ON ISOLATING POINTS USING UNIT DISKS*

Matt Gibson,[†] *Gaurav Kanade,*[‡] *Rainer Penninger,*[§] *Kasturi Varadarajan*[¶] *and Ivo Vigan*[‖]

Abstract. Given a set of points in the plane and a set of disks which separate the points, we consider the problem of selecting a minimum size subset of the disks such that any path between any pair of points is intersected by at least one of the selected disks. We present a $(9 + \epsilon)$-approximation algorithm for this problem and show that it is NP-complete even if all disks have unit radius and no disk contains any points. Using a similar reduction, we further show that the Multiterminal Cut problem [9] remains NP-complete on unit disk graphs. Lastly, we prove that removing a minimum subset of a collection of unit disks, such that the plane minus the arrangement of the remaining disks consists of a single connected region is also NP-complete.

1 Introduction

Wireless sensors are being extensively used in applications to provide barriers as a defense mechanism against intruders at important buildings, estates, national borders etc. Monitoring the area of interest by this type of coverage is called *barrier* coverage [13]. Such sensors are also being used to detect and track moving objects such as animals in national parks, enemies in a battlefield, forest fires, crop diseases etc. In such applications it might be prohibitively expensive to attain full coverage, where each point of a given region is within sensing distance of at least one sensor. It suffices to ensure that the object under consideration cannot travel too far before it is detected. Such coverage is called *trap* coverage [5, 16]. Inspired by such applications, we consider the problem of isolating a set of points by a minimum-size subset of a given set of unit radius disks. A unit disk crudely models the region sensed by a sensor, and the algorithm presented in the next section readily generalizes to disks of arbitrary, different radii.

Definition 1. *A set $\mathcal{D}$ of disks embedded in the plane,* separates *a set P of points, if for any two points $p, q \in P$, every path between p and q intersects at least one disk in $\mathcal{D}$, as shown in Figure 1.*

Problem 2 (Point Isolation)**.** *Given a set P of k points and a set $\mathcal{D}$ of n unit disks in the plane, separating P, find a minimum cardinality subset of $\mathcal{D}$ that still separates P.*

*This material is based upon work supported by the National Science Foundation under Grant No. 0915543 and Grant No. 1017539.

[†] *Dept. of Computer Science, University of Texas at San Antonio, San Antonio*, gibson@cs.utsa.edu

[‡] *Dept. of Computer Science, The University of Iowa, Iowa City*, gaurav-kanade@uiowa.edu

[§] *Dept. of Computer Science I, University of Bonn, Bonn, Germany*, penninge@cs.uni-bonn.de

[¶] *Dept. of Computer Science, The University of Iowa, Iowa City*, kvaradar@iowa.uiowa.edu

[‖] *Dept. of Computer Science, The City University of New York, The Graduate Center, New York*, ivigan@gc.cuny.edu

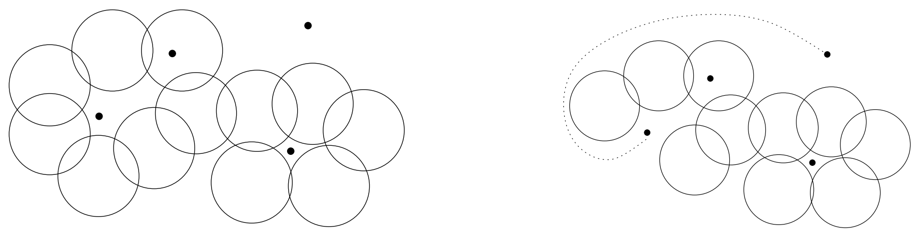

(a) (b)

Figure 1: (a) A set of disks separating four points, since every path connecting any two points intersects a disk. (b) A set of disks that does not separate the points.

For several variants of the geometric set cover problem, approximation algorithms have been designed [4, 8, 14] that improve upon the best guarantees for the combinatorial set cover problem. For the problem of covering points by the smallest subset of a given set of unit disks, there exists a $(9+\epsilon)$-approximation algorithm [1], with $0 < \epsilon \leq 6$, and a PTAS [14]. The PTAS can even be used for disks of arbitrary different radii. Our problem can be viewed as a set cover problem where the elements that need to be covered are not points, but paths. However, since the set system consists of all subsets of $\mathcal{D}$, known results only imply a trivial $O(n)$-approximation when viewed through this set cover lens.

Contribution and Organization. In Section 2, we present a polynomial time algorithm that guarantees a $(9+\epsilon)$-approximation for the Point Isolation problem. We first cover all points contained in disks using existing set cover algorithms. We then recursively apply to the remaining points the two-point separation algorithm of [7] which solves the Point Isolation problem optimally when $k = 2$, to find the smallest subset B of $\mathcal{D}$ that separates some pair of points in P. Observe that the arrangement of the disks in B induces exactly two faces in the plane (one bounded and one unbounded) and it thus partitions P into P_1 and P_2, such that all points in P_1 are separated from the points in P_2. The algorithm then recursively finds a separator for P_1 and for P_2, and returns the union of these separators and B. In Section 3, we prove that the Point Isolation problem is NP-complete, even if no disk contains any points. We believe that our hardness construction can be used to show hardness for a variety of unit disk problems and we show two of them: in Section 4, we show that the NP-complete Multiterminal Cut problem [9] (defined below) remains NP-complete on unit disk graphs. Lastly, in Section 5, we show that removing a minimum subset of a given collection of unit disks, such that the plane minus the arrangement of the remaining disks consists of a single connected region is also NP-complete.

Problem 3 (Multiterminal Cut [9]). *Given a graph $G = (V, E)$ and set $S \subseteq V$ of k terminals, find a minimum cardinality set $E' \subseteq E$ of edges such that in $G' = (V, E \setminus E')$ there is no path between any two terminals in S.*

Related Work. Sankararaman et al. [16] investigate a notion of coverage which they call *weak coverage*. Given a region $\mathcal{R}$ of interest (which they take to be a square in the plane) and a set $\mathcal{D}$ of unit disks (sensors), the region is said to be k-weakly covered if each connected component of $\mathcal{R} - \bigcup_{D\in\mathcal{D}} D$ has diameter at most k. They consider the situation when a given set $\mathcal{D}$ of unit disks *completely* covers $\mathcal{R}$, and address the problem of partitioning $\mathcal{D}$ into as many subsets as possible so that $\mathcal{R}$ is k-weakly covered by every subset.

Given two regions in the complement of the sensor arrangement, the *resilience* [13] with respect to these two regions is the minimum number of sensors that need to be deactivated so that there is a path between the two regions, not intersecting any sensor region. In [6], a 3-approximation algorithm for computing the resilience for unit disk sensors is presented; the computational complexity of this problem remains open. In [12], it is shown that computing the resilience for certain types of fat sensors, such as axis aligned rectangles is NP-hard. In [3], it is shown that computing the resilience of a set of line segment sensors is NP-hard. This was extended in [20] to hold even for unit length line segments. The reductions of both [12] of [20] have some resemblance to our NP-hardness reduction, since they both reduce optimization problems on graphs to the resilience problem in a sensor network modeling the graph instances. In the context of separation, [3] addresses the two-point separation problem for a set of line segments. They show that the problem in fact admits a polynomial-time exact algorithm. This work was later extended in [7] to include an exact $O(n^3)$ time algorithm for solving the two-point separation problem on unit disks. Furthermore, they present a hardness result similar to ours, which holds for unit circles. They reduce a version of PLANAR-3-SAT to the circle separation problem using variable gadgets consisting of intersecting circles. But since they place points into the faces of these intersecting circles, their reduction cannot be applied to unit disks. This holds since any two points contained in a disk are by definition separated and disks thus do not provide enough degrees freedom for the variable gadgets to work.

2 Approximation Algorithm

We will refer to the standard notions of vertices, arcs, and faces in arrangements of disks [2]. In particular, for a set $\mathcal{D}$ of n disks, we are interested in the faces in the complement of the union of the disks in $\mathcal{D}$. These are the connected regions in the arrangement of $\mathbb{R}^2 \setminus \bigcup_{D\in\mathcal{D}} D$, which, by slightly abusing notation, we will write as $\mathbb{R}^2 \setminus \mathcal{D}$.

Definition 4. *For any disk in an arrangement of disks $\mathcal{D}$, we refer to each maximal connected subset of its boundary which is not contained in any other disk as a* boundary arc*; see Figure 2. We denote the collection of all boundary arcs of all disks in $\mathcal{D}$ by $\mathcal{B}(\mathcal{D})$.*

Theorem 5 (Theorem 3.1 in [11], (see also [15]))**.** *If $|\mathcal{D}| \geq 3$, it holds that $|\mathcal{B}(\mathcal{D})| \leq 6|\mathcal{D}| - 12$.*

Observation 6. *Let $\mathcal{D}$ be a set of at least 3 disks in the plane, and Q a set of points so that (a) no point from Q is contained in any disk from $\mathcal{D}$, and (b) no face in $\mathbb{R}^2 \setminus \mathcal{D}$ contains more than one point of Q. It then holds that $|Q| \leq 2|\mathcal{D}| - 4$.*

Proof. Since each connected region in $\mathbb{R}^2 \setminus \mathcal{D}$ is bounded by at least three boundary arcs and no boundary arc appears in two regions, it follows from Theorem 5 that $|Q| \leq 2|\mathcal{D}| - 4$. □

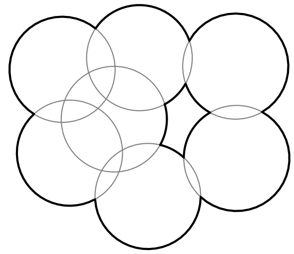

Figure 2: Illustration of the boundary arcs (in black) of an arrangement of disks (in gray).

Covering vs. Separating. The input to our problem is a set $\mathcal{D}$ of n unit disks, and a set P of k points such that $\mathcal{D}$ separates P. Let $P_c \subseteq P$ denote those points contained in some disk of $\mathcal{D}$ and let P_s denote the remaining points. Note that it follows from our definition of separation that any point contained in a disk is separated from all other points. Computing a minimum set cover of $\mathcal{D}$ for P_c is thus equivalent to solving the Point Isolation problem for the points in P_c. We first compute an α-approximation for the smallest subset of $\mathcal{D}$ that covers P_c. We then compute a β-approximation for the smallest subset of $\mathcal{D}$ that separates P_s. We claim that the combination of the two solutions is an $(\alpha + \beta)$-approximation to the Point Isolation problem. To see this, let $OPT \subseteq \mathcal{D}$ denote an optimal subset that separates P. Suppose that OPT covers k_1 of the points in P_c and let $k_2 = |P_c| - k_1$. By Observation 6, it holds that $k_2 \leq 2|OPT|$. Picking one disk to cover each of the k_2 points of P_c not covered by OPT, we see that there exists a cover for P_c of size at most $|OPT| + k_2 \leq 3|OPT|$. Thus, an α-approximate solution to this set cover problem has size at most $3\alpha|OPT|$. Since OPT also separates P_s, a β-approximation for P_s uses at most $\beta|OPT|$ disks. Since the α-approximate set cover algorithm separates each point in P_c from all the points in P and the β-approximate algorithm separates each point in P_s from all the points in P, the two algorithm combined provide a solution to the Point Isolation problem that has size $(3\alpha + \beta)|OPT|$.

Using the PTAS of [14] for the points in P_c results in $\alpha = 1 + \epsilon$, with $\epsilon > 0$; its running time is $O(n^{2(\frac{8\sqrt{2}}{\epsilon})^2+1}|P_c|)$, for $0 < \epsilon \leq 2$, (see [1]). In the rest of this section, we assume that no point of P is contained in any disk of $\mathcal{D}$ and we present a 6-approximation algorithm for the Point Isolation problem in this setting, i.e., we show that $\beta = 6$. Thus combining the two algorithms yields a $(9 + \epsilon)$-approximation algorithm.

Algorithm For a set Q of points separated by a set $\mathcal{D}$ of unit disks, such that no point of Q is contained in any disk of $\mathcal{D}$, we will show that the following algorithm yields a 6-approximation for the Point Isolation problem in this setting.

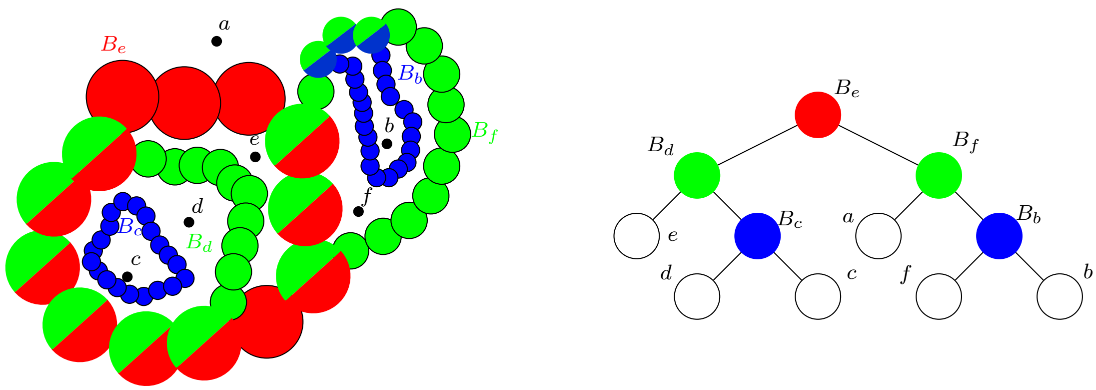

(a) (b)

Figure 3: Illustration of the charging scheme. (a) A separating arrangement computed by **recSep**, where the colors of a disk encode the separators it is contained in. Disks of two colors appear in two separators (For example, the two-colored disks in the lower left appear in both B_e and B_d). (b) The execution tree of **recSep** and the corresponding charging of the separators to the leaves.

recSep$(Q,\ \mathcal{D})$

1. If $|Q| \leq 1$, return $\emptyset$.
2. For every pair of points $s, t \in Q$, invoke the algorithm of [7] to find a minimum cardinality subset $B_{s,t} \subseteq \mathcal{D}$ such that $B_{s,t}$ separates s and t.
3. Let B denote a minimum size separator $B_{s,t}$ over all pairs $s, t \in Q$.
4. Let Q_1 and Q_2 be the partition of Q into two subsets such that each subset corresponds to points in the same face induced by the arrangement of disks in B.

   return $B \cup$ **recSep**$(Q_1,\ \mathcal{D})$ $\cup$ **recSep**$(Q_2,\ \mathcal{D})$

Since in any recursive call, both Q_1 and Q_2 contain fewer points than Q, the **recSep**$(P,\ \mathcal{D})$ algorithm indeed yields a separator for P. Computing a single separator $B_{s,t}$ takes time $O(n^3)$ (see [7]), and thus computing B takes time $O(k^2n^3)$. Since there are at most k separation steps, the total running time of **recSep**$(P,\ \mathcal{D})$ is $O(k^3n^3)$.

In each recursive step, Q is partitioned into two sets Q_1 and Q_2, thus the execution-tree of **recSep** is a binary tree where the leaves correspond to the points in Q and the internal nodes correspond to the computed separators. Each such separator separates all points in its left subtree from the points in its right subtree. It is easy to see that a tree, where each internal node has at least 2 children, has more leaves than internal nodes. After the execution of the algorithm, we charge each separator B_p to an arbitrary point $p \in P$ appearing as a leaf in the subtree rooted at B_p, such that each point in P gets charged at most once (see Figure 3). In the optimal separating arrangement OPT, we let F_p be the disks contributing to the boundary of the face containing only $p \in P$. For any $p \in P$, it holds that $|B_p| \leq |F_p|$, since otherwise the smallest separator in the corresponding recursive call would be F_p, not B_p, contradicting optimality of B_p. Letting $\mathcal{B}(OPT)$ denote the set of boundary arcs of the disks from OPT, and $P_B \subseteq P$ be the set of points to which the separators were charged, we bound the number of disks returned by **recSep** using the following inequalities

$$|\textbf{recSep}(P,\ \mathcal{D})| \leq \sum_{p \in P_B} |B_p| \leq \sum_{p \in P} |F_p| \leq |\mathcal{B}(OPT)| \leq 6|OPT|.$$

The first inequality holds because each disk used by the algorithm appears in at least one separator, the second inequality was argued above. To see the third inequality, observe that if a disk of OPT appears in $F_{p_1}, \ldots, F_{p_i}$, then its boundary appears in at least i different regions induced by the disk arrangement of OPT and thus by definition, this disk contributes at least i boundary arcs to $\mathcal{B}(OPT)$. The last inequality holds, by Theorem 5.

This proves the following theorem.

Theorem 7. *The Point Isolation problem can be approximated within a factor of* $9 + \epsilon$ *in general and within a factor of* 6 *if no point is contained in a disk of* $\mathcal{D}$.

3 NP-Completeness of the Point Isolation Problem

We show NP-completeness of the (decision version of the) Point Isolation problem, by reducing the Planar Subdivision problem (which we introduce in Problem 10) to it. This problem will be shown to be NP-complete in Theorem 11, by a reduction from the Planar Multiterminal Cut problem which is NP-hard according to Theorem 9.

Problem 8 (Planar Multiterminal Cut [9])**.** *Given a simple planar graph* $G = (V, E)$ *and a set* $S \subseteq V$ *of* k *terminals, find a minimum cardinality set* $E' \subseteq E$ *such that in* $G' = (V, E \setminus E')$*, there is no path between any two nodes in* S.

Theorem 9 ([9])**.** *If* k *is not fixed, the Planar Multiterminal Cut problem is* NP*-complete.*

Problem 10 (Planar Subdivision)**.** *Given a simple planar graph* $G = (V, E)$*, its planar embedding and a set* S *of* k *faces of* G*, find a minimum cardinality set* $E' \subseteq E$ *such that in the planar embedding of* $G' = (V, E')$ *- where the vertices of* G' *are embedded in the same way as those of* G *- any path between any two points in two faces of* S *gets intersected by at least one embedded edge of* E'.

Theorem 11. *The Planar Subdivision problem is* NP*-complete, even on connected graphs.*

Proof. Given an instance $I_1 = (G_1, S_1)$ of the Planar Multiterminal Cut problem, with $G_1 = (V_1, E_1)$, we embed G_1 in the plane (using for example the linear time algorithm of [17]) and we build an instance $I_2 = (G_2, S_2)$ of the Planar Subdivision problem with $G_2 = (V_2, E_2)$ as follows: We let $\overline{G}_2 = (\overline{V}_2, \overline{E}_2)$ be the geometric dual multigraph of the embedded graph G_1. The geometric dual is constructed by placing a vertex in each face of G_1 and adding an edge between two such vertices if the corresponding faces in G_1 had an edge in common. Then we create a simple graph G_2 from $\overline{G}_2$ by subdividing each edge $\{u,v\} \in \overline{E}_2$ into $\{u,x\}$ and $\{x,v\}$ by adding a new vertex x to V_2. We embed G_2 in the plane and let S_2 be the set of faces of G_2 whose dual vertices are in S_1. Observe that an optimal solution $OPT \subseteq E_2$ for I_2 uses an edge $\{u,x\}$ if and only if it uses edge $\{x,v\}$, with $\{u,v\} \in \overline{E}_2$, since only using one of the two subdivision edges does not change the partition of the plane, with respect to the embedding of G_2. Let $OPT_\cup \subseteq \overline{E}_2$ denote the edges of $\overline{G}_2$ obtained from OPT by merging any subdivision edges $\{u,x\}$, $\{x,v\} \in OPT$ into $\{u,v\}$, thus $|OPT| = 2|OPT_\cup|$. We denote by $OPT_\cup^* \subseteq E_1$ the duals of the edges in $OPT_\cup$. We claim that OPT is an optimal solution for I_2 if and only if $OPT_\cup^*$ is an optimal solution for I_1. To see this, let E_1' be an arbitrary subset of E_1, let $\overline{E}_2' \subseteq \overline{E}_2$ be the dual edges of E_1' and let $E_2' \subseteq E_2$ be the subdivision edges in G_2 corresponding to $\overline{E}_2'$. Two vertices $u,v \in S_1$ are connected by a path $(u,v_1,\dots,v_l,v)$ in $G_1' = (V_1, E_1 \setminus E_1')$ if and only if there is a sequence $u^*, v_1^*, \dots, v_l^*, v^*$ of adjacent faces in G_2, which, in $G_2' = (V_2, E_2')$, are merged into one face. Therefore, I_2 has a solution of size $2M$ if and only if I_1 has a solution of size M; the factor 2 stemming from subdividing each edge of $\overline{G}_2$. Since the Planar Multiterminal Cut problem is NP-complete on connected graphs and the dual of a connected graph is connected, the Planar Subdivision problem is NP-complete, even on connected graphs. Containment in NP follows, since we can transform a potential solution for I_2 back to I_1 (by constructing its dual graph) and test in polynomial time if the terminal nodes are pairwise separated. □

Corollary 12. *Solving the Point Isolation problem for line segments instead of disks is* NP-*complete.*

The reduction now takes an instance $I_2 = (G_2, S_2)$ of the Planar Subdivision problem, with G_2 being a connected embedded graph and transforms it in polynomial time to an instance $I_1 = (\mathcal{D}, S_1)$ of the Point Isolation problem. We do this by first transforming the embedding of G_2 to an equivalent straight line embedding on an $n \times n$ integer grid. We will use Lemmas 13, 14 and 15 below to argue that there is enough space in the construction to replace each edge by a path of unit disks such that no two such paths of different edges intersect.

In [17], a linear time algorithm is presented which constructs a drawing of a planar graph on n vertices, crossing free using straight line segments and having its vertices lie on an $n \times n$ grid[1]; we call such an embedding a *grid embedding*.

Lemma 13. *For any Planar Subdivision instance (G,S), with $G = (V,E)$ being a connected graph on n vertices, there exists a grid embedding of G, such that every solution in the original embedding is a solution in the grid embedding and vice versa.*

Proof. We say that two embeddings of a graph in the plane are *equivalent*, if they are homeomorphic to each other, i.e., one can be continuously transformed into the other. It is well known that every maximal plane graph on at least four vertices is three-connected (Corollary 4.4.7 of [10]) and that every three-connected graph has a unique embedding according to Whitney's Theorem (modulo the choice of the outer face). Given an embedding

[1]In [17] it is shown that this is even possible on an $(n-2) \times (n-2)$ grid, but for us an $n \times n$ grid suffices

Σ of G, we make G maximally plane by adding $3|V| - 6 - |E|$ edges to G, obtaining a new graph G' having a corresponding embedding Σ'. Embedding G' on an $n \times n$ grid and removing the $3|V| - 6 - |E|$ additional edges thus results in a grid embedding $\overline{\Sigma}$ of G which is equivalent to Σ. It thus holds that every solution for a Planar Subdivision instance in the embedding Σ of G is a solution in the grid embedding $\overline{\Sigma}$ of G and vice versa. □

We now present two lemmas which will be useful for arguing that in a grid embedding, each edge can be replaced by a path of disks such that no disks of different paths intersect.

Lemma 14. *In an $n \times n$ grid, the minimum distance between any line l through two grid points and any grid point not on l is $(2n^2 - 2n + 1)^{-\frac{1}{2}}$.*

Proof. W.l.o.g., we fix one point on l to $(0, 0)$. Denoting the second point on l by (a, b), we get a line equation of $bx - ay = 0$. Thus, the distance from a point (c, d) to l is $\frac{|bc-ad|}{\sqrt{b^2+a^2}}$. Furthermore, we can assume that $gcd(a, b) = 1$ since otherwise we can divide both coordinates by $gcd(a, b)$. Thus, setting $a = n$ and $b = n - 1$ maximizes $a^2 + b^2$, given $gcd(a, b) = 1$ and the minimum non-zero distance is thus at least $(2n^2 - 2n + 1)^{-\frac{1}{2}}$. □

Lemma 15. *In an $n \times n$ grid, for any grid point p, the minimum angle between any two distinct lines, each going through p and at least one other grid point respectively, is larger than $2 \arctan 1/(6n^2)$.*

Proof. Let g and h denote two lines through p and (a, b) and (c, d) respectively, having minimum angle, and let the slope of g be larger than the slope of h. First, observe that the slopes of g and h have the same sign, since otherwise their angle is not minimal. Therefore, it is easy to see that the minimum angle between g and h is obtained when $p = (0, 0)$. Due to symmetry we can further restrict g and h to be contained in the lower right triangle portion $\{(i, j) \mid 0 < j \le i \le n\}$ of the grid. Since $b/a > d/c$, due to monotonicity of arctan, it holds that $\arctan b/a - \arctan d/c = \arctan \frac{b/a - d/c}{1+bd/(ac)} = \arctan \frac{bc-ad}{ac+bd} \ge \arctan \frac{1}{2n^2}$. The last inequality holds since all coordinates are integers. Thus, $bc - ad \ge 1$ and $ac + bd \le 2n^2$. The Lemma then follows from the fact that $\arctan(x) > 2 \arctan(x/3)$ holds for all $0 < x < \sqrt{3}$. □

In the remainder of this section, we are going to prove the following theorem.

Theorem 16. *The Point Isolation problem is* NP*-complete, even if no point is contained in a disk.*

In order to prove Theorem 16, we reduce an instance $I_2 = (G_2, S_2)$ of the Planar Subdivision problem, with G_2 being a connected embedded graph in polynomial time to an instance $I_1 = (\mathcal{D}, S_1)$ of the Point Isolation problem; note that solving I_2 is NP-complete according to Theorem 11. We do this by first transforming the embedding of G_2 to an equivalent straight line embedding on an $n \times n$ integer grid as argued in Lemma 13.

We then replace each edge in the embedding by an edge gadget of Definition 17. Such an edge gadget is depicted in Figure 5(a) and consists of a path of disks constructed in such

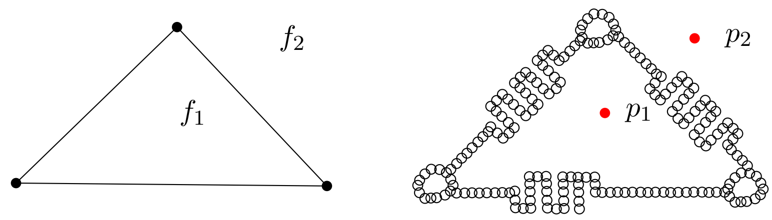

Figure 4: (a) An instance $I_2 = (G_2, S_2)$ of the Planar Subdivision problem, with $S_2 = \{f_1, f_2\}$. (b) The corresponding instance $I_1 = (\mathcal{D}, S_1)$ of the Point Isolation problem using the Vertex- and Edge gadgets defined in this section, with $S_1 = \{p_1, p_2\}$.

a way that every edge gadget contains the same amount of disks, regardless of the length of the corresponding embedded edge. Furthermore, the dimensions of each edge gadget are carefully chosen so that no two disks of different edge gadgets intersect. Having replaced each edge by an edge gadget, we replace each vertex by a vertex gadget of Definition 18, shown in Figure 5(b). A vertex gadget for a vertex v consists of a cycle of disks which are circularly arranged around v. An edge gadget for an edge $\{u, v\}$ will have non-empty intersection with the vertex gadgets for both u and v, but does not intersect any other vertex or edge gadgets. Furthermore, all vertex gadgets are pairwise disjoint. If we denote the collection of all disks contained in the vertex- and edge gadgets by $\mathcal{D}$, then each face in the embedding of G_2 has a corresponding connected region in the arrangement $\mathbb{R}^2 \setminus \bigcup \mathcal{D}$. For each face $s_2 \in S_2$ in the embedding of G_2, we place a point s_1 into this corresponding region and add s_1 to S_1. We thus obtain an instance $I_1 = (\mathcal{D}, S_1)$ of the Point Isolation problem. The main task of the reduction is to choose the radius of the disks and the dimension of the gadgets such that every edge gadget consists of the same amount of disks and all edge gadgets are disjoint. Thus, on I_1, removing any disk D from an edge gadget merges the two adjacent regions in $\mathbb{R}^2 \setminus (\bigcup \mathcal{D} \setminus \{D\})$. Note that this is not true for vertex gadgets, i.e., removing disks from vertex gadgets does not necessarily merge any regions containing points. In order to infer the solution size for I_2 from the solution size of I_1, i.e., retrieve the number of edges removed from G_2 from the number of disks removed in the solution for I_1, we have to choose the number of disks contained in a single edge gadget to be larger than the number of disks contained in all the vertex gadgets together. We will make this argument formal in Lemma 19.

Definition 17. *An edge gadget (see Figure 5(a)) for a grid-embedded edge $e = \{u, v\}$ of length $2s + 2a + b$ is a path of disks which can be thought of as being placed in an elongated octagon of height h and length $2a + b$. Every edge gadget consists of a path of C_E many radius r disks which are arranged as a straight path at the beginning and the end of the gadget, which we call* hallways*, and as an up-down path in the middle part which we call a* cabin*. While C_E, a, h and s are constant for all edge gadgets, b may vary from $1-(2s+2a)$ to $\sqrt{2}n - (2s + 2a)$, depending on the length of the embedded edge e.*

The constants a, h and s will be fixed later in this section in such a way that no two

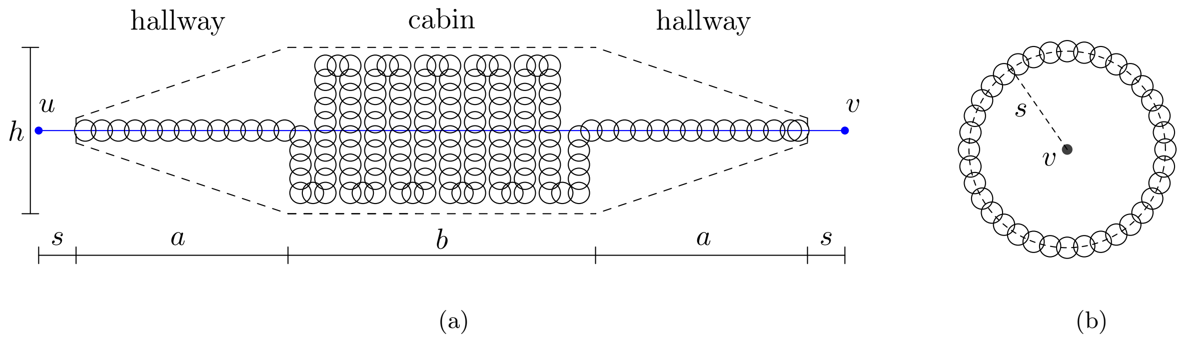

Figure 5: (a) An illustration of an edge gadget replacing the edge $\{u, v\}$, (b) an illustration of a vertex gadget.

disks in two different edge gadgets intersect and that at least one disk of an edge gadget intersects the incident vertex gadgets respectively. We define $C_E = \lceil \frac{\sqrt{2}n-2s}{2r} \rceil$, since this amounts to the number of disks of radius r needed to represent the longest edge in an $n \times n$ grid embedding as a straight line path of touching disks. Note that we take r as the unit measure for the disks. Since each edge gadget needs to contain an equal amount of disks, we have to place $C_E - a/r$ disks into the cabin of any edge gadget. Arranging the disks in cabins (for edges of length $< \sqrt{2}n$) as an up-down path as described in Definition 17 (and shown in Figure 5(a)) allows us to put a path consisting of up to $\lfloor \frac{h}{2r} \rfloor \cdot \lfloor \frac{b}{2r} - 1 \rfloor$ many disks into the cabin.

Definition 18. *A vertex gadget, shown in Figure 5(b), for an embedded vertex v consists of a cycle of $C_V = \lceil \pi s/r \rceil$ many intersecting disks of radius r which are centered on a circle of radius s centered at v*

Since the first and last disks of an edge gadget for an edge $e = \{u, v\}$ have a point at distance s from u and v respectively, it is easy to see that if $C_V \geq 4$, these disks have a non-empty intersection with the disks of the vertex gadgets for u and v respectively.

In the following lemma we show how to retrieve a solution for an instance I_2 of the Planar Subdivision problem from the solution for I_1, where I_1 is built using the construction described above.

Lemma 19. *An instance of $I_2 = (G_2, S_2)$ of the Planar Subdivision problem has a solution of size at most k_2 if and only if I_1 of the Point Isolation problem has a solution of size at most $C_E(k_2 + 1) - 1$, where I_1 is built out of I_2 using the construction described above.*

Proof. Removing any disk D from an edge gadget merges the two adjacent regions in $\mathbb{R}^2 \setminus (\bigcup \mathcal{D} \setminus \{D\})$. Thus, I_2 has a solution consisting of the corresponding k_2 edges, if and only if I_1 uses all the disks of k_2 edge gadgets. On the other hand, removing a disk from a vertex gadget does not necessarily merge two regions. Thus, a solution for I_1 may contain no disks or all disks of any vertex gadget. Since according to constraint 7, it holds that $nC_V < C_E$, it follows that if I_2 has a solution of size at most k_2, then I_1 has a solution of size at most $C_E k_2 + nC_V \leq C_E(k_2+1) - 1$. On the other hand, if I_1 has a solution consisting of at most $k_1 = C_E(k_2+1) - 1$ disks, if follows that I_2 uses at most $k_1/C_E < k_2 + 1$ edges and thus the lemma follows. □

Since the two-point separation algorithm of [7] can be used to test whether all points of P are separated in a potential solution $\mathcal{D}' \subseteq \mathcal{D}$ of the Point Isolation problem, it follows that its decision version is indeed contained in NP.

In order to finish the proof of Theorem 16, we need to show that the reduction is geometrically feasible. We thus choose the radius r of the disks and the height h of the edge gadgets such that the seven constraints below hold. These constraints ensure that no two vertex gadgets intersect (1), no two edges gadgets intersect (3, 4, 5), no edge gadget intersects any non-incident vertex gadget (2) and all edge gadgets have an equal amount of disks (6), and a single edge gadget contains more disks than all the vertex gadgets combined (7). If we set $a = 1/4$ it follows that b is at least $1/2 - 2s$ in every edge gadget. We furthermore fix s to $r(6n^2 - 1)$. Choosing the radius r to be $\frac{1}{40n^4}$ and the height h to be $\frac{1}{12n^2}$ it is easy to see that the following seven constraints are satisfied for all $n \geq 2$.

1. No two vertex gadgets intersect: $2(r+s) < 1$
2. No edge gadget intersects any vertex gadget other than the ones at its two endpoints: $r + s + h/2 < (2n^2 - 2n + 1)^{-\frac{1}{2}}$, using Lemma 14
3. No edge gadget intersects another edge gadget if the corresponding edges do not share a common vertex: $2\frac{h}{2} < (2n^2 - 2n + 1)^{-\frac{1}{2}}$
4. No disk placed in the cabin of an edge gadget intersects any disk in any incident edge gadget: $h/2 < \frac{s+a}{6n^2}$, using Lemma 15
5. No two disks contained in the hallways of two incident edge gadgets intersect each other: $s \geq \frac{r}{\sin(2\arctan(1/(6n^2))/2)} - r > r(6n^2-1)$, using Lemma 15 and the fact that the minimum distance of two disjoint lines segments occurs at an endpoint of one
6. The cabin of every edge gadget is big enough so that the whole gadget contains a path of C_E many disks: $\left\lceil \frac{\sqrt{2}n - 2s}{2r} \right\rceil - 2\left\lceil \frac{a}{2r} \right\rceil \leq \left\lfloor \frac{h}{2r} \right\rfloor \cdot \left\lfloor \frac{1-2(s+a)}{2r} - 1 \right\rfloor$
7. An edge gadget contains more disks than all the vertex gadgets combined: $C_E > nC_V$

Plugging the calculated values $r = \frac{1}{40n^4}$ and $s = r(6n^2 - 1)$ into C_E, which we earlier defined to be $\lceil \frac{\sqrt{2}n - 2s}{2r} \rceil$ yields that an edge gadget consists of $\left\lceil (20\sqrt{2}n^5 - 6n^2) + 1 \right\rceil$ disks. Analogously, plugging the calculated values into C_V, which is defined as $\lceil \pi s/r \rceil$, yields that a vertex gadget consists of $\left\lceil \pi(6n^2 - 1) \right\rceil$ many disks. We can conclude that the above construction can be computed in polynomial time.

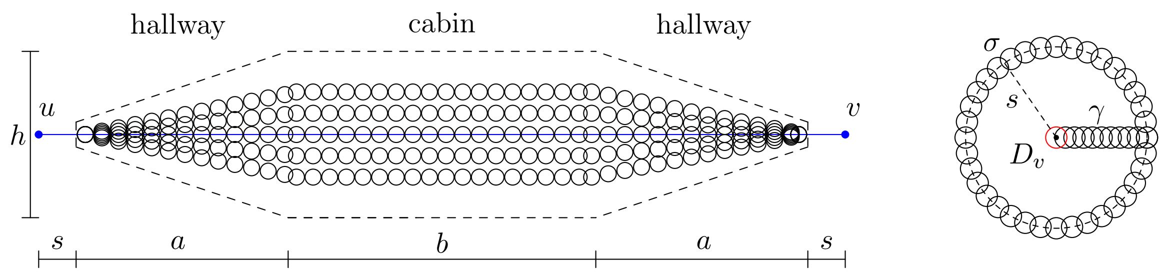

(a) (b)

Figure 6: (a) An example of an Edge Gadget for an edge of weight 5 in the proof of Theorem 20. (b) The vertex gadget for vertex v, where γ is a path of unit disks and σ is a cycle of disks. Each disk shown in γ and σ corresponds to 16 perturbed copies of it; the red disk is the centroid disk D_v.

4 Multiterminal Cut Problem on Unit Disk Graphs

In this section we are going to prove the following theorem about unit disk graphs, i.e., intersection graphs of a collection of unit disks in the plane.

Theorem 20. *The (Unweighted) Multiterminal Cut problem remains* NP*-complete on unit disk graphs, if k is not fixed.*

Proof. We make a reduction from the following weighted version of the planar Multiterminal Cut problem, which was proven to be NP-complete in [9].

Problem 21 ([9])**.** *Given an edge-weighted planar graph $G = (V, E)$, where each edge has a weight in $\{1, \dots, 5\}$ and each vertex has degree at most 3, and given set $S \subseteq V$ of k terminals, find a minimum weight set $E' \subseteq E$ of edges such that in $G' = (V, E \setminus E')$ there is no path between any two terminals in S.*

Let $I_2 = (G_2, S_2)$ be an instance of the restricted version of the planar Multiterminal Cut problem described in Problem 21. We embed $G_2 = (V_2, E_2)$ crossing free into an $n \times n$ grid, replace each edge by an edge gadget of Definition 22 (see Figure 6(a)) and each vertex by a vertex gadget of Definition 23 (see Figure 6(b)). We then construct the embedded unit disk graph $G_1 = (V_1, E_1)$ of the disks in all vertex and edge gadgets.

Definition 22 (Edge Gadget)**.** *An edge gadget for an embedded edge $e = \{u, v\}$ of weight $w \in \{1, \dots, 5\}$ consists of w paths of radius r disks, which are pairwise disjoint in the cabin of the gadget but start and end from the same point respectively, as shown in Figure 6(a). Both, start and end point are w slightly perturbed copies of a single disk touching the boundary of the gadget.*

It thus holds that inside an edge gadget for an edge $\{u, v\}$ of weight w, removing w many edges (from the cabin) in G_1 disconnects u from v inside the subgraph of G_1 corresponding to this edge gadget, but removing fewer than w edges leaves u and v connected in this subgraph. Note that if the w paths in the cabin weren't disjoint, it didn't necessarily hold that removing w edges from the unit disk graph in the cabin separates the two endpoint of the gadgets (since arbitrarily intersecting disks can form highly connected unit disk graphs.)

Definition 23 (Vertex Gadget)**.** *For a vertex v, the vertex gadget consists of 16 slightly perturbed copies of a cycle of $C_V = \lceil \pi s/r \rceil$ disks of radius r which are centered on a circle of radius s around v as shown in Figure 6(b). We denote the arrangement of those 16 copies by σ. Furthermore, we place a* centroid disk *D_v centered at the coordinates of v and connect it to σ by 16 slightly perturbed copies of a path of radius r disks which we denote by γ in Figure 6(b).*

We build a set $S_1 \subseteq V_1$ of terminals, by putting the vertex corresponding to the centroid disk D_v into S_1 for every vertex gadget, corresponding to a vertex $v \in S_2 \subseteq V_2$. From this set S_1, together with the unit disk graph G_1 constructed above, we obtain an instance I_1 of the Unweighted Multiterminal Cut problem on unit disk graphs.

Since both edge and vertex gadgets have the same dimension as the gadgets of the previous section, it holds that all edge gadgets are pairwise disjoint, all vertex gadgets are pairwise disjoint and no edge gadget intersects any vertex gadget other than the ones of its end vertices. There, each of the w copies of the last disks intersect at least one 16-disk cluster of the vertex gadget. Note that in the vertex gadget, each vertex in the corresponding subgraph of G_1 has degree at least 16. Thus removing less than 16 edges in this subgraph does not disconnect any two vertices in this subgraph; therefore, an optimal solution for I_1 does not remove any edge from the subgraphs of the vertex gadgets. Furthermore, it is easy to see that $E_2' \subseteq E_2$ is a solution of weight k for I_2 if and only if S_1 gets disconnected in I_1, by removing a total of k edges in the edge gadgets corresponding to the edges in E_2'. Since the reduction from I_2 to I_1 can be done in polynomial time, Theorem 20 follows. □

5 All-Cells-Connection Problem for Unit Disks

While the Point Isolation problem can be interpreted as asking for the minimum number of sensors which have to be turned *on* to detect any transition between the input points, the All-Cells-Connection problem, as defined below, can be interpreted as asking for the minimum number of sensors which have to be turned *off* so that an intruder can transition between *any* connected region induced by the sensors (see Figure 7). In the context of line segment sensors, the All-Cells-Connection problem was shown to be NP-complete in [3] and we reuse their reduction idea in this section for unit disk sensors.

Problem 24 (All-Cells-Connection for unit disks)**.** *Given a set $\mathcal{D}$ of unit disks embedded in the plane, remove a minimum cardinality subset $\mathcal{D}'$ of $\mathcal{D}$ such that the plane minus the arrangement of the remaining disks, i.e., $\mathbb{R}^2 \setminus \bigcup(\mathcal{D} \setminus \mathcal{D}')$, consists of a single connected region.*

Theorem 25. *The All-Cells-Connection problem for unit disks is NP-complete.*

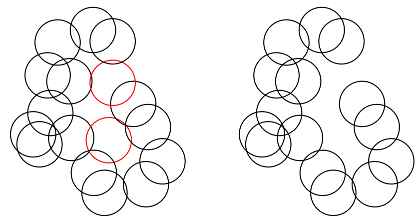

Figure 7: An example of the All-Cells Connection problem for unit disks. (a) A set $\mathcal{D}$ of unit disks with its *merging* subset $\mathcal{D}'$ in red. (b) Illustration of $\mathbb{R}^2 \setminus \bigcup(\mathcal{D} \setminus \mathcal{D}')$, consisting of a single connected region.

Proof. For a given graph, the Feedback Vertex Set problem (FVS) asks for the minimum cardinality subset of vertices to be removed such that the remaining subgraph is acyclic. In order to prove Theorem 25, we are going to reduce a restricted version of FVS to the All-Cells Connection problem for unit disks. In [18] (see also [19]), it is shown that FVS is NP-complete in planar graphs of maximum degree 4. Given such an instance $G = (V, E)$, we embed it into an $n \times n$ grid and replace each edge with the edge gadget[2] of Definition 17. We further replace each vertex with a vertex gadget, described below; we refer to the collection of all disks used in the vertex and edge gadgets as $\mathcal{D}$.

The idea of the vertex gadget for a vertex v is to center a radius r disk D_v at v and connect the, at most 4, incident edge gadgets to D_v, using simple paths which are pairwise disjoint; below we will argue that this is always possible. It then follows that removing disk D_v merges the up to 4 regions adjacent to D_v. Removing a disk in an edge gadget for an edge $\{u, v\}$ never merges more regions than removing D_v (or D_u). Therefore, G has a FVS of size k if and only if the corresponding All-Cells-Connection Problem for $\mathcal{D}$ has a solution of size k.

We now first provide a conceptual description of the *vertex gadget* and then we will replace its conceptual paths by pairwise disjoint paths of intersecting disks of radius r. For an embedded vertex v, which has $1 \leq i \leq 4$ incident edges $e_1, \ldots e_i$, where in the grid embedding edge e_j has an angle α_j with the x-axis. We center a *cross* at v, which consists of four perpendicular rays $a_1, \ldots, a_4$ introducing four quadrants I, II, III, IV in clockwise order, as shown in Figure 8(a). The cross is oriented such that (1) at least one of the incident edges lies on ray a_1, (2) its clockwise quadrant I does not contain any edge incident to v, (3) quadrant II contains at most one edge incident to v and (4) quadrant II and III together contain at most two edges incident to v. To see that this is always feasible, center the cross on an edge such that the first quadrant is empty, this is always possible by the pigeon hole principle. Now at most one constraint, either (3) or (4), is violated. If constraint (3) is violated, this implies that quadrants III and IV together contain only one edge, which if chosen to lie on the axis (i.e., used to satisfy constraint I) results in a cross satisfying all constraints. If, on the other hand, constraint (4) is violated, this means that quadrants III

[2] The fact that all edge gadgets contain the same amount of disks is irrelevant for the reduction. The only relevant property is that the disks in each edge gadget form a simple path.

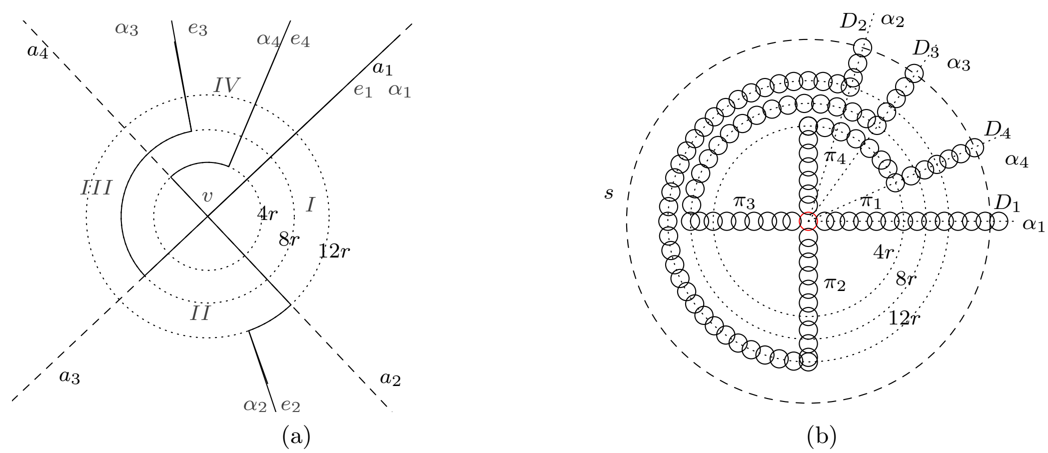

Figure 8: (a) A conceptual illustration of the Vertex gadget. The cross, satisfying the four constraints below, is depicted as dashed rays $a_1, \ldots, a_4$. Edge e_1 lies on a_1. Axes $a_2, \ldots, a_4$ are connected to the edges $e_2, \ldots, e_4$, which are truncated by $12r, 8r, 4r$ from v respectively, by arcs of radius $12r, 8r, 4r$ until reaching $\alpha_2, \alpha_3, \alpha_4$ and thereby $e_2, \ldots, e_4$. (b) An illustration of an actual realization of a vertex gadget, having three edges in quadrant IV. The red disk D_v is centered at v. The dashed circle has radius s and the dotted circles have radii $12r, 8r, 4r$ respectively. D_v is connected by the paths $\pi_1, \ldots, \pi_4$ to the last disks $D_1, \ldots, D_4$ of the edge gadget for the edges $e_1, \ldots, e_4$ respectively, which have an angle of $\alpha_1, \ldots, \alpha_4$ with respect to the x-axis.

and IV together contain 3 edges. Choosing the last of these edges in clockwise orientation to lie on the axis (i.e., used to satisfy constraint I) again results in a cross satisfying all constraints.

Next, we truncate edges $e_2, \dots, e_i$ to end at a distance of $12r, 8r, 4r$ respectively from v. We then connect the truncated edges $e_2, \dots, e_i$ to the rays $a_2, \dots, a_i$ respectively, by extending an arc of radius $12r, 8r, 4r$ from $a_2, \dots, a_i$ until we reach $\alpha_2, \dots \alpha_i$ respectively and thereby orthogonally connecting all the edges $e_1, \dots e_i$ to v, as shown in Figure 8(a). Denoting by π_j the path connecting edge e_j to v, we realize the construction of the vertex gadget by first centering a radius r disk D_v at v. We then center radius r disks on the paths $\pi_1, \dots, \pi_i$ such that consecutive disks intersect. Note that any point in the last disk of any incident edge gadget has distance $s - r = r(6n^2 - 2)$ to v, which is greater than $13r$ for any $n \geq 2$, and thus the whole vertex gadget is contained inside the circle of radius s around v. Furthermore, it is easy to see that all paths of disks $\pi_1, \dots, \pi_i$ are pairwise disjoint.

To show containment in NP, we put a point in each connection region of $\mathbb{R}^2 \setminus \mathcal{D}$. As in the proof of Lemma 19, we use the two-point separation algorithm of [7] to test in polynomial time for each pair of points whether there exists a path not intersecting any disks in $\mathcal{D} \setminus \mathcal{D}'$. It follows that the decision version of the All-Cells Connection problem for unit disks is indeed contained in NP.

Since this reduction can be done in polynomial time, Theorem 25 follows. □

Corollary 26. *Given a set P of k points and a set $\mathcal{D}$ of unit disks in the plane separating P, it is* NP*-complete to find a minimum cardinality subset $\mathcal{D}'$ of $\mathcal{D}$ such that there is a path between any two point in P not intersecting any disk in $\mathcal{D} \setminus \mathcal{D}'$.*

Note that the problem of this corollary is a generalization of the Barrier Resilience problem from 2 to k points. The computational complexity of the Barrier Resilience problem for 2 points is still open in the unit disk setting [6].

Acknowledgments. We thank Alon Efrat for discussions that led to the formulation of the problem, and Sariel Har-Peled for discussions that led to the algorithm described here. We would also like to thank the two anonymous referees for their valuable comments on an earlier version of this manuscript.

References

[1] Rashmisnata Acharyya, Manjanna Basappa, and Gautam Das. Unit disk cover problem in 2d. In *Computational Science and Its Applications, ICCSA 2013*, volume 7972 of *Lecture Notes in Computer Science*, pages 73–85. Springer Berlin Heidelberg, 2013.

[2] Pankaj K. Agarwal and Micha Sharir. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, 1998.

[3] Helmut Alt, Sergio Cabello, Panos Giannopoulos, and Christian Knauer. Minimum cell connection and separation in line segment arrangements. *CoRR*, abs/1104.4618, 2011.

[4] Boris Aronov, Esther Ezra, and Micha Sharir. Small-size epsilon-nets for axis-parallel rectangles and boxes. *SIAM J. Comput*, 39(7):3248–3282, 2010.

[5] Paul Balister, Zizhan Zheng, Santosh Kumar, and Prasun Sinha. Trap coverage: Allowing coverage holes of bounded diameter in wireless sensor networks. In *In Proc. of IEEE INFOCOM, Rio de Janeiro*, 2009.

[6] Sergei Bereg and David G. Kirkpatrick. Approximating barrier resilience in wireless sensor networks. In *Proc. 5th ALGOSENSORS*, pages 29–40, 2009.

[7] Sergio Cabello and Panos Giannopoulos. The complexity of separating points in the plane. In *Proceedings of the twenty-ninth annual symposium on Computational geometry*, pages 379–386. SCG '13, 2013.

[8] Kenneth L. Clarkson and Kasturi Varadarajan. Improved approximation algorithms for geometric set cover. In *Proc. Symposium on Computational Geometry*, pages 135–141. SCG '05, 2005.

[9] Elias Dahlhaus, David Johnson, Christos Papadimitriou, Paul Seymour, and Mihalis Yannakakis. The complexity of multiterminal cuts. *SIAM J. Comput*, 23(4):864–894, 1994.

[10] Reinhard Diestel. *Graph Theory*. Number 173 in Graduate Texts in Mathematics. Springer, 2010.

[11] 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(1):59–71, 1986.

[12] Matias Korman, Maarten Löffler, Rodrigo I. Silveira, and Darren Strash. On the complexity of barrier resilience for fat regions. In Paola Flocchini, Jie Gao, Evangelos Kranakis, and Friedhelm Meyer auf der Heide, editors, *Algorithms for Sensor Systems*, Lecture Notes in Computer Science, pages 201–216. Springer Berlin Heidelberg, 2014.

[13] Santosh Kumar, Ten H. Lai, and Anish Arora. Barrier coverage with wireless sensors. in. In *MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking*, pages 284–298. New York, 2005.

[14] Nabil Mustafa and Saurabh Ray. Improved results on geometric hitting set problems. *Discrete & Computational Geometry*, 44(4):883–895, 2010.

[15] János Pach. On the complexity of the union of geometric objects. In Jin Akiyama, Mikio Kano, and Masatsugu Urabe, editors, *Discrete and Computational Geometry*, volume 2098 of *Lecture Notes in Computer Science*, pages 292–307. Springer Berlin Heidelberg, 2001.

[16] Swaminathan Sankararaman, Alon Efrat, Srinivasan Ramasubramanian, and Javad Taheri. Scheduling sensors for guaranteed sparse coverage. *CoRR*, abs/0911.4332, 2009.

[17] Walter Schnyder. Embedding planar graphs on the grid. In *Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '90, pages 138–148. Society for Industrial and Applied Mathematics, 1990.

[18] Ewald Speckenmeyer. *Untersuchungen zum Feedback Vertex Set Problem in ungerichteten Graphen*. PhD thesis, Paderborn, 1983.

[19] Ewald Speckenmeyer. On feedback vertex sets and nonseparating independent sets in cubic graphs. *Journal of Graph Theory*, 12(3), 1988.

[20] Kuan Tseng, Robert Chieh, and David Kirkpatrick. On barrier resilience of sensor networks. In Thomas Erlebach, Sotiris Nikoletseas, and Pekka Orponen, editors, *Algorithms for Sensor Systems*, volume 7111 of *Lecture Notes in Computer Science*, pages 130–144. Springer Berlin Heidelberg, 2012.