POINTS is closely related to the aforementioned point-separation problem. In fact, it is a special case of the *generalized* point-separation problem studied in [18], in which we

*University of Illinois at Urbana-Champaign, USA. tmc@illinois.edu.

†University of Illinois at Urbana-Champaign, USA. qizheng6@illinois.edu.

‡New York University Shanghai, China. jiexue@nyu.edu.

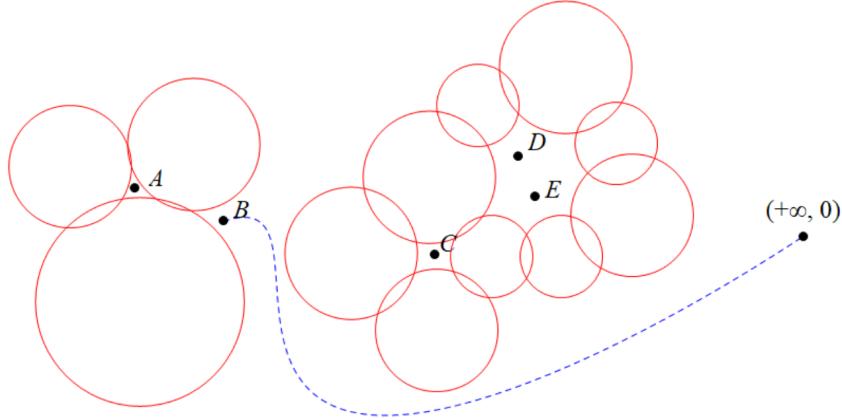


Figure 1: The set of red disks enclose points A , C , D and E , but does not enclose the point B .

want to use a minimum number of obstacles to separate a set of point-pairs $\{(s_1, t_1), \dots, (s_r, t_r)\}$ in the plane, i.e., block all paths between s_i and t_i for all $i \in [r]$. Enclosing a point-set X is equivalent to separating the point-pairs in $\{(x, z) : x \in X\}$, where $z = (+\infty, 0)$ is the point at infinity. Kumar et al. [18] showed that when the number of point-pairs to be separated is fixed, the generalized point-separation problem is polynomial-time solvable for any connected obstacles. However, this does not solve the ENCLOSING-ALL-POINTS problem, since in our problem the set X is a part of the input (when the number of point-pairs is not fixed, no nontrivial algorithms for the generalized point-separation problem were known). To the best of our knowledge, ENCLOSING-ALL-POINTS has not been studied before.

In this paper, we propose two general algorithmic frameworks to design polynomial-time approximation algorithms for ENCLOSING-ALL-POINTS.

Approach based on sparsification and min-cut. Our first framework solves ENCLOSING-ALL-POINTS via two steps. The first step computes a *sparse* subset $\mathcal{O}' \subseteq \mathcal{O}$ of obstacles such that an optimal solution using the obstacles in \mathcal{O}' is an $O(1)$ -approximation solution of the original problem for \mathcal{O} . Roughly speaking, here “sparse” means that any point in the plane stabs $O(1)$ obstacles in \mathcal{O}' . The second step reduces the problem on \mathcal{O}' to a min-cut problem in a graph (which is polynomial-time solvable) by losing a constant factor in cost. This framework results in $O(1)$ -approximation algorithms for the problem with similarly-sized fat pseudo-disks, e.g., unit-disks, unit squares, etc.

Approach based on LP rounding. Our second framework is based on LP rounding. Compared to the first one, this framework applies to more general types of obstacles, while giving worse approximation ratios. Specifically, it results in an $O(\alpha(n) \log n)$ -approximation algorithm for ENCLOSING-ALL-POINTS with line segments, where $\alpha(n)$ denotes the inverse Ackermann function, and an $O(\log n)$ -approximation algorithm for ENCLOSING-ALL-POINTS with (general) disks. The framework also generalizes to the problem with curves that pairwise intersect $s = O(1)$ times, with an approximation ratio $\alpha(n)^{O(\alpha(n)^{s-1})} \log n$. Although LP rounding is a common technique widely used for designing approximation algorithms, how to apply it for ENCLOSING-ALL-POINTS is totally non-obvious. In fact, for our problem, new ideas are required in both the LP formulation and the rounding scheme.

1.1 Related work

To the best of our knowledge, **ENCLOSING-ALL-POINTS** has not been studied before. Here we briefly summarize the literature for two problems that are closely related to **ENCLOSING-ALL-POINTS**, the point-separation problem and the obstacle-removal problem.

Point separation. Given a set X of points and a set \mathcal{O} of obstacles in \mathbb{R}^2 , the point-separation problem asks for a minimum subset $\mathcal{O}^* \subseteq \mathcal{O}$ that separates all points in X . This problem has applications in barrier coverage with wireless sensors. Gibson et al. [11] showed the NP-hardness of the problem for unit disks by reducing from planar multiterminal-cut, and presented an $(9 + \varepsilon)$ -approximation algorithm that works for (general) disks. Later, Cabello and Giannopoulos [6] designed the first polynomial-time exact algorithm for point-separation with arbitrary connected curves for the case $|X| = 2$, which runs in $O(n^3)$ time. They also showed NP-hardness of the problem with unit circles or orthogonal segments, via a reduction from planar-3-SAT. Recently, Kumar et al. [18] generalizes the algorithm of Cabello and Giannopoulos [6] and showed that point-separation with arbitrary connected curves can be solved in $n^{O(|X|)}$ time.

Obstacle removal. Given two points s, t and a set \mathcal{O} of obstacles in \mathbb{R}^2 , the obstacle problem asks for a minimum subset $\mathcal{O}^* \subseteq \mathcal{O}$ such that s and t are not separated by $\mathcal{O} \setminus \mathcal{O}^*$. This problem is also sometimes called computing the *barrier resilience* [7, 19], where a real-world application is to compute the minimum number of sensors that need to be deactivated, so that there is a path between s and t that does not intersect any of the active sensors. It also has applications in robotics [8, 9]. Obstacle-removal was shown to be NP-hard for arbitrary line segments [2], for unit segments [25, 26], and for certain types of fat regions with bounded ply (such as axis-aligned rectangles of aspect ratio $1 : 1 + \varepsilon$ and $1 + \varepsilon : 1$) [16]. For approximation results, Bereg and Kirkpatrick provided a 3-approximation algorithm for unit disks [5]. Bandyapadhyay et al. [4] presented an $O(\sqrt{n})$ -approximation algorithm for pseudodisks and rectilinear polygons. Kumar et al. [17] further designed an $O(1)$ -approximation algorithm for any well-behaved objects, such as polygons or splines. Their arguments were later simplified by Kumar et al. [18].

2 Preliminaries

Curves and homotopy. A *curve* in the plane is a continuous map $\gamma : [0, 1] \rightarrow \mathbb{R}^2$. For convenience, sometimes the term “curve” also refers to the image of such a map. If $\gamma(0) = \gamma(1)$, we say γ is a *closed curve* and define its *base point* as $\gamma(0) = \gamma(1)$. A segment between two points a and b naturally defines two curves, one from a to b and the other from b to a . For two curves $\gamma, \gamma' : [0, 1] \rightarrow \mathbb{R}^2$ with $\gamma(1) = \gamma'(0)$, we can concatenate them to obtain a curve from $\gamma(0)$ to $\gamma'(1)$.

Two curves $\gamma, \gamma' : [0, 1] \rightarrow \mathbb{R}^2$ with $\gamma(0) = \gamma'(0)$ and $\gamma(1) = \gamma'(1)$ are *homotopic* in a region $R \subseteq \mathbb{R}^2$ if the images of γ, γ' lie in R and one can deform γ continuously to γ' inside R without changing the two endpoints $\gamma(0)$ and $\gamma(1)$. A closed curve γ is *contractible* in $R \subseteq \mathbb{R}^2$ if γ and the trivial curve μ are homotopic in R , where $\mu(i) = \gamma(0) = \gamma(1)$ for all $i \in [0, 1]$; and *non-contractible* in R otherwise.

Union complexity of geometric objects. For a type of geometric objects, its *union complexity* is defined as the combinatorial complexity of the boundary of the union of n such objects. It is known that disks have linear union complexity [15].

For our applications, it suffices to look at the combinatorial complexity of the outer face in the arrangement of the objects, which is upper-bounded by the union complexity of the objects. So for

disks, this complexity is $O(n)$. Pollack et al. [21] proved that for n arbitrary line segments in \mathbb{R}^2 , the complexity of the outer face in the arrangement is bounded by $O(n\alpha(n))$ (by Davenport-Schinzel sequences), where $\alpha(n)$ is the inverse Ackermann function.

More generally, for curves that pairwise intersect only at most $s = O(1)$ times, a similar bound holds: the combinatorial complexity of the outer face in the arrangement is $O(\lambda_{s+2}(n))$, where the function $\lambda_s(n) = n\alpha(n)^{O(\alpha(n)^{s-3})}$ is the maximum length of an (n, s) Davenport-Schinzel sequence [1, 12]. For any constant s , $\lambda_s(n)$ is almost linear in n .

Point-in-polygon check. Given a point q and a polygon P , the point-in-polygon problem asks whether q lies inside, outside, or on the boundary of P . A standard method for solving this problem is to use the ray-casting algorithm [24]: cast a ray \vec{r} starting from q and going in any fixed direction, if \vec{r} intersects P an even number of times, then q is outside P ; otherwise if \vec{r} intersects P an odd number of times, then q is inside P . (For our applications, we consider the points on the boundary of P as inside P .)

Another method for point-in-polygon check is by computing the *winding number* of q with respect to P , denote as $\text{wind}(q, P)$. The idea is to modify the ray-casting algorithm as follows: first orient the edges of P in counter-clockwise direction. Initialize a counter c_q with value 0. Cast a ray \vec{r}_q starting from q and going in any fixed direction, for each edge \vec{e} of P intersecting \vec{r}_q , if \vec{e} crosses \vec{r}_q in counter-clockwise direction, then increase c_q by 1; otherwise if \vec{e} crosses \vec{r}_q in clockwise direction, decrease c_q by 1. In the end, q is inside P iff the counter c_q is nonzero. In particular, when P is a simple polygon, q is inside P iff $c_q = 1$. Intuitively, if q is outside P , then the weighted crossings between the ray \vec{r}_q and the polygon P will cancel each other. An example for this process is shown in Fig. 2.

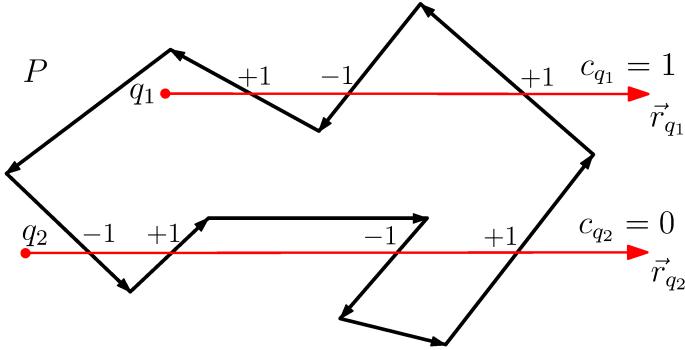


Figure 2: The winding numbers with respect to a (simple) polygon P . The point q_1 has winding number $c_{q_1} = 1$, and is inside P ; the point q_2 has winding number $c_{q_2} = 0$, and is outside P .

3 Approach based on sparsification and min-cut

In this section, we discuss our algorithmic framework based on sparsification and min-cut. We present our algorithm for unit disks. It generalizes to any similarly-sized fat pseudo-disks, e.g., unit squares or more generally translates of a fixed convex body; see Appendix A.

Let X be a set of points in \mathbb{R}^2 and \mathcal{D} be a set of unit disks. For a subset $\mathcal{D}' \subseteq \mathcal{D}$, denote by $G[\mathcal{D}']$ the intersection graphs of the unit disks in \mathcal{D}' and by $\text{Int}(\mathcal{D}')$ the edge set of $G[\mathcal{D}']$, i.e., the set of pairs (D, D') of unit disks in \mathcal{D}' that intersect each other. For each unit disk $D \in \mathcal{D}$, let $\text{ctr}(D)$ denote the center of D . For every edge $(D, D') \in \text{Int}(\mathcal{D})$, we have $D \cap D' \neq \emptyset$ and let

$\sigma(D, D')$ be the curve from $\text{ctr}(D)$ to $\text{ctr}(D')$ defined by the segment between $\text{ctr}(D)$ and $\text{ctr}(D')$; note that $\sigma(D, D') \subseteq D \cup D'$. With a bit abuse of notation, we also use $\sigma(D, D')$ to denote the segment between $\text{ctr}(D)$ and $\text{ctr}(D')$. For a path $\phi = (D_1, \dots, D_r)$ in $G[\mathcal{D}]$, we define γ_ϕ as the curve obtained by concatenating the curves $\sigma(D_1, D_2), \dots, \sigma(D_{r-1}, D_r)$ in order. If ϕ is a cycle, i.e., $D_1 = D_r$, then γ_ϕ is a closed curve. The following lemma gives a characterization of the subsets of \mathcal{D} enclosing a point.

Lemma 1. *Let $\mathcal{D}_0 \subseteq \mathcal{D}$ and $x \in \mathbb{R}^2$. The following statements are equivalent.*

- (i) \mathcal{D}_0 encloses x .
- (ii) $G[\mathcal{D}_0]$ contains a cycle ϕ such that γ_ϕ is non-contractible in $\mathbb{R}^2 \setminus \{x\}$.
- (iii) $\{\sigma(D, D') : (D, D') \in \text{Int}(\mathcal{D}_0)\}$ encloses x .

Proof. Clearly, (iii) implies (i) since $\bigcup_{(D, D') \in \text{Int}(\mathcal{D}_0)} \sigma(D, D') \subseteq \bigcup_{D \in \mathcal{D}_0} D$. We shall show that (i) \Rightarrow (ii) \Rightarrow (iii). To see (i) implies (ii), suppose \mathcal{D}_0 encloses x . Let U denote the union of all unit disks in \mathcal{D}_0 . The outer boundary of U consists of several simple curves; one of these curves encloses x and we denote it by ξ . Note that ξ is non-contractible in $\mathbb{R}^2 \setminus \{x\}$ by Jordan curve theorem. Suppose ξ consists of circular arcs ξ_1, \dots, ξ_r (sorted in the order they appear on ξ). Set $\xi_0 = \xi_r$ for convenience. Let D_0, D_1, \dots, D_r be the unit disks in \mathcal{D}_0 contributing the circular arcs $\xi_0, \xi_1, \dots, \xi_r$, respectively. Then $\phi = (D_0, D_1, \dots, D_r)$ is a cycle in $G[\mathcal{D}_0]$. It suffices to show that γ_ϕ is non-contractible in $\mathbb{R}^2 \setminus \{x\}$. For $i \in [r]$, let p_i be the intersection point of ξ_{i-1} and ξ_i , and μ_i (resp., μ'_i) be the curve from p_i to $\text{ctr}(D_i)$ (resp., from $\text{ctr}(D_i)$ to p_i) defined by the segment between $\text{ctr}(D_i)$ and p_i . Consider the closed curve η obtained by concatenating $\xi_0, \mu_1, \mu'_1, \xi_1, \mu_2, \mu'_2, \xi_2, \dots, \xi_{r-1}, \mu_r, \mu'_r$ in order. Clearly, η is homotopic to ξ in $\mathbb{R}^2 \setminus \{x\}$ (when picking the same point as the base points of η and ξ), since concatenating μ_i and μ'_i results in a contractible closed curve for every $i \in [r]$. Since ξ is non-contractible in $\mathbb{R}^2 \setminus \{x\}$, so is η . On the other hand, we observe that η is homotopic to γ_ϕ in $\mathbb{R}^2 \setminus \{x\}$. To see this, let η_i be the curve from $\text{ctr}(D_i)$ to $\text{ctr}(D_{i+1})$ obtained by concatenating μ'_i, ξ_i, μ_{i+1} in order, for $i \in [r]$. Then η is the concatenation of η_1, \dots, η_r . Now note that η_i is homotopic to $\sigma(D_i, D_{i+1})$ in $D_i \cup D_{i+1}$ (and hence in $\mathbb{R}^2 \setminus \{x\}$), as $D_i \cup D_{i+1}$ is simply-connected. Thus, η is homotopic to γ_ϕ in $\mathbb{R}^2 \setminus \{x\}$, which implies that γ_ϕ is also non-contractible in $\mathbb{R}^2 \setminus \{x\}$.

To see (ii) implies (iii), suppose $G[\mathcal{D}_0]$ contains a cycle ϕ such that γ_ϕ is non-contractible in $\mathbb{R}^2 \setminus \{x\}$. It suffices to show that the image of γ_ϕ encloses x , since the image of γ_ϕ is contained in $\bigcup_{(D, D') \in \text{Int}(\mathcal{D}_0)} \sigma(D, D')$. Consider a point $o \in \mathbb{S}^2$ on the sphere \mathbb{S}^2 . Let $f : \mathbb{R}^2 \rightarrow \mathbb{S}^2 \setminus \{o\}$ be the natural homeomorphism. Here o is viewed as the point at infinity of the plane. The curve γ_ϕ corresponds to a curve on $\mathbb{S}^2 \setminus \{o\}$ by the homeomorphism f , which we denote by γ'_ϕ . Assume x is not enclosed by the image of γ_ϕ . Then on \mathbb{S}^2 there exists a simple curve τ connecting $f(x)$ and o , which is disjoint from the image of γ'_ϕ . Note that $\mathbb{S}^2 \setminus \tau$ is simply-connected. Therefore, γ'_ϕ is contractible in $\mathbb{S}^2 \setminus \tau$ and hence contractible in $\mathbb{S}^2 \setminus \{x, o\}$. It follows that γ_ϕ is contractible in $\mathbb{R}^2 \setminus \{x\}$, contradicting our assumption. \square

3.1 Sparsification step

We construct a grid Γ in the plane with $\frac{1}{2} \times \frac{1}{2}$ square cells. For each cell $\square \in \Gamma$, we write $\mathcal{D}_\square = \{D \in \mathcal{D} : \text{ctr}(D) \in \square\}$. For $\mathcal{D}' \subseteq \mathcal{D}$, let $\text{opt}(X, \mathcal{D}')$ denote the minimum number of obstacles in \mathcal{D}' needed to enclose X . The goal of this section is to prove the following lemma.

Lemma 2. *There exists a constant $c > 0$ such that one can compute in polynomial time a subset $\mathcal{D}' \subseteq \mathcal{D}$ satisfying $|\mathcal{D}' \cap \mathcal{D}_\square| \leq c$ for all $\square \in \Gamma$ and $\text{opt}(X, \mathcal{D}') \leq c \cdot \text{opt}(X, \mathcal{D})$.*

A pair (\square, \square') of grid cells in Γ is *relevant* if there exist $D \in \mathcal{D}_\square$ and $D' \in \mathcal{D}_{\square'}$ such that $D \cap D' \neq \emptyset$. Note that for each $\square \in \Gamma$, there are $O(1)$ cells $\square' \in \Gamma$ such that (\square, \square') is relevant. To construct the desired \mathcal{D}' in Lemma 2, we consider all relevant pairs. For each relevant pair (\square, \square') , we include in \mathcal{D}' some unit disks in $\mathcal{D}_\square \cup \mathcal{D}_{\square'}$ as follows. Let $Y \subseteq X$ be the set of points enclosed by $\mathcal{D}_\square \cup \mathcal{D}_{\square'}$. Also, let $\text{Int}^* \subseteq \text{Int}(\mathcal{D}_\square \cup \mathcal{D}_{\square'})$ consist of the edges (D, D') where $D \in \mathcal{D}_\square$ and $D' \in \mathcal{D}_{\square'}$. If $Y = \emptyset$, we arbitrarily pick $(D, D') \in \text{Int}^*$, and include D, D' in \mathcal{D}' . If $Y \neq \emptyset$, we shall show the existence of at most four unit disks in $\mathcal{D}_\square \cup \mathcal{D}_{\square'}$ which enclose all points in Y ; then we include them in \mathcal{D}' .

Consider an edge $(D, D') \in \text{Int}^*$. The line containing the segment $\sigma(D, D')$ partitions the plane into two halfplanes H_L and H_R , where H_L (resp., H_R) is to the left (resp., right) of the line with respect to the direction from $\text{ctr}(D)$ to $\text{ctr}(D')$. Define $Y_L(D, D') = Y \cap H_L$ and $Y_R(D, D') = Y \cap H_R$. We have the following key observation.

Observation 3. *Let $(D, D'), (E, E') \in \text{Int}^*$.*

- (i) *Either $Y_L(D, D') \subseteq Y_L(E, E')$ or $Y_L(E, E') \subseteq Y_L(D, D')$.*
- (ii) *Either $Y_R(D, D') \subseteq Y_R(E, E')$ or $Y_R(E, E') \subseteq Y_R(D, D')$.*

Proof. Let Q be the convex hull of $\square \cup \square'$. We first show that $Y \subseteq Q \setminus (\square \cup \square')$. Note that $Y \cap (\square \cup \square') = \emptyset$, since \square is contained in any unit disk in \mathcal{D}_\square and \square' is contained in any unit disk in $\mathcal{D}'_{\square'}$. So it suffices to show $Y \subseteq Q$. Suppose there exists a point $a \in Y \setminus Q$. Since Q is convex, there exists a line ℓ such that a and Q lie on different sides of ℓ . Consider the ray ψ shot from a that is perpendicular to ℓ and does not intersect ℓ . For any point b lying on the other side of ℓ than a , we have $\text{dist}(a, b) \geq \text{dist}(a', b)$ for all $a' \in \psi$. Therefore, if a unit disk centered at the other side of ℓ than a does not contain a , then it does not contain any point on ψ and is thus disjoint from ψ . Since a is not contained in any unit disk in $\mathcal{D}_\square \cup \mathcal{D}_{\square'}$ and the centers of the unit disks in $\mathcal{D}_\square \cup \mathcal{D}_{\square'}$ all lie on the other side of ℓ than a , we know that ψ is disjoint from all unit disks in $\mathcal{D}_\square \cup \mathcal{D}_{\square'}$. This implies that a is not enclosed by $\mathcal{D}_\square \cup \mathcal{D}_{\square'}$, which contradicts the fact $a \in Y$. Thus, $Y \subseteq Q$ and $Y \subseteq Q \setminus (\square \cup \square')$.

To prove the lemma, notice that (i) and (ii) are symmetric. Thus, it suffices to show (i). The segment $\sigma(D, D')$ partitions the region $Q \setminus (\square \cup \square')$ into two parts, the left part L_1 and the right part R_1 (with respect to the direction from $\text{ctr}(D)$ to $\text{ctr}(D')$), where L_1 contains $Y_L(D, D')$. See the left figure of Figure 3. Similarly, $\sigma(E, E')$ also partitions $Q \setminus (\square \cup \square')$ into the left part L_2 and the right part R_2 , where L_2 contains $Y_L(E, E')$. If $\sigma(D, D') \setminus (\square \cup \square')$ and $\sigma(E, E') \setminus (\square \cup \square')$ do not intersect, then either $L_1 \subseteq L_2$ and or $L_2 \subseteq L_1$. See the middle figure of Figure 3. We then have either $Y_L(D, D') \subseteq Y_L(E, E')$ or $Y_L(E, E') \subseteq Y_L(D, D')$, which implies (i). Next, suppose $\sigma(D, D') \setminus (\square \cup \square')$ and $\sigma(E, E') \setminus (\square \cup \square')$ intersect at the point o . Let \triangle be the triangle with vertices $\text{ctr}(D), \text{ctr}(E), o$ and \triangle' be the triangle with vertices $\text{ctr}(D'), \text{ctr}(E'), o$. See the right figure of Figure 3. One of \triangle and \triangle' contains $L_1 \setminus L_2$, while the other one contains $L_2 \setminus L_1$. Since $o \in \sigma(D, D')$, we have $o \in D$ or $o \in D'$; without loss of generality, assume $o \in D$. Then we have $\triangle \subseteq D$, because all the three vertices of \triangle are contained in D . It follows that $Y \cap \triangle = \emptyset$, since no point in X lies in D . Therefore, we have either $Y \cap (L_1 \setminus L_2) = \emptyset$ (if \triangle contains $L_1 \setminus L_2$) or $Y \cap (L_2 \setminus L_1) = \emptyset$ (if \triangle contains $L_2 \setminus L_1$). The former implies $Y_L(D, D') \subseteq Y_L(E, E')$ while the latter implies $Y_L(E, E') \subseteq Y_L(D, D')$. \square

By (i) of the above observation, there exists $(D_L, D'_L) \in \text{Int}^*$ such that $Y_L(D, D') \subseteq Y_L(D_L, D'_L)$ for $(D, D') \in \text{Int}^*$. Similarly, by (ii) of the above observation, there exists $(D_R, D'_R) \in \text{Int}^*$ such that $Y_R(D, D') \subseteq Y_R(D_R, D'_R)$ for all $(D, D') \in \text{Int}^*$. We then include the four unit disks D_L, D'_L, D_R, D'_R in \mathcal{D}' .

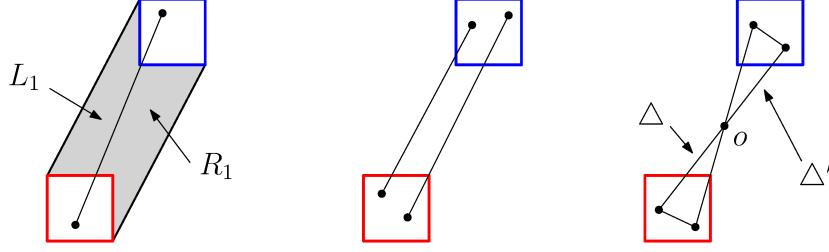


Figure 3: Illustration of the proof of Observation 3 — \square is red, \square' is blue, and $Q \setminus (\square \cup \square')$ is grey.

Observation 4. All points in Y are enclosed by $\{D_L, D'_L, D_R, D'_R\}$.

Proof. Consider a point $y \in Y$. We first show that $y \in Y_L(D_L, D'_L)$ and $y \in Y_R(D_R, D'_R)$. Without loss of generality, we only need to show $y \in Y_L(D_L, D'_L)$. By the choice of (D_L, D'_L) , we have $Y_L(D, D') \subseteq Y_L(D_L, D'_L)$ for all $(D, D') \in \text{Int}^*$. Therefore, it suffices to show the existence of $(D, D') \in \text{Int}^*$ such that $y \in Y_L(D, D')$. Since y is enclosed by $\mathcal{D}_\square \cup \mathcal{D}_{\square'}$, it is also enclosed by $\{\sigma(D, D') : (D, D') \in \text{Int}(\mathcal{D}_\square \cup \mathcal{D}_{\square'})\}$ by Lemma 1. As observed in the proof of Observation 3, we have $y \in Q \setminus (\square \cup \square')$ where Q is the convex hull of $\square \cup \square'$. The segments $\sigma(D, D')$ for $(D, D') \in \text{Int}^*$ decompose the region $Q \setminus (\square \cup \square')$ into small faces, among which there are two faces incident to the boundary of Q (which we call *boundary faces*). If $y \in Y_R(D, D')$ for all $(D, D') \in \text{Int}^*$, then y lies in a boundary face. In this case, y is not enclosed by $\{\sigma(D, D') : (D, D') \in \text{Int}^*\} \cup \{\square, \square'\}$, and thus not enclosed by $\{\sigma(D, D') : (D, D') \in \text{Int}(\mathcal{D}_\square \cup \mathcal{D}_{\square'})\}$ because $\sigma(D, D') \subseteq \square \cup \square'$ for all $(D, D') \in \text{Int}(\mathcal{D}_\square \cup \mathcal{D}_{\square'}) \setminus \text{Int}^*$. But this contradicts the fact that $y \in Y$. Therefore, $y \in Y_L(D, D')$ for some $(D, D') \in \text{Int}^*$ and hence $y \in Y_L(D_L, D'_L)$.

To further see that y is enclosed by $\{D_L, D'_L, D_R, D'_R\}$, observe that y is enclosed by

$$\{\sigma(D_L, D'_L), \sigma(D_R, D'_R), \square, \square'\},$$

since $y \in Q \setminus (\square \cup \square')$ and $y \in Y_L(D_L, D'_L) \cap Y_R(D_R, D'_R)$. We have $\sigma(D_L, D'_L) \subseteq D_L \cup D'_L$, $\sigma(D_R, D'_R) \subseteq D_R \cup D'_R$, $\square \subseteq D_L$, and $\square' \subseteq D'_L$. Hence, y is enclosed by $\{D_L, D'_L, D_R, D'_R\}$. \square

Our construction of \mathcal{D}' satisfies $|\mathcal{D}' \cap \mathcal{D}_\square| = O(1)$ for all $\square \in \Gamma$, simply because every grid cell in Γ is only involved in a constant number of relevant pairs. The following observation further shows that $\text{opt}(X, \mathcal{D}') = O(\text{opt}(X, \mathcal{D}))$, which completes the proof of Lemma 2.

Observation 5. $\text{opt}(X, \mathcal{D}') \leq c \cdot \text{opt}(X, \mathcal{D})$, where $c = \max_{\square \in \Gamma} |\mathcal{D}' \cap \mathcal{D}_\square|$.

Proof. Consider a minimum subset $\mathcal{D}_{\text{opt}} \subseteq \mathcal{D}$ enclosing X . We construct a subset $\mathcal{D}'_{\text{opt}} \subseteq \mathcal{D}'$ as follows. For every $\square \in \Gamma$ with $\mathcal{D}_{\text{opt}} \cap \mathcal{D}_\square \neq \emptyset$, we include in $\mathcal{D}'_{\text{opt}}$ all unit disks in $\mathcal{D}' \cap \mathcal{D}_\square$. Clearly, we have $|\mathcal{D}'_{\text{opt}}| \leq c \cdot |\mathcal{D}_{\text{opt}}|$. It suffices to show that $\mathcal{D}'_{\text{opt}}$ encloses X .

Consider a point $x \in X$. We say a cycle $\phi = (D_0, D_1, \dots, D_r)$ in $G[\mathcal{D}_{\text{opt}} \cup \mathcal{D}'_{\text{opt}}]$ is *good* if it satisfies (i) γ_ϕ is non-contractible in $\mathbb{R}^2 \setminus \{x\}$ and (ii) r is even and for every $i \in [\frac{r}{2}]$ there exists $\square_i \in \Gamma$ such that $D_{2i-1}, D_{2i} \in \mathcal{D}_{\square_i}$. The *cost* of ϕ is defined as the number of indices $i \in [r]$ satisfying $D_i \notin \mathcal{D}'_{\text{opt}}$. We claim that $G[\mathcal{D}_{\text{opt}} \cup \mathcal{D}'_{\text{opt}}]$ contains at least one good cycle. Since \mathcal{D}_{opt} encloses x , by Lemma 1, there exists a cycle $\phi = (D_0, D_1, \dots, D_r)$ in $G[\mathcal{D}_{\text{opt}}]$ (and thus in $G[\mathcal{D}_{\text{opt}} \cup \mathcal{D}'_{\text{opt}}]$) such that γ_ϕ is non-contractible in $\mathbb{R}^2 \setminus \{x\}$. This cycle satisfies condition (i) for good cycles but not condition (ii). Now define $\psi = (D_0, D_0, D_1, D_1, \dots, D_r, D_r)$, which is a cycle in $G[\mathcal{D}_{\text{opt}} \cup \mathcal{D}'_{\text{opt}}]$ satisfying condition (ii). Note that ψ also satisfies condition (i) since $\gamma_\psi = \gamma_\phi$. Thus, ψ is good and $G[\mathcal{D}_{\text{opt}} \cup \mathcal{D}'_{\text{opt}}]$ contains at least one good cycle.

To prove that $\mathcal{D}'_{\text{opt}}$ encloses x , let $\phi = (D_0, D_1, \dots, D_r)$ be a good cycle in $G[\mathcal{D}_{\text{opt}} \cup \mathcal{D}'_{\text{opt}}]$ with minimum cost. By condition (i), γ_ϕ is non-contractible in $\mathbb{R}^2 \setminus \{x\}$. If $D_1, \dots, D_r \in \mathcal{D}'_{\text{opt}}$, then we are done. Indeed, in this case, ϕ is also a cycle in $G[\mathcal{D}'_{\text{opt}}]$ and Lemma 1 implies that $\mathcal{D}'_{\text{opt}}$ encloses x . So suppose $D_j \notin \mathcal{D}'_{\text{opt}}$ for some $j \in [r]$. We assume $j = 2i - 1$ for some $i \in [\frac{r}{2}]$; the case where j is even can be handled in the same way. Since ϕ is good, by condition (ii), there exists $\square_i \in \Gamma$ (resp., $\square_{i-1} \in \Gamma$) such that $D_j, D_{j+1} \in \mathcal{D}_{\square_i}$ (resp., $D_{j-2}, D_{j-1} \in \mathcal{D}_{\square_{i-1}}$). Note that $(\square_{i-1}, \square_i)$ is a relevant pair, because $D_{j-1} \cap D_j \neq \emptyset$. Recall that when considering $(\square_{i-1}, \square_i)$, we included in \mathcal{D}' (at most) four unit disks in $\mathcal{D}_{\square_{i-1}} \cup \mathcal{D}_{\square_i}$, and among them there exist two intersecting unit disks $D'_{j-1} \in \mathcal{D}_{\square_{i-1}}$ and $D'_j \in \mathcal{D}_{\square_i}$. We have $D'_{j-1}, D'_j \in \mathcal{D}'_{\text{opt}}$, since $\mathcal{D}_{\text{opt}} \cap \mathcal{D}_{\square_{i-1}} \neq \emptyset$ and $\mathcal{D}_{\text{opt}} \cap \mathcal{D}_{\square_i} \neq \emptyset$. As $D_{j-2}, D'_{j-1} \in \mathcal{D}_{\square_{i-1}}$ and $D'_j, D_{j+1} \in \mathcal{D}_{\square_i}$, we have $D_{j-2} \cap D'_{j-1} \neq \emptyset$ and $D'_j \cap D_{j+1} \neq \emptyset$. Therefore, $\phi' = (D_0, D_1, \dots, D_{j-2}, D'_{j-1}, D'_j, D_{j+1}, \dots, D_r)$ is also a cycle in $G[\mathcal{D}_{\text{opt}} \cup \mathcal{D}'_{\text{opt}}]$. Observe that ϕ' cannot be a good cycle. Indeed, if ϕ' is good, the cost of ϕ' is strictly smaller than the cost of ϕ as $D'_{j-1}, D'_j \in \mathcal{D}'_{\text{opt}}$ but $D_j \notin \mathcal{D}'_{\text{opt}}$, which contradicts the fact that ϕ is a good cycle with minimum cost. However, ϕ' satisfies condition (ii) for good cycles, because $D_{j-2}, D'_{j-1} \in \mathcal{D} \cap \mathcal{D}_{\square_{i-1}}$ and $D'_j, D_{j+1} \in \mathcal{D}_{\square_i}$. Thus, ϕ' does not satisfy condition (i), i.e., $\gamma_{\phi'}$ is contractible in $\mathbb{R}^2 \setminus \{x\}$. It follows that γ_ϕ and $\gamma_{\phi'}$ are not homotopic in $\mathbb{R}^2 \setminus \{x\}$. Now consider the paths $\psi = (D_{j-2}, D_{j-1}, D_j, D_{j+1})$ and $\psi' = (D_{j-2}, D'_{j-1}, D'_j, D_{j+1})$ in $G[\mathcal{D}_{\text{opt}} \cup \mathcal{D}'_{\text{opt}}]$. Their corresponding curves γ_ψ and $\gamma_{\psi'}$ share the same endpoints. Note that if γ_ψ and $\gamma_{\psi'}$ are homotopic in $\mathbb{R}^2 \setminus \{x\}$, then γ_ϕ and $\gamma_{\phi'}$ are also homotopic in $\mathbb{R}^2 \setminus \{x\}$ (when picking the same point as their base points). Hence, γ_ψ and $\gamma_{\psi'}$ are not homotopic in $\mathbb{R}^2 \setminus \{x\}$. This further implies γ_ξ is non-contractible in $\mathbb{R}^2 \setminus \{x\}$, where $\xi = (D_{j-2}, D_{j-1}, D_j, D_{j+1}, D'_j, D'_{j-1}, D_{j-2})$ is a cycle in $G[\mathcal{D}_{\text{opt}} \cup \mathcal{D}'_{\text{opt}}]$. Since all vertices of ξ are in $\mathcal{D}_{\square_{i-1}} \cup \mathcal{D}_{\square_i}$, by Lemma 1, $\mathcal{D}_{\square_{i-1}} \cup \mathcal{D}_{\square_i}$ encloses x . Observation 4 then implies that the four unit disks we include in \mathcal{D}' for the relevant pair $(\square_{i-1}, \square_i)$ also enclose x . These unit disks are all in $\mathcal{D}'_{\text{opt}}$ as $\mathcal{D}_{\text{opt}} \cap \mathcal{D}_{\square_{i-1}} \neq \emptyset$ and $\mathcal{D}_{\text{opt}} \cap \mathcal{D}_{\square_i} \neq \emptyset$. Thus, $\mathcal{D}'_{\text{opt}}$ encloses x . \square

3.2 Approximation for the sparse case via min-cut

Thanks to Lemma 2, we can now assume $|\mathcal{D}_\square| = O(1)$ for all $\square \in \Gamma$, and design a constant-approximation algorithm under this assumption.

First, we observe that the maximum degree of the graph $G[\mathcal{D}]$ is $O(1)$. Let $D \in \mathcal{D}_\square$. For every neighbor D' of D in $G[\mathcal{D}]$, we must have $D' \in \mathcal{D}_{\square'}$ for a cell $\square' \in \Gamma$ with constant distance from \square . The number of such cells is $O(1)$ and $|\mathcal{D}_{\square'}| = O(1)$ for each such cell \square' . Thus, D has $O(1)$ neighbors and the maximum degree of $G[\mathcal{D}]$ is $O(1)$.

Consider the following drawing of the graph $G[\mathcal{D}]$ in the plane. We draw each vertex $D \in \mathcal{D}$ at the point $\text{ctr}(D)$. Then we draw each edge $(D, D') \in \text{Int}(\mathcal{D})$ as the segment $\sigma(D, D')$. This drawing is not necessarily planar because (the images of) the edges can cross. However, it has a nice and important property: each edge crosses with at most $O(1)$ other edges. Again, this follows from the sparsity of \mathcal{D} . Let $(D, D') \in \text{Int}(\mathcal{D})$ where $D \in \mathcal{D}_\square$ and $D' \in \mathcal{D}_{\square'}$. If $\sigma(D, D')$ crosses with $\sigma(E, E')$ for $(E, E') \in \text{Int}(\mathcal{D})$, then the cells containing $\text{ctr}(E)$ and $\text{ctr}(E')$ must be with constant distance from \square and \square' . There are $O(1)$ such cells each of which contains $O(1)$ centers of the unit disks in \mathcal{D} . Thus, the number of edges whose images cross with $\sigma(D, D')$ is $O(1)$.

From the drawing of $G[\mathcal{D}]$, we create a planar graph H as follows. Let C be the set of crossing points of the drawing of $G[\mathcal{D}]$. The vertex set of H is $\{\text{ctr}(D) : D \in \mathcal{D}\} \cup C$. The points in C subdivide the image of each edge of $G[\mathcal{D}]$ into $O(1)$ pieces. These pieces, which have endpoints in $\{\text{ctr}(D) : D \in \mathcal{D}\} \cup C$, are the edges of H . The drawing of $G[\mathcal{D}]$ induces a planar drawing of H . See Figure 4. We have the following important observation.

Observation 6. *There exists a subset $E \subseteq E(H)$ with $|E| = O(\text{opt}(X, \mathcal{D}))$ which encloses X .*

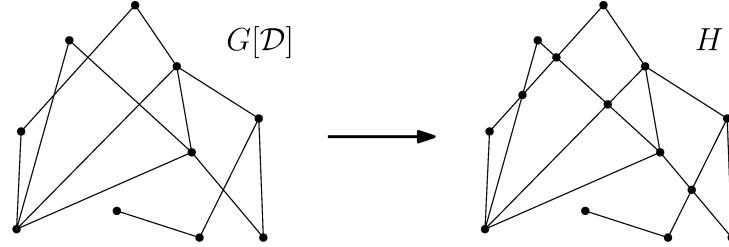


Figure 4: Illustration of the construction of H from the drawing of $G[\mathcal{D}]$.

Furthermore, given a subset $E \subseteq E(H)$ that encloses X , one can compute in polynomial time a subset $\mathcal{D}_0 \subseteq \mathcal{D}$ enclosing X such that $|\mathcal{D}_0| \leq 2|E|$.

Proof. Let $\mathcal{D}_{\text{opt}} \subseteq \mathcal{D}$ be a subset of size $\text{opt}(X, \mathcal{D})$ that encloses X . Since the maximum degree of $G[\mathcal{D}]$ is $O(1)$, $|\text{Int}(\mathcal{D}_{\text{opt}})| = O(\text{opt}(X, \mathcal{D}))$. Each edge $(D, D') \in \text{Int}(\mathcal{D}_{\text{opt}})$ corresponds to $O(1)$ edges in H (which form a subdivision of $\sigma(D, D')$). Let $E \subseteq E(H)$ consist of all edges in H correspond to the edges in $\text{Int}(\mathcal{D}_{\text{opt}})$. As $|\text{Int}(\mathcal{D}_{\text{opt}})| = O(\text{opt}(X, \mathcal{D}))$, $|E| = O(\text{opt}(X, \mathcal{D}))$. The union of (the images) of the edges in E is exactly equal to the union of $\sigma(D, D')$ for $(D, D') \in \text{Int}(\mathcal{D}_{\text{opt}})$, while the latter encloses X by Lemma 1. Thus, E encloses X .

Next, suppose we are given $E \subseteq E(H)$ that encloses X . Each $e \in E$ is a piece of an edge $f(e) \in \text{Int}(\mathcal{D})$. We simply define $\mathcal{D}_0 \subseteq \mathcal{D}$ as the subset consisting of the endpoints of $f(e)$ for all $e \in E$. Clearly, $|\mathcal{D}_0| \leq 2|E|$. The union of $\sigma(D, D')$ for $(D, D') \in \text{Int}(\mathcal{D}_0)$ contains the union of the edges in E , and hence encloses X . By Lemma 1, \mathcal{D}_0 encloses X . \square

It now suffices to compute a minimum-size $E_{\text{opt}} \subseteq E(H)$ that encloses X . The above observation then implies that $|E_{\text{opt}}| = O(\text{opt}(X, \mathcal{D}))$ and one can compute in polynomial time a subset $\mathcal{D}_0 \subseteq \mathcal{D}$ enclosing X such that $|\mathcal{D}_0| \leq 2|E_{\text{opt}}| = O(\text{opt}(X, \mathcal{D}))$. We show that computing E_{opt} can be reduced to the following *minimum S-T cut* problem.

Minimum S-T Cut

Input: A graph G and two disjoint sets $S, T \subseteq V(G)$.

Output: A minimum subset $E \subseteq E(G)$ such that s and t lie in different connected components of $G - E$ for every $s \in S$ and every $t \in T$.

Before discussing the reduction, we first observe that the above minimum *S-T cut* problem is polynomial-time solvable. Indeed, we can add to the input graph G a source vertex s connecting to all vertices in S and a target vertex t connecting to all vertices in T . Then we give the original edges of G weights 1 and give the edges incident to s (resp., t) weights ∞ . Let G^+ be the resulting edge-weighted graph. Now a minimum *S-T cut* in the original G is equivalent to a minimum-weight *s-t cut* in G^+ . The latter can be computed in polynomial time by the well-known duality between min-cut and max-flow.

To reduce our problem to minimum *S-T cut*, we consider the dual graph H^* of the planar graph H . Each vertex of H^* corresponds to a face f of H , which is called the *dual vertex* of f and is denoted by f^* . Each edge of H^* corresponds to an edge e of H , which is called the *dual edge* of e and is denoted by e^* ; here e^* connects the dual vertices of the two faces of H incident to e . Let o be the outer face of H . We say a face of H is *nonempty* if it contains at least one point in X . We have the following observation.

Observation 7. A subset $E \subseteq E(H)$ encloses X iff for every nonempty face f of H , f^* and o^* lie in different connected components of $H^* - \{e^* : e \in E\}$.

Proof. To see the “if” direction, assume E does not enclose X . Let $x \in X$ be a point not enclosed by E , and f be the (nonempty) face of H containing x . As E does not enclose x , there exists a curve γ in the plane connecting x and a point y in the outer face o of H that does not intersect any edge in E . Without loss of generality, we may assume that γ does not intersect any vertex of H . Suppose the faces of H visited by γ in order are f_1, \dots, f_r (where $f_1 = f$ and $f_r = o$), and when γ enters f_{i+1} from f_i it goes across the edge e_i . Now $e_1, \dots, e_{r-1} \notin E$. Thus, there is a path in $H^* - \{e^* : e \in E\}$ from f^* to o^* , which consists of the edges e_1, \dots, e_{r-1} . So f^* and o^* lie in the same connected component of $H^* - \{e^* : e \in E\}$. To see the “only if” direction, assume that E encloses X . Consider a nonempty face f of H . To see f^* and o^* lie in different connected components of $H^* - \{e^* : e \in E\}$, let π be a path from f^* to o^* in H^* . Our goal is to show that at least one edge on π is in $\{e^* : e \in E\}$. Suppose the edges on π are e_1^*, \dots, e_r^* . Then one can connect f and o by a curve γ in the plane that does not intersect any edge of H except e_1, \dots, e_r . Since E encloses X and f contains at least one point in X , γ must intersect at least one edge in E . Thus, $e_i \in E$ for some $i \in [r]$ and $e_i^* \in \{e^* : e \in E\}$. \square

By the above observation, computing a minimum-size $E_{\text{opt}} \subseteq E(H)$ enclosing X is equivalent to computing a minimum-size $E_{\text{opt}}^* \subseteq E(H^*)$ such that f^* and o^* lie in different connected components of $H^* - E_{\text{opt}}^*$ for any nonempty face f of H . Note that the latter is in turn equivalent to the minimum S - T cut instance on H^* with $S = \{o^*\}$ and $T = \{f^* : f \text{ is a nonempty face of } H\}$. Thus, E_{opt} can be computed in polynomial time, and we finally obtain our algorithm for unit disks.

Theorem 8. There exists a polynomial-time $O(1)$ -approximation algorithm for **ENCLOSING-ALL-POINTS** with unit disks, where n is the total number of points and unit disks.

4 Approach based on LP rounding

In this section, we present our algorithmic framework based on LP rounding. We present our algorithm for line segments. It generalizes to curves pairwise intersecting a constant number of times, and also general disks; see Appendix B.

Formulation of the LP relaxation. Let X be a set of points in \mathbb{R}^2 and \mathcal{S} be a set of line segments. We first formalize an LP relaxation of **ENCLOSING-ALL-POINTS**. Intuitively, the idea is as follows: a point q is enclosed by a set $\mathcal{S}' \subseteq \mathcal{S}$ of segments iff the outer boundary of these segments enclose q . So it suffices to specify the outer boundary \mathcal{B} of the union of segments that we select in the solution. \mathcal{B} is a set of disjoint simple polygons (as shown in Fig. 5), which can also be viewed as a set of cycles in a graph G : the vertices of G are intersection points between two segments in \mathcal{S} , and the edges of G are subsegments between a pair of vertices in G that lie on the same segment in \mathcal{S} . To get a fractional solution, our goal is to select a set of fractional cycles in G to enclose all points.

We now formally describe our LP. We start with specifying the variables in the LP. The arrangement of \mathcal{S} can be viewed as a planar graph G_{arr} , which has $O(|\mathcal{S}|^2)$ vertices and edges. Each vertex in G_{arr} is an intersection point of two line segments in \mathcal{S} . We then create a new (directed) graph G , which has the same set of vertices as G_{arr} . For each line segment $\ell \in \mathcal{S}$, let $V(\ell)$ denote the set of all vertices in G that lie on ℓ . For each pair of points $u, v \in V(\ell)$, we include in G a directed edge $e_{uv}^{(\ell)}$ from u to v , and also create a variable $x_{uv}^{(\ell)}$ where $0 \leq x_{uv}^{(\ell)} \leq 1$, indicating in what

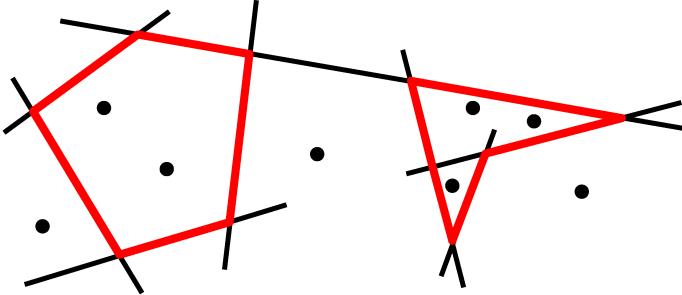


Figure 5: The outer boundary \mathcal{B} of the union of segments is a set of disjoint simple polygons (shown in red).

degree the directed subsegment uv of ℓ (with direction from u to v) is selected in the fractional solution. We call $x_{uv}^{(\ell)}$ the weight of the directed subsegment uv .

We emphasize that we need to create two variables (and respectively, two directed edges in G) for each undirected subsegment uv , because to ensure all points are enclosed according to the winding number constraints to be defined later, we need to know whether uv appears on the boundary of the solution in clockwise or counter-clockwise order.

We next present the constraints of our LP. There are two types of constraints that need to be handled: flow constraints and winding number constraints.

Flow constraints. To ensure that the solution to the LP is a set of fractional cycles, we design the flow constraints as follows. For each vertex u of G , we create the constraint

$$\sum_{v, \ell: e_{uv}^{(\ell)} \in E(G)} x_{uv}^{(\ell)} = \sum_{v, \ell: e_{vu}^{(\ell)} \in E(G)} x_{vu}^{(\ell)}.$$

This guarantees that a solution to the LP is a circulation on G . Using standard techniques, this circulation can be decomposed into $O(n^3)$ fractional cycles in polynomial time: repeatedly select the smallest variable $x_{uv}^{(\ell)}$ that is strictly positive, find a path P_{vu} from v to u in G using only edges with strictly positive values (such path always exists, and the value of each edge on the path is at least $x_{uv}^{(\ell)}$), and subtract from the solution the fractionally weighted cycle C formed by concatenating $e_{uv}^{(\ell)}$ and P_{vu} . The fractional cycle C has weight $w_C = x_{uv}^{(\ell)}$. Each round sets the value of at least one variable $x_{uv}^{(\ell)}$ to 0, so the process will terminate in $O(n^3)$ steps. Each cycle in G corresponds to a polygon (which is not necessarily simple) in the plane, so the LP solution can be decomposed into $O(n^3)$ fractionally weighted polygons.

Winding number constraints. To ensure that each point in X is enclosed by the outer boundary of the union of the chosen subsegments, the idea is to constrain their winding number. After orienting each (unknown) simple polygon on the outer boundary of the optimal solution \mathcal{S}^* in counter-clockwise order, the winding number of each point $q \in X$ is exactly 1. So it suffices to ensure a similar inequality when we form the constraints for the fractional solution to the LP.

For each point $q \in X$, let \vec{r}_q be an arbitrary ray starting from q . Define $E_q^+ = \{(u, v) : e_{uv}^{(\ell)} \text{ crosses } \vec{r}_q \text{ in counter-clockwise order}\}$ and $E_q^- = \{(u, v) : e_{uv}^{(\ell)} \text{ crosses } \vec{r}_q \text{ in clockwise order}\}$. Inspired by the winding-number algorithm, we add the following constraint to ensure the winding

number of q with respect to the LP solution that we selected is at least 1:

$$\sum_{(u,v) \in E_q^+} x_{uv}^{(\ell)} - \sum_{(u,v) \in E_q^-} x_{uv}^{(\ell)} \geq 1, \quad \forall q \in X.$$

In other words, the (weighted) winding number of q equals to the total weight of the edges that cross the ray \vec{r}_q in counter-clockwise direction, minus the total weight of the edges that cross the ray \vec{r}_q in clockwise direction. We remark that we require the winding number of each point to be at least 1 instead of exactly 1, since in the LP solution, each polygon may not be simple, which means the winding numbers can be greater than 1.

The winding number constraints together with the flow constraints ensure that each point is enclosed by the fractional cycles in the LP solution. Namely, for each point $q \in X$, we have

$$\sum_{\text{cycle } C} w_C \cdot \text{wind}(q, C) \geq 1.$$

The objective function to be minimized is just $\sum_{(u,v) \in E(G)} x_{uv}^{(\ell)}$, the total weight of the (fractionally) selected subsegments.

Rounding the LP. After solving the LP to compute the optimal fractional solution in polynomial time, we represent the LP solution as a set \mathcal{C} of $O(n^2)$ fractionally weighted cycles in the graph G as discussed above, and use these cycles to guide our rounding procedure.

Our rounding scheme is similar to the randomized rounding used in [23] (which is commonly used for approximating set cover solutions [27]). Basically speaking, for each fractional cycle $C \in \mathcal{C}$ of weight w_C , we randomly and independently select it with probability $\min\{10w_C \log n, 1\}$. Here “selecting a cycle” means selecting all subsegments corresponding to the edges on the cycle, and if a subsegment $e_{uv}^{(\ell)}$ of an input line segment ℓ is selected, then the whole segment ℓ is included in our integral solution. Note that a line segment ℓ may be included because of multiple subsegments of it are selected during the rounding process; although it suffices to include ℓ just once, in the analysis below, we treat as if ℓ would be included multiple times. We show this will not affect the approximation factor by much, as the combinatorial complexity of the outer boundary of the objects we consider is near linear.

Our goal is to ensure all points in X have winding number at least 1 w.h.p. after the rounding process, which implies that they are enclosed by the integral solution. But a technical issue arises for cycles in \mathcal{C} that are not necessarily simple: they may wind around a point multiple times, resulting in a large positive winding number for the point. In this case, even giving a small weight to the cycle can still ensure the (weighted) winding number of that point is at least 1 in the fractional solution to the LP, but that means the cycle will only be selected with a small probability during the rounding process (selecting the cycle with a larger probability would be too costly).

To handle the above issue, the idea is to “unwind” the non-simple cycles in \mathcal{C} first before the rounding, i.e., decompose them into a set of simple cycles, and then randomly select each simple cycle independently. This ensures that each cycle contributes at most once to the winding number of any point q .

To unwind a non-simple cycle $C \in \mathcal{C}$, we repeatedly choose a self-intersection point of C and split C into two cycles at that point. Specifically, if two edges u_1v_1 and u_2v_2 of C cross each other, then we create a new vertex v at their intersection point, subdividing u_1v_1 (resp., u_2v_2) into two edges u_1v and vv_1 (resp., u_2v and vv_2). In this way, we decompose the cycle C into two cycles C_1 and C_2 by splitting at v (see Figure 6). The total length of the cycles increases by 2, since the two

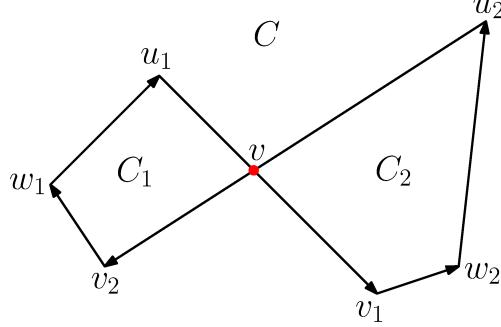


Figure 6: During the unwinding process, the non-simple cycle $C = u_1 v_1 w_1 u_2 v_2 w_1$ with length 6 is decomposed into two simple cycles $C_1 = v v_2 w_1 u_1$ and $C_2 = v v_1 w_2 u_2$ each with length 4.

crossing edges are split into four edges. Recursively subdivide the cycles C_1 and C_2 if they are not simple.

To upper bound the approximation ratio, we show that it only loses a constant approximation factor at the step of decomposing non-simple cycles into simple cycles. Let $f(n)$ denote the maximum total number of edges that a non-simple cycle of length n can decompose into, which satisfies the recurrence $f(n) \leq \max_{3 \leq n_1 \leq n-1} (f(n_1) + f(n - n_1 + 2))$, since a simple cycle formed by line segments would have length at least 3. The base case is $f(n) = n$ for $n \leq 3$. Solving the recurrence, we get $f(n) \leq 3n - 6 = O(n)$.

Let \mathcal{C}' be the set of fractional simple cycles obtained by unwinding the cycles in \mathcal{C} . It still holds that for each point $q \in X$,

$$\sum_{\text{cycle } C \in \mathcal{C}'} w_C \cdot \text{wind}(q, C) = \sum_{\text{cycle } C \in \mathcal{C}} w_C \cdot \text{wind}(q, C) \geq 1.$$

We can ignore the simple cycles in \mathcal{C}' that appear in clockwise order, because they only contribute negatively to the winding number constraints. In other words, we can assume without loss of generality that all cycles in \mathcal{C}' are counter-clockwise.

Now we analyze the probability of failure. For each $q \in X$, let $\mathcal{C}_q \subseteq \mathcal{C}'$ consist of the cycles that enclose q . We then have

$$\begin{aligned} \Pr[q \text{ is not enclosed}] &= \prod_{C \in \mathcal{C}_q} \Pr[C \text{ is not selected}] \\ &= \prod_{C \in \mathcal{C}_q} (1 - \min\{10w_C \log n, 1\}) \\ &\leq \lim_{k \rightarrow \infty} \left(1 - \frac{10 \log n}{k}\right)^k \leq e^{-10 \log n} \leq \frac{1}{n^{10}}. \end{aligned}$$

By union bound, the probability of existing any point $q \in X$ that is not enclosed is at most

$$\Pr[\exists q \in X, q \text{ is not enclosed}] \leq n \cdot \frac{1}{n^{10}} \leq \frac{1}{n^9}.$$

Therefore, the rounding process succeeds w.h.p. We remark that the approach can be derandomized, by the standard method of conditional probabilities [22].

Analyzing the approximation ratio. It is easy to see that the optimal solution \mathcal{S}^* corresponds to a feasible solution to the LP: if we set the variables corresponding to the directed subsegments on the outer boundary of the arrangement of \mathcal{S}^* to be 1, and all other variables to be 0, we will obtain a feasible solution to the LP. From the known combinatorial complexity of the outer face in the arrangement of line segments, the outer boundary of \mathcal{S}^* can be decomposed into $O(|\mathcal{S}^*| \cdot \alpha(|\mathcal{S}^*|))$ subsegments. Therefore, the optimal LP value satisfies $\text{opt} = O(|\mathcal{S}^*| \cdot \alpha(|\mathcal{S}^*|))$.

As we have shown before, after the unwinding step, the total weight of edges among all cycles only increases by a constant factor, i.e., $\sum_{C \in \mathcal{C}'} w_C \cdot \text{len}(C) = O(\text{opt})$, where $\text{len}(C)$ is the length of C . So after performing randomized rounding on the unwound fractional solution, the expected size of the resulting integral solution is at most $\sum_{C \in \mathcal{C}'} 10w_C \log n \cdot \text{len}(C) = O(\text{opt} \log n) = O(|\mathcal{S}^*| \cdot \alpha(|\mathcal{S}^*|) \log n)$. So we conclude the following.

Theorem 9. *There exists a polynomial-time $O(\alpha(n) \log n)$ -approximation algorithm for ENCLOSING-ALL-POINTS with segments, where n is the total number of points and segments.*

References

- [1] Pankaj K. Agarwal and Micha Sharir. Davenport-Schinzel sequences and their geometric applications. In *Handbook of Computational Geometry*, pages 1–47. North Holland / Elsevier, 2000. [4](#), [17](#)
- [2] Helmut Alt, Sergio Cabello, Panos Giannopoulos, and Christian Knauer. Minimum cell connection in line segment arrangements. *International Journal of Computational Geometry & Applications*, 27(03):159–176, 2017. [3](#)
- [3] Takao Asano, Tetsuo Asano, Leonidas Guibas, John Hershberger, and Hiroshi Imai. Visibility of disjoint polygons. *Algorithmica*, 1:49–63, 1986. [1](#)
- [4] Sayan Bandyapadhyay, Neeraj Kumar, Subhash Suri, and Kasturi R. Varadarajan. Improved approximation bounds for the minimum constraint removal problem. *Comput. Geom.*, 90:101650, 2020. [3](#)
- [5] Sergey Bereg and David G. Kirkpatrick. Approximating barrier resilience in wireless sensor networks. In *Algorithmic Aspects of Wireless Sensor Networks, 5th International Workshop (ALGOSENSORS)*, volume 5804 of *Lecture Notes in Computer Science*, pages 29–40. Springer, 2009. [3](#)
- [6] Sergio Cabello and Panos Giannopoulos. The complexity of separating points in the plane. *Algorithmica*, 74(2):643–663, 2016. [1](#), [3](#)
- [7] David Yu Cheng Chan and David G. Kirkpatrick. Approximating barrier resilience for arrangements of non-identical disk sensors. In *Algorithms for Sensor Systems, 8th International Symposium on Algorithms for Sensor Systems, Wireless Ad Hoc Networks and Autonomous Mobile Entities (ALGOSENSORS)*, volume 7718 of *Lecture Notes in Computer Science*, pages 42–53. Springer, 2012. [1](#), [3](#)
- [8] Eduard Eiben, Jonathan Gemmell, Iyad A. Kanj, and Andrew Youngdahl. Improved results for minimum constraint removal. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, pages 6477–6484. AAAI Press, 2018. [3](#)

[9] Lawrence H. Erickson and Steven M. LaValle. A simple, but NP-hard, motion planning problem. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*. AAAI Press, 2013. [3](#)

[10] Subir Kumar Ghosh and David M Mount. An output-sensitive algorithm for computing visibility graphs. *SIAM Journal on Computing*, 20(5):888–910, 1991. [1](#)

[11] Matt Gibson, Gaurav Kanade, Rainer Penninger, Kasturi R. Varadarajan, and Ivo Vigan. On isolating points using unit disks. *J. Comput. Geom.*, 7(1):540–557, 2016. [1, 3](#)

[12] Leonidas J. Guibas, Micha Sharir, and Shmuel Sifrony. On the general motion-planning problem with two degrees of freedom. *Discret. Comput. Geom.*, 4:491–521, 1989. [4, 17](#)

[13] John Hershberger and Subhash Suri. An optimal algorithm for euclidean shortest paths in the plane. *SIAM Journal on Computing*, 28(6):2215–2256, 1999. [1](#)

[14] Sanjiv Kapoor and SN Maheshwari. Efficient algorithms for euclidean shortest path and visibility problems with polygonal obstacles. In *Proceedings of the fourth annual symposium on computational geometry*, pages 172–182, 1988. [1](#)

[15] Klara Kedem, Ron Livne, János Pach, and Micha Sharir. On the union of jordan regions and collision-free translational motion amidst polygonal obstacles. *Discrete & Computational Geometry*, 1:59–70, 1986. [3](#)

[16] Matias Korman, Maarten Löffler, Rodrigo I Silveira, and Darren Strash. On the complexity of barrier resilience for fat regions and bounded ply. *Computational Geometry*, 72:34–51, 2018. [3](#)

[17] Neeraj Kumar, Daniel Lokshtanov, Saket Saurabh, and Subhash Suri. A constant factor approximation for navigating through connected obstacles in the plane. In *Proceedings of the 32nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 822–839. SIAM, 2021. [1, 3](#)

[18] Neeraj Kumar, Daniel Lokshtanov, Saket Saurabh, Subhash Suri, and Jie Xue. Point separation and obstacle removal by finding and hitting odd cycles. In *Proceedings of the 38th Symposium on Computational Geometry (SoCG)*, volume 224 of *LIPICS*, pages 52:1–52:14, 2022. [1, 2, 3](#)

[19] Santosh Kumar, Ten-Hwang Lai, and Anish Arora. Barrier coverage with wireless sensors. *Wirel. Networks*, 13(6):817–834, 2007. [1, 3](#)

[20] Joseph SB Mitchell. Shortest paths among obstacles in the plane. In *Proceedings of the ninth annual symposium on Computational geometry*, pages 308–317, 1993. [1](#)

[21] Ricky Pollack, Micha Sharir, and Shmuel Sifrony. Separating two simple polygons by a sequence of translations. *Discrete & Computational Geometry*, 3:123–136, 1988. [4](#)

[22] Prabhakar Raghavan. Probabilistic construction of deterministic algorithms: Approximating packing integer programs. *J. Comput. Syst. Sci.*, 37(2):130–143, 1988. [13](#)

[23] Prabhakar Raghavan and Clark D. Thompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Comb.*, 7(4):365–374, 1987. [12](#)

[24] M. Shimrat. Algorithm 112: Position of point relative to polygon. *Commun. ACM*, 5(8):434, 1962. [4](#)

- [25] Kuan-Chieh Robert Tseng. *Resilience of wireless sensor networks*. PhD thesis, University of British Columbia, 2011. [3](#)
- [26] Kuan-Chieh Robert Tseng and David G. Kirkpatrick. On barrier resilience of sensor networks. In *Algorithms for Sensor Systems - 7th International Symposium on Algorithms for Sensor Systems, Wireless Ad Hoc Networks and Autonomous Mobile Entities (ALGOSENSORS)*, volume 7111 of *Lecture Notes in Computer Science*, pages 130–144. Springer, 2011. [3](#)
- [27] Vijay V. Vazirani. *Approximation Algorithms*. Springer, 2001. [12](#)
- [28] Haitao Wang. A new algorithm for euclidean shortest paths in the plane. *Journal of the ACM*, 70(2):1–62, 2023. [1](#)

A More results based on sparsification and min-cut

Our algorithmic technique in Section 3 can be generalized to any set \mathcal{O} of obstacles that are similarly-sized fat pseudo-disks. Formally, these obstacles are required to satisfy the following conditions. First, all obstacles are convex, and there exist constants r^- and r^+ such that for every $O \in \mathcal{O}$, there exist two concentric disks D_O^- and D_O^+ with radii r^- and r^+ respectively satisfying that $D_O^- \subseteq O \subseteq D_O^+$. Second, the boundaries of any two obstacles intersect at most twice.

To generalize our algorithm, we need to choose the center $\text{ctr}(O)$ of each obstacle $O \in \mathcal{O}$ and specify a curve $\sigma(O, O')$ for every edge $\text{Int}(\mathcal{O})$. By definition, for every $O \in \mathcal{O}$, we have concentric disks D_O^- and D_O^+ with radii r^- and r^+ respectively satisfying $D_O^- \subseteq O \subseteq D_O^+$. We simply define $\text{ctr}(O)$ as the common center of D_O^- and D_O^+ . To define $\sigma(O, O')$, we choose an arbitrary point $o \in O \cap O'$. Then we let $\sigma(O, O')$ be the poly-line obtained by concatenating the segment between $\text{ctr}(O)$ and o with the segment between o and $\text{ctr}(O')$. Clearly, $\sigma(O, O') \subseteq O \cup O'$.

The arguments are almost the same as those in Section 3. The only place we need to slightly adjust is the argument for Observation 3. Consider a relevant pair (\square, \square') , and let $Y \subseteq X$ consist of the points enclosed by $\mathcal{O}_\square \cup \mathcal{O}_{\square'}$ where $\mathcal{O}_\square = \{O \in \mathcal{O} : \text{ctr}(O) \in \square\}$ and $\mathcal{O}_{\square'} = \{O \in \mathcal{O} : \text{ctr}(O) \in \square'\}$. As in Section 3, let $\text{Int}^* \subseteq \text{Int}(\mathcal{O}_\square \cup \mathcal{O}_{\square'})$ consist of the edges (O, O') where $O \in \mathcal{O}_\square$ and $O' \in \mathcal{O}_{\square'}$. Here we need a definition for $Y_L(O, O')$ and $Y_R(O, O')$ for every $(O, O') \in \text{Int}^*$ so that the property in Observation 3, which can in turn allow us to use the same argument to show the existence of four objects in $\mathcal{O}_\square \cup \mathcal{O}_{\square'}$ enclosing Y . Denote by z and z' the centers of \square and \square' , respectively. For convenience, we move $\text{ctr}(O)$ for all $O \in \mathcal{O}_\square$ (resp., $O \in \mathcal{O}_{\square'}$) to the point z (resp., z'). This is fine because we still have $\text{ctr}(O) \in \square$ for all $O \in \mathcal{O}_\square$ and $\text{ctr}(O) \in \square'$ for all $O \in \mathcal{O}_{\square'}$. Now for every $(O, O') \in \text{Int}^*$, $\sigma(O, O')$ is a two-piece poly-line connecting z and z' . Let ℓ be the line through z and z' . For $(O, O') \in \text{Int}^*$, let $\ell(O, O')$ be the poly-line obtained from ℓ by replacing the segment $\overline{zz'}$ with $\sigma(O, O')$. Then $\ell(O, O')$ partitions the plane into two parts H_L and H_R , where H_L (resp., H_R) is to the left (resp., right) of $\ell(O, O')$ with respect to the direction from z to z' . Now we can define $Y_L(O, O') = Y \cap H_L$ and $Y_R(O, O') = Y \cap H_R$.

To see the property in Observation 3 holds, consider $(O, O'), (P, P') \in \text{Int}^*$. If $\sigma(O, O')$ and $\sigma(P, P')$ do not intersect at any point other than z and z' , the same proof of Observation 3 works. So assume $\sigma(O, O')$ and $\sigma(P, P')$ intersect at a point v other than z and z' . Let $o \in O \cap O'$ (resp., $p \in P \cap P'$) be the middle vertex of $\sigma(O, O')$ (resp., $\sigma(P, P')$). Without loss of generality, assume \overline{zo} and \overline{zp} intersect at the point v . Then $v \in O$ and $v \in P'$. As in the proof of Observation 3, our goal is to show that either $Y \cap \Delta zvp = \emptyset$ or $Y \cap \Delta z'vo = \emptyset$, which implies either $Y_L(O, O') \subseteq Y_L(P, P')$ or $Y_L(P, P') \subseteq Y_L(O, O')$. If $p \in O$, then $z, v, p \in O$ and thus $\Delta zvp \subseteq O$, which implies $Y \cap \Delta zvp = \emptyset$. Also, if $z' \in O$, then $z', v, o \in O$ and thus $\Delta z'vo \subseteq O$, which implies $Y \cap \Delta z'vo = \emptyset$. So assume

$p, z' \notin O$. For the same reason, we can assume $o, z \notin P'$. But this contradicts the pseudo-disk property of O and P' . To see this, let ℓ_1 (resp., ℓ_2) be the line through o and z (resp., p and z'). The lines ℓ_1 and ℓ_2 intersect at v (which is contained in both O and P'), and partition the plane into four wedges. Since the boundaries of O and P' intersect at most twice, there exists one wedge W in which the boundaries of O and P' do not intersect and hence either $O \cap W \subseteq P' \subseteq W$ or $P' \cap W \subseteq O \subseteq W$. However, observe that every wedge contains a point in $O \setminus P'$ and a point in $P' \setminus O$, and hence cannot satisfy this condition. For example, the wedge whose boundary containing \overline{zv} and \overline{pv} contains $z \in O \setminus P'$ and $p \in P' \setminus O$. Therefore, we must have either $Y \cap \Delta zvp = \emptyset$ or $Y \cap \Delta z'vo = \emptyset$, and everything follows.

Theorem 10. *There exists a polynomial-time $O(1)$ -approximation algorithm for ENCLOSING-ALL-POINTS with similarly-sized fat pseudo-disks (in particular, unit disks and unit squares), where n is the total number of points and obstacles.*

B More results based on LP

The LP-based technique can be modified to work for more general types of curves with minor changes. For curves that pairwise intersect only $s = O(1)$ times, one technical issue is during unwinding, a simple cycle may only have length 2. If we decompose a cycle C into a cycle C_1 with length 2 and another cycle C_2 , then the length of C_2 is the same as C , therefore not getting a good recurrence for the maximum total size $f(n)$ of the simple cycles that we decompose into.

To fix this issue, we first try to unwind at a self-intersection point v of C , such that the two subcycles have length strictly less than $\text{len}(C)$. If such point exists, then our previous analysis for the recurrence of $f(n)$ still works. Otherwise if there are no such points, then it means only pairs of consecutive edges on the cycle C can intersect. In this case, there are only at most $s \cdot \text{len}(C)$ self-intersections on C , so we can decompose C into $O(s \cdot \text{len}(C))$ subcycles, increasing the total number of edges by only a factor of $O(s) = O(1)$.

For curves that pairwise intersect at most s times, it is known that the combinatorial complexity of the outer face in the arrangement is bounded by $O(\lambda_{s+2}(n))$ [12, 1] as mentioned earlier in Sec. 2. So our approximation algorithm yields the following result.

Theorem 11. *There exists a polynomial-time $O(\frac{\lambda_{s+2}(n)}{n} \log n)$ -approximation algorithm for ENCLOSING-ALL-POINTS with curves that pairwise intersect at most $s = O(1)$ times, where n is the total number of points and segments.*

In particular, the approach also applies to disks. Since it is known that the outer boundary for disks has complexity $O(n)$, we get the following result for disks.

Theorem 12. *There exists a polynomial-time $O(\log n)$ -approximation algorithm for ENCLOSING-ALL-POINTS with disks, where n is the total number of points and disks.*