



On the complexity of barrier resilience for fat regions and bounded ply [☆]



Matias Korman ^a, Maarten Löffler ^b, Rodrigo I. Silveira ^{c,*}, Darren Strash ^d

^a Tohoku University, Japan

^b Dept. of Computing and Information Sciences, Utrecht University, The Netherlands

^c Dept. de Matemàtiques, Universitat Politècnica de Catalunya, Spain

^d Department of Computer Science, Colgate University, USA

ARTICLE INFO

Article history:

Received 2 June 2017

Accepted 5 February 2018

Available online 27 February 2018

Keywords:

Barrier resilience and coverage

Fat regions

Bounded ply

Approximation algorithms

Parameterized complexity and algorithms

ABSTRACT

In the *barrier resilience* problem (introduced by Kumar et al., *Wireless Networks* 2007), we are given a collection of regions of the plane, acting as obstacles, and we would like to remove the minimum number of regions so that two fixed points can be connected without crossing any region. In this paper, we show that the problem is NP-hard when the collection only contains fat regions with bounded ply Δ (even when they are axis-aligned rectangles of aspect ratio $1 : (1 + \varepsilon)$). We also show that the problem is fixed-parameter tractable (FPT) for unit disks and for similarly-sized β -fat regions with bounded ply Δ and $O(1)$ pairwise boundary intersections. We then use our FPT algorithm to construct an $(1 + \varepsilon)$ -approximation algorithm that runs in $O(2^{f(\Delta, \varepsilon, \beta)} n^5)$ time, where $f \in O(\frac{\Delta^4 \beta^8}{\varepsilon^4} \log(\beta \Delta / \varepsilon))$.

© 2018 Elsevier B.V. All rights reserved.

1. Introduction

The *barrier resilience* problem asks for the minimum number of spatial regions from a collection \mathcal{D} that need to be removed, such that two given points p and q are in the same connected component of the complement of the union of the remaining regions. This problem was posed originally in 2005 by Kumar et al. [15,16], motivated from sensor networks. In their formulation, the regions are unit disks (sensors) in some rectangular strip $B \subset \mathbb{R}^2$, where each sensor is able to detect movement inside its disk. The question is then how many sensors need to fail before an entity can move undetected from one side of the strip to the opposite one (that is, how *resilient* to failure the sensor system is). Kumar et al. present a polynomial time algorithm to compute the resilience in this case. They also consider the case where the regions are disks in an annulus, but their approach cannot be used in that setting.

1.1. Related work

Despite the seemingly small change from a rectangular strip to an annulus, the second problem still remains open, even for the case in which regions are unit disks in \mathbb{R}^2 . There has been partial progress towards settling the question: Bereg and

[☆] A preliminary version of this work appeared at the 9th International Symposium on Algorithms and Experiments for Sensor Systems, Wireless Networks and Distributed Robotics [14].

* Corresponding author.

E-mail addresses: mati@dais.is.tohoku.ac.jp (M. Korman), m.loffler@uu.nl (M. Löffler), rodrigo.silveira@upc.edu (R.I. Silveira), dstrash@cs.colgate.edu (D. Strash).

Kirkpatrick [3] present a factor $5/3$ -approximation algorithm for the unit disk case. This result was very recently improved to a 1.5 -approximation by Chan and Kirkpatrick [5]. On the negative side, Alt et al. [2], Tseng and Kirkpatrick [23], and Yang [26, Section 5.1] independently showed that if the regions are line segments in \mathbb{R}^2 , the problem is NP-hard. Tseng and Kirkpatrick [23] also sketched how to extend their proof for the case in which the input consists of (translated and rotated) copies of a fixed square or ellipse.

The problem of covering barriers with sensors has received a lot of attention in the sensor network community (e.g., [6, 7,12]). In the algorithms community, closely related problems involving region intersection graphs have also become quite popular. Gibson et al. [11] study a problem that is, in a sense, opposite of ours: given a set of points and disks separating them (i.e., every path between two points intersects some disk), compute the maximum number of disks one can remove while keeping the points separated. They present a constant-factor approximation algorithm for this problem. Later, Penninger and Vigan showed that the problem is NP-complete [21]. Recently, Cabello and Giannopoulos [4] gave a cubic-time algorithm for the case where only two points have to be kept separated, for barriers that are arbitrary connected curves (under some mild assumptions).

1.2. Results

We present constructive results for two natural restricted variants of the problem. In Section 3 we show that the problem is fixed-parameter tractable on the resilience when the regions are unit disks. We then extend this approach to other shapes that resemble unit disks. This resemblance is measured with the following three restrictions: all regions are of similar size, region boundaries have $O(1)$ pairwise intersections, and the collection of regions have bounded *ply* [19] (that is, no point of the plane is covered by too many sensors). Such restrictions are similar in spirit to previous results that bound the union complexity of fat (and non-fat) regions [8,9,24]. Formal definitions of fatness, ply, and more detailed descriptions of our restrictions are given in Section 3.2. In Section 4 we also show that the FPT result can be used to obtain an approximation scheme. In particular, the constructive results apply to the original unit disk coverage setting when the collection of disks (or in general fat objects) has bounded ply.

As a complement to these algorithms, in Section 5 we show that the problem is NP-hard even when the input is a collection of fat regions of arbitrary shape in \mathbb{R}^2 . The result holds even if regions consist of axis-aligned rectangles of aspect ratio $1 : 1 + \varepsilon$ and $1 + \varepsilon : 1$. Our results rely on tools and techniques from both computational geometry and graph theory.

2. Preliminaries

We denote with p and q the points that need to be connected, and with \mathcal{D} the set of regions that represent the sensors. To simplify the presentation of our results, we make the following general position assumption: all intersections between boundaries of regions in \mathcal{D} consist of isolated points. We say that a collection of objects in the plane are *pseudodisks* if the boundaries of any two of them intersect at most twice.

We formally define the concepts of *resilience* and *thickness* introduced in [3]. The *resilience of a path* π between two points p and q , denoted $r(\pi)$, is the number of regions of \mathcal{D} intersected by π . Given two points p and q , the *resilience of p and q* , denoted $r(p, q)$, is the minimum resilience over all paths connecting p and q . In other words, the resilience between p and q is the minimum number of regions of \mathcal{D} that need to be removed to have a path between p and q that does not intersect any region of \mathcal{D} . Note that sometimes we will assume that neither p nor q are contained in any region of \mathcal{D} , since such regions must always be counted in the minimum resilience paths, hence we can ignore them (and update the resilience we obtain accordingly).

Often it will be useful to refer to the arrangement (i.e., the subdivision of the plane into faces; see [18, Section 6.2]) for a formal definition) induced by the regions of \mathcal{D} , which we denote by $\mathcal{A}(\mathcal{D})$. Based on this arrangement we define a weighted dual graph $G_{\mathcal{A}(\mathcal{D})}$ as follows (refer to Fig. 1). There is one vertex for each face (i.e., 2-dimensional cell) of $\mathcal{A}(\mathcal{D})$. Each pair of neighboring cells A, B is connected in $G_{\mathcal{A}(\mathcal{D})}$ by two directed edges, (A, B) and (B, A) . The weight of an edge is 1 if, when traversing from the starting cell to the destination one, we enter a region of \mathcal{D} (or 0 if we leave a region¹).

The *thickness* of a path π between p and q , denoted $t(\pi)$, equals the number of times π enters a region of \mathcal{D} when traveling from p to q (possibly counting the same region multiple times). Given two points p and q , the *thickness of p and q* , denoted $t(p, q)$, is the value $|\mathfrak{P}_{G_{\mathcal{A}(\mathcal{D})}}(p, q)| + \Delta(p)$, where $\mathfrak{P}_{G_{\mathcal{A}(\mathcal{D})}}(p, q)$ is a shortest path in $G_{\mathcal{A}(\mathcal{D})}$ from the cell of p to the cell of q , and $\Delta(p)$ equals the number of regions that contain p . Also note that the resilience (or thickness) between two points only depends on the cells to which the points belong. Hence, we can naturally extend the definitions of thickness to encompass two cells of $\mathcal{A}(\mathcal{D})$, or a cell and a point. Unless otherwise stated, we will use ρ to denote a path with minimum resilience, and τ for one of minimum thickness.

Note that thickness and resilience can be different (since entering the same region several times has no impact on the resilience, but is counted every time for the thickness). In fact, the thickness between two points can be efficiently computed in polynomial time using any shortest path algorithm for weighted graphs (for example, using Dijkstra's algorithm). However, as we will see later, the thickness (and the associated shortest path) will help us find a path of low resilience.

¹ Note that no other option is possible under our general position assumption.

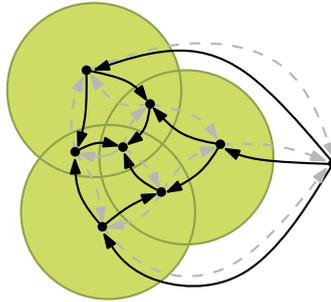


Fig. 1. The graph $G_{\mathcal{A}(\mathcal{D})}$ for an arrangement of three disks. Solid edges have weight 1, dashed edges have weight 0.

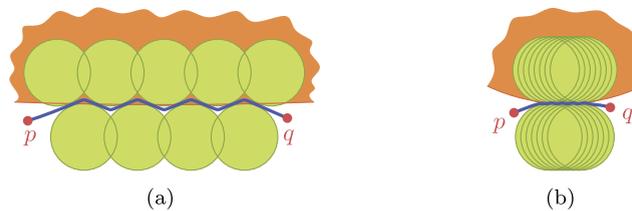


Fig. 2. (a) With n unit disks and one arbitrarily large disk (orange), the optimal tour may be forced to enter and leave the same region up to $\Omega(n)$ times, even when the ply of the disks is at most 3. (b) When we move the yellow disks closer together, the radius of the orange disk can be made arbitrarily close to 1, at the cost of increasing the ply (i.e., having many disks covering the same point). (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

Throughout the paper we often use the following fundamental property of disks, already observed in [3]. In the statement below, “well-separated” is in the sense used in [3]—that is, the distance between p and q is at least $2\sqrt{3}$.²

Lemma 1 ([3], Lemma 1). *Let \mathcal{D} be a set of unit disks, and let ρ be a path from p to q of minimum resilience. If p, q are well-separated, then ρ encounters no disk of \mathcal{D} more than twice.*

Corollary 1 ([3]). *When the regions of \mathcal{D} are unit disks, the thickness between two well-separated points is at most twice their resilience.*

Note that a crucial property in the above results is that all disks have the same size. In Fig. 2(a) we show problem instances with a single large disk that has to be traversed a linear number of times in any minimum resilience path. The same instance is then modified in Fig. 2(b) so that the radius of the larger disk is only $1 + \varepsilon$ times larger than the radius of the other disks (at the expense of concentrating all disks at the same point).

3. Fixed-parameter tractability

In this section we introduce a single-exponential fixed-parameter tractable (FPT) algorithm, where the parameter is the resilience of the optimal solution. Thus, our aim is to obtain an algorithm that given a problem instance, determines whether or not there is a path of resilience r between p and q , and runs in $O(2^{f(r)}n^c)$ time for some constant c and some polynomial function f .

For clarity we first explain the algorithm for the special case of unit disks. Afterwards, in Section 3.2, we show how to adapt the solution to the case in which \mathcal{D} is a collection of β -fat objects. Note that for treating the case of unit disk regions we assume that p and q are well-separated, so we can apply Lemma 1. This requirement is afterwards removed in Section 3.2.

First we give a quick overview of the method of Kumar et al. [15] for open belt regions. Their idea consists of considering the intersection graph of \mathcal{D} together with two additional artificial vertices s_a, t_a with some predefined adjacencies. There is a path from the bottom side to the top side of the belt if and only if there is no path between s_a and t_a in the graph. Hence, computing the resilience of the network is equal to finding a minimum vertex cut between s_a and t_a .

We start by giving a bird’s-eye view of our algorithm. Let ρ be a path of minimum resilience from p to q , and let π be any known path that starts at p , passes through q , and reaches an unbounded region. Assume that somehow we know that

² Note that the well-separatedness of p and q is used to prove a factor 2 instead of 3. Everything still works for points that are not well-separated, at a slight increase of the constants. Our most general statements for β -fat regions do not make this requirement.

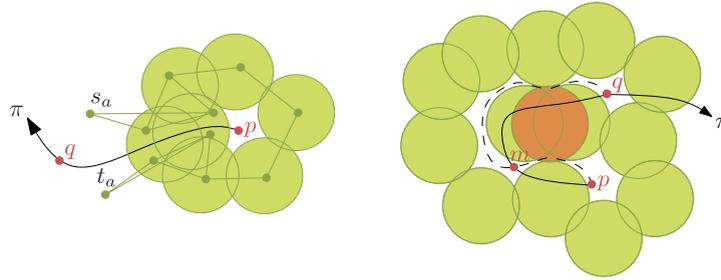


Fig. 3. (left) If we are given an infinitely long path π (in black in the figure) that goes through p and q , and is not crossed by ρ , we can cut open through it and obtain an open belt instance. The resulting graph (with the artificial vertices) is shown for clarity. (right) When the two paths intersect (ρ denoted with a dashed path) we obtain several open belt problem instances. However, these problems are not independent, since the removal of the highlighted disk makes the paths from p to m and from m to q feasible.

ρ and π do not cross (other than at p and q). Then, we can cut open through π effectively splitting the regions of $\mathcal{A}(\mathcal{D})$ traversed by the path into two. Topologically speaking, we get something that is homeomorphic to an open belt region, and thus we can solve the problem as such: construct the intersection graph, connect the split regions of $\mathcal{A}(\mathcal{D})$ to either of the artificial vertices depending on which side of the cut they lie in, and look for a minimum vertex cut (see Fig. 3, left). Note that, when doing this cut, it is possible that a disk is split into more than one component. Whenever this happens, we must identify the portions as one (i.e., when one portion is entered, then entering the other portions of the same disk is *for free*).

Thus, the problem is easy once we have a path π that does not cross with ρ . Unfortunately, finding such a path is difficult. Instead, we use several observations to compute a (possibly non-simple) path that cannot have many crossings with ρ , and guess where (if any) these crossings happen. Naturally, we don't know the way in which the two paths interact, but we will try all possibilities and return the one whose resulting resilience is smallest. A fixed crossing pattern decomposes ρ into subpaths whose endpoints are in π (see Fig. 3, right). Although the subpaths are unknown, we can compute them via the usual open belt region approach. The main problem is that the different sub-problems are not independent (removing a single region may be useful for several subpaths). Thus, rather than finding a vertex cut that isolates the single source to the single sink, we are given a list of sources and sinks that need to be pairwise disconnected from each other. In the literature, this problem is known as the *vertex multicut problem* [25], and several FPT approaches are known.

We now present some observations that will allow us to have a nice choice of π (i.e., find a path in which the number of crossings with ρ does not depend on n). Consider a minimum resilience path ρ of shortest length between the cells containing p and q in $G_{\mathcal{A}(\mathcal{D})}$, and let t be the number of disks traversed by ρ . Since ρ has shortest length, it does not enter and leave the same region unless it helps reduce resilience. Since we assumed that p is not contained in any region, t is exactly the thickness of p and q . We observe that cells with high thickness to p or q can be ignored when we look for low resilience paths.

Lemma 2. *The minimum resilience path ρ between p and q cannot traverse cells whose thickness to p or q is larger than $1.5t$.*

Proof. We argue about thickness to p ; the argument with respect to q is analogous. Let ρ be a path of minimum resilience between p and q , and let r be the resilience of ρ . Also, let τ be a minimum-thickness path from p to q . Recall that ρ does not enter a disk more than twice, hence the thickness of ρ is at most $2r \leq 2t$. Assume, for the sake of contradiction, that the thickness of some cell C traversed by ρ is greater than $1.5t$. Let ρ_C be the portion of ρ from C to q . Since the thickness of ρ from p to q is at most $2t$, the triangular inequality implies that the thickness of ρ_C is less than $0.5t$.

Now, by concatenating τ and ρ_C , we would obtain a path that connects p with C whose thickness is less than $1.5t$, giving a contradiction with the thickness of cell C . \square

For simplicity in the exposition, we will also bound the region to consider (thus, we discard regions with very high resilience since they will not be traversed by ρ). Let R be the union of the cells of the arrangement that have thickness from p at most $1.5t$; we call R the *domain* of the problem. Observe that R is connected, but need not be simple (see Fig. 4(a)).

For simplicity in the explanation, we add additional discs surrounding \mathcal{D} so as to make sure that the unbounded face has thickness more than $1.5t$. This does not affect the asymptotic behavior of our algorithm, but it removes the need of considering some degenerate situations. Note that the number of cells remaining in R might still be quadratic, hence asymptotically speaking the instance size has not decreased (the purpose of this pruning will become clear later).

Lemma 3. *There exists a point q' on the outer boundary of R and a tree that spans p , q , and q' that has total thickness³ $2.5t$.*

³ The thickness of the tree is defined as the thickness of the paths that compose the tree.

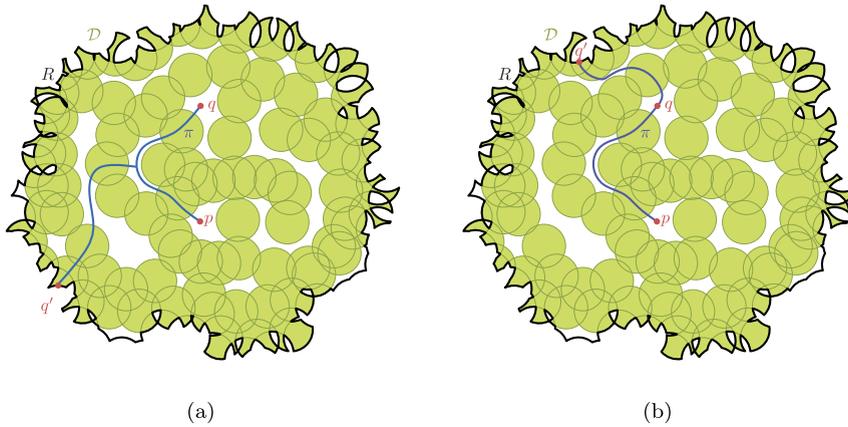


Fig. 4. In order to transform our problem to one that resembles an open belt, we remove all cells of high thickness and cut through the tree formed by the union of two shortest paths. Figures (a) and (b) show two examples of the result.

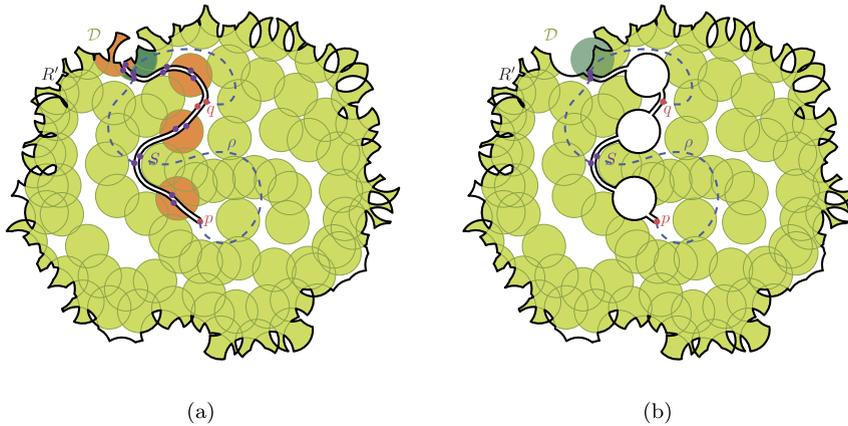


Fig. 5. (a) Illustration of the situation in Fig. 4(b) after cutting along π . We get the domain R' , add a set S of extra vertices on the boundary of R' , and end up with two copies of q . A crossing pattern, consisting of a topological path ρ (defined by the sequence of points of S it passes). The disks of \mathcal{D} intersected by π are shown green if they are crossed by ρ , and orange otherwise. (b) Domain after removing the disks traversed by π that are not crossed by ρ . The green disk (shown transparent) is added to the solution, and thus ignored from now on.

Proof. Pick any point q' in the outer boundary of R and consider the tree obtained by joining the shortest paths from q' to p , and p to q . Note that the two paths may go through the same cell of R , see Fig. 4(a). The exact paths chosen are not important provided that they have no proper crossings. By definition, the thickness of each of these paths cannot exceed $1.5t$ and t , respectively, hence the lemma is shown. \square

Let π be the path from q' to q' that traverses the tree from the previous lemma. We “cut open” through π , removing it from our domain. Note that cells that are traversed by π are split into two copies (or even three, if the tree has a vertex of degree three) of the same Jordan curve, see Fig. 5(a).

Consider now a minimum resilience path ρ , and let $r = r(\rho)$ denote its resilience. This path can cross π several times, and it can even coincide with π in some parts (shared subpaths). Although we do not know how and where these crossings occur, we can *guess* (i.e., try all possibilities) the topology of ρ with respect to π . For each disk that π passes through, we consider two cases: if ρ goes through it, it will be part of the solution, and can be ignored from now on (increasing by one the total resilience). Otherwise, we make it an obstacle, removing it from the domain, see Fig. 5(b). In that way we know the exact behavior of ρ in the regions traversed by π . Additionally, we guess how many times ρ and π share part of their paths (either for a single crossing in one cell, or for a longer shared subpath). For each shared subpath, we guess from which cell ρ arrives and leaves.

We call each such configuration a *crossing pattern* between π and ρ . More formally, a single crossing is described by a tuple of four cells: the first cell C that the two paths have in common for that crossing, the cell that ρ visits right before entering C . Similarly, we add the last cell that the two paths have in common and the cell that is afterwards entered by ρ . A crossing pattern is described by a sorted list of all the crossings that π and ρ have.

Lemma 4. For any problem instance \mathcal{D} , there are at most $2^{4r \log r + o(r \log r)}$ crossing patterns between π and ρ , where $r = r(\rho)$.

Proof. First, for all disks in π , we guess whether or not they are also traversed by ρ . By Lemma 3, π has thickness at most $2.5t$, there are at most such many disks (hence up to $2^{2.5t}$ choices for which disks are traversed by ρ).

Now observe that π cannot traverse many cells of $\mathcal{A}(\mathcal{D})$: when moving from a cell to an adjacent one, we either enter or leave a disk of \mathcal{D} . Since we cannot leave a disk we have not entered and π has thickness at most $2.5t$, we conclude that at most $5t$ cells will be traversed by π (other than the starting and ending cells).

We now bound the number of (maximal) shared subpaths between ρ and π : recall that ρ passes through exactly $r = r(\rho)$ disks, and visits each disk at most twice. Hence, there cannot be more than $2r$ shared subpaths. For each shared subpath we must pick two of the cells traversed in π (as candidates for first and last cell in the subpath). By the previous observation there are at most $5t$ candidates for first and last cell (since that is the maximum number of cells traversed by π). Additionally, for each shared subpath we must determine from which side ρ entered and left the subpath; in most cases we have two options for entering and leaving (since most cells are split into two by π). However, it could happen that the first, last (or even both cells) are the cell containing m . The cell containing m was split into three, and thus we have three options on which part of the cell ρ enters or leaves. That is, on the worst case there are three possibilities where ρ enters and three possibilities where ρ leaves the path, which gives a total of nine options overall. Since these choices are independent, in total we have at most $2r \times (5t \times 5t \times 9)^{2r} = 101250^r \cdot t^{4r} r$ possibilities.

That is, in order to determine a crossing pattern, we must fix which disks of π are traversed by ρ as well as how many and where do the crossings between ρ and π happen. The bounds for each of these terms are $2^{2.5t}$ and $101250^r \cdot t^{4r} r$, respectively. Since these choices are independent, and using the fact that $t \leq 2r$, we obtain:

$$\begin{aligned} 2^{2.5t} \cdot 101250^r \cdot t^{4r} r &\leq 2^{5r} \cdot 101250^r \cdot (2r)^{4r} r \\ &= 2^{5r+r \log 101250+4r \log 2r+\log r} \\ &= 2^{4r \log r + o(r \log r)} \quad \square \end{aligned}$$

Note that the bound is very loose, since most of the choices will lead to an invalid crossing pattern. However, the importance of the lemma is in the fact that the total number of crossing patterns only depends on r .

Our FPT algorithm works by considering all possible crossing patterns, finding the optimal solution for a fixed crossing pattern, and returning the solution of smallest resilience. From now on, we assume that a given pattern has been fixed, and we want to obtain the path of smallest resilience that satisfies the given pattern. If no path exists, we simply discard it and associate infinite resilience to it.

3.1. Solving the problem for a fixed crossing pattern

Recall that the crossing pattern gives us information on how to deal with the disks traversed by π . Thus, we remove all cells of the arrangement that contain one or more disks that are forbidden to ρ . Similarly, we remove from \mathcal{D} the disks that ρ must cross. After this removal, several cells of our domain may be merged.

Since we do not use the geometry, we may represent our domain by a disk W (possibly with holes). After the transformation, each remaining region of \mathcal{D} becomes a pseudodisk, and ρ becomes a collection of disjoint partial paths, each of which has its endpoints on the boundary of W (see Fig. 6(a)), but is otherwise not yet fixed. To solve the subproblem associated with the crossing pattern we must remove the minimum number of disks so that all partial paths are feasible.

We consider the intersection graph G_I between the remaining regions of \mathcal{D} . That is, each vertex represents a region of \mathcal{D} , and two vertices are adjacent if and only if their corresponding regions intersect. Similarly to [15], we must augment the graph with boundary vertices. The partial paths split the boundary of R into several components. We add a vertex for each component (these vertices are called *boundary vertices*). We connect each such vertex to vertices corresponding to pseudodisks that are adjacent to that piece of boundary (Fig. 6(b)). Let $G_{\mathcal{X}} = (V_{\mathcal{X}}, E_{\mathcal{X}})$ be the resulting graph associated to crossing pattern \mathcal{X} . Note that no two boundary vertices are adjacent.

We now create a secondary graph H as follows: the vertices of H are the boundary vertices of $G_{\mathcal{X}}$. We add an edge between two vertices if there is a partial path that separates the vertices in $G_{\mathcal{X}}$ (Fig. 6(c)). Two vertices connected by an edge of H are said to form a *forbidden pair* (each partial path that would create the edge is called a *witness partial path*). We first give a bound on the number of forbidden pairs that H can have.

Lemma 5. Any crossing pattern has at most $2r^2 + r$ forbidden pairs.

Proof. By definition, $G_{\mathcal{X}}$ only adds edges between boundary vertices. Thus, it suffices to show that $G_{\mathcal{X}}$ has at most $2r + 1$ boundary vertices. Since partial paths cannot cross, each such path creates a single cut of the domain. This cut introduces a single additional boundary vertex (except the first partial path that introduces two vertices). Recall that we can map the partial paths to crossings between paths π and ρ and, as argued in the proof of Lemma 4, these paths can cross at most $2r$ times. Thus, we conclude that there cannot be more than $2r + 1$ boundary vertices. \square

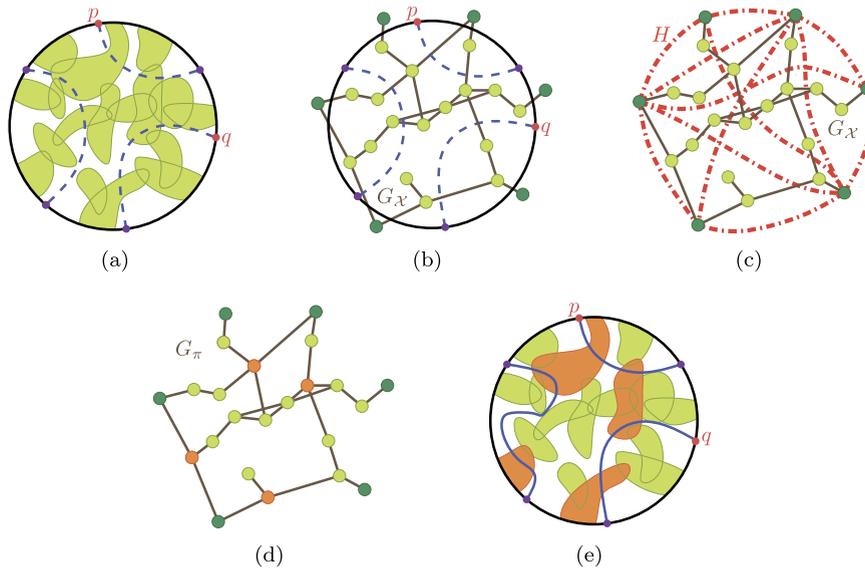


Fig. 6. (a) We may schematically represent W as a circle, since the geometry no longer plays a role. Partial paths are dashed (note that we do not know through which disks these paths will traverse). (b) The intersection graph of the regions after adding extra vertices for boundary pieces between points of $S \cup \{p, q\}$, shown in green. (c) The secondary graph H , representing the forbidden pairs. (d) A possible solution of the vertex multicut problem (highlighted in orange). (e) The corresponding cut for the original problem. Once the orange disks have been removed, the endpoints of the partial paths belong to the same region, and thus we can connect them without entering any additional disk (solid paths).

The following lemma shows the relationship between the vertex multicut problem and the minimum resilience path for a fixed pattern.

Lemma 6. *There are k vertices of $G_{\mathcal{X}}$ whose removal disconnects all forbidden pairs if and only if there are k disks in \mathcal{D} whose removal creates a path between p and q that obeys the crossing pattern \mathcal{X} .*

Proof. Consider the regions of $\mathcal{A}(\mathcal{D})$ inside R that are not covered by any disk after the k disks have been removed and let R' be their union. By definition, there is a path between p and q with the fixed crossing pattern if all partial paths are feasible (i.e., there exists a path connecting the two endpoints that is totally within R'). The reasoning for each partial path is analogous to the one used by Kumar et al. [15]. If all partial paths are possible, then no forbidden pair can remain connected in $G_{\mathcal{X}}$, since—by definition—each forbidden pair disconnects at least one partial path (the witness path). On the other hand, as soon as one forbidden pair remains connected, there must exist at least one partial path (the witness path) that crosses the forbidden pair. Thus if a forbidden path is not disconnected, there can be no path connecting p and q for that crossing pattern. \square

Using Lemma 6, we can transform the barrier resilience problem to the following one: given two graphs $G = (V, E)$, and $H = (V, E')$ on the same vertex set, find a set $D \subset V$ of minimum size so that no pair $(u, v) \in E'$ is connected in $G \setminus D$. This problem is known as the (vertex) multicut problem [25]. Although the problem is known to be NP-hard if $|E'| > 2$ [13], there exist several FPT algorithms on the size of the cut and on the size of the set E' [17,25]. Among them, we distinguish the method of Xiao ([25], Theorem 5) that solves the vertex multicut problem in roughly $O((2k)^{k+\ell/2}n^3)$ time, where k is the number of vertices to delete, $\ell = |E'|$, and n is the number of vertices of G .

Theorem 1. *Let \mathcal{D} be a collection of unit disks in \mathbb{R}^2 , and let p and q be two well-separated points. There exists an algorithm to test whether $r(p, q) \leq r$, for any value r , and if so, to compute a path with that resilience, in $O(2^{f(r)}n^3)$ time, where $f(r) = r^2 \log r + o(r^2 \log r)$.*

Proof. Recall that our algorithm considers all possible crossings between ρ and π . For any fixed crossing pattern \mathcal{X} , our algorithm computes $G_{\mathcal{X}}$, and all associated forbidden pairs. We then execute Xiao's FPT algorithm [25] for solving the vertex multicut problem. By Lemma 6, the number of removed vertices (plus the number of disks that were forced to be deleted by \mathcal{X}) will give the minimum resilience associated with \mathcal{X} .

Regarding the running time, the most expensive part of the algorithm is running an instance of the vertex multicut problem for each possible crossing pattern. Observe that the parameters k and ℓ of the vertex multicut problem are bounded by functions of r as follows: $k \leq r$ and $\ell \leq 2r^2 + r$ (the first claim is direct from the definition of resilience, and the second one follows from Lemma 5). Hence, a single instance of the vertex multicut problem will need

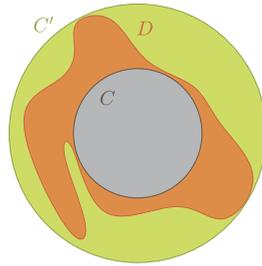


Fig. 7. A β -fat region D is contained in a big disk, but contains a smaller disk; in this example, $\beta = 2$.

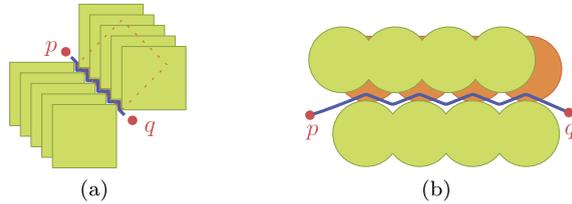


Fig. 8. If we eliminate any one of our restrictions we can construct a problem instance whose minimum resilience path must leave and reenter the same (orange) region $\Theta(n)$ times. Here are constructions when removing one of our three restrictions: (a) bounded ply, and (b) bounded region complexity. Note that the case of distinct size was already discussed in Figs. 2(a) and 2(b).

$O((2r)^{r+(2r^2+r)/2}n^3) = O(2^{(1+\log r)(r^2+1.5r)}n^3) = O(2^{r^2 \log r + o(r^2 \log r)}n^3)$ time. By Lemma 4 the number of crossing patterns is bounded by $2^{4r \log r + o(r \log r)}$. Thus, by multiplying both expressions we obtain the bound on the running time, and the theorem is shown. \square

We remark that the importance of this result lies in the fact that an FPT algorithm exists. Hence, although the dependency on r is high, we emphasize that the bounds are rather loose. We also note that both the minimum resilience path and the disks to be deleted can be reported.

3.2. Extension to fat regions

We now generalize the algorithm to consider more general shapes. A region D is β -fat if there exist two concentric disks C and C' whose radii differ by at most a factor β , such that $C \subseteq D \subseteq C'$ (whenever the constant β is not important, the region D is simply called *fat*). Fig. 7 shows an example of a 2-fat region. However, for our algorithms, it is not sufficient for us to assume that the regions are fat. We impose three restrictions on our fat regions, which make them more like disks: (1) the collection of regions has bounded ply Δ , (2) all regions have similar size, allowing us to assume the radius of C is 1, and the radius of C' is β , and (3) any two regions have $O(1)$ intersections between their boundaries. Together, these three restrictions ensure that no minimum resilience path traverses a given region more than a constant number of times, making thickness within a constant factor of resilience. We formally describe each restriction, and illustrate how its removal impacts the path complexity.

Bounded ply The arrangement formed by a collection of regions \mathcal{D} is said to have bounded ply Δ if no point $p \in \mathbb{R}^2$ is contained in more than Δ elements of \mathcal{D} . As we illustrate in Fig. 8(a), we can place regions of similar size and bounded region complexity (but no bounded ply) forming a corridor. In particular, the minimum resilience path between s and t may be forced to leave and reenter another similarly-sized region $\Theta(n)$ times. Note that this construction is not possible for unit disks, and therefore unit disk instances do not require bounded ply; however, as soon as we allow a disk with larger radius (e.g., a disk of radius $1 + \epsilon$, $\epsilon > 0$), the bounded ply restriction is required.

Similar size We assume without loss of generality that the radius of C is 1 and the radius of C' is β ; in this case we will call D a β -fat unit region. As previously shown in Figs. 2(a) and 2(b), with the existence of a single larger region we can create a corridor of $\Theta(n)$ small interlocking regions with constant ply, and partially cover it with a large region to force the optimal resilience path to leave and reenter the large region $\Theta(n)$ times.

Bounded region complexity Our final assumption is that the fat regions cannot be too complex. In particular, we assume that any two region boundaries have $O(1)$ pairwise intersections, ensuring that the intersection between any two regions has $O(1)$ connected components. As shown in Fig. 8(b), we can create a corridor with two regions that have $\Theta(n)$ pairwise boundary intersections with a third region, forcing the minimum resilience path to leave and reenter this third region $\Theta(n)$ times. Note that such complex regions can be formed, for example, by taking the union of $\Theta(n)$ circles with radius 1, with centers that are spaced $(\beta - 1)/n$ apart on a line.

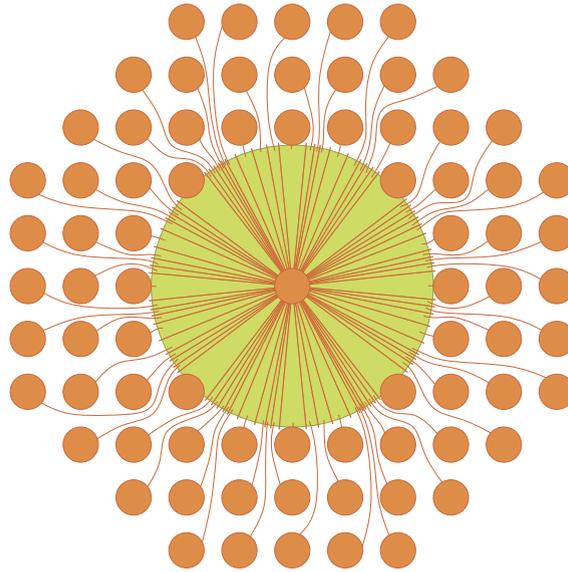


Fig. 9. Example showing we can have $\Theta(\beta^2)$ pairwise disjoint β -fat regions (in orange) that intersect a fixed region (green). Most orange regions consist of the union of a disk and a thin curve-like shape. By placing a constant number of regions, we can force the minimum resilience path to follow around the boundary of the green region, causing it to enter and leave $\Omega(\beta^2)$ times. This construction has overall constant ply, so we can repeat it until we reach the maximum ply Δ and get the $\Omega(\Delta\beta^2)$ lower bound.

Although these restrictions may seem excessive, previous results have made similar assumptions on input regions, and for the same reason we do here: worst-case configurations are possible even with the simplest inputs. For example, to bound the union complexity of fat (α, β) -covered regions, Efrat [9] assumes constant *algebraic complexity*—that region boundaries can be represented by $O(1)$ algebraic polynomials, implying that the region boundaries have at most $O(1)$ pairwise intersections. Whitesides and Zhao [24], when defining k -admissible curves, impose further restrictions on their (non-fat) regions, requiring the difference of any two regions to be connected, in order to guarantee linear-size union boundary (see also [1,20] for alternative proofs of this result). Lastly, de Berg [8] assumes constant density, which bounds the number of regions that can intersect any small disk, similar in spirit to ply.

To our knowledge, no definition of fatness meets any of our three assumptions. Fortunately, our assumptions are not overly restrictive. Indeed, they are representative of cases that we are likely to encounter in practice, as it is inefficient to place sensors so that many of them cover the same region, sensor ranges are typically of similar size, and limiting the boundary intersections encompasses both unit disks and pseudodisks as special cases.

The main workings of the algorithm remain unchanged. We start by extending Lemmas 1, 2, 3, 4 and 5 to consider β -fat unit regions.

Lemma 7. *Let \mathcal{D} be a set of β -fat unit regions forming an arrangement with ply Δ , and bounded region complexity. Let $S \subset \mathcal{D}$ be an optimal solution. In the sequence of regions of S found when going from p to q in an optimal way, no region of S appears more than $O((2\beta + 1)^2 \Delta)$ times.*

Proof. Let D be a region in S , and consider its containing disk C' with center c . Analogously to the original argument by Bereg and Kirkpatrick [3], we note that every time the optimal path visits and leaves D , it must do so to avoid some other region. This other region must intersect D , and since it is β -fat unit, it must contain a unit disk centered at distance at most β from D .

Therefore all regions intersecting D have their unit-disks centered at distance at most 2β from c . In particular, their unit-disks are totally contained in a disk of radius $2\beta + 1$ centered at c . A simple area argument shows that at most $(2\beta + 1)^2$ disjoint unit-disks fit into a disk of radius $(2\beta + 1)$. Since the ply is bounded by Δ , overall there can be up to $\Delta(2\beta + 1)^2$ regions intersecting D . Recall that, by our fatness assumption, two regions can intersect only in $O(1)$ connected components. Therefore, the number of times an optimal path can reenter region D is, proportional to the number of other regions that intersect D which is bounded by $\Delta(2\beta + 1)^2$. \square

We note that our bound is asymptotically tight. Fig. 9 illustrates how a matching lower bound.

Corollary 2. *When the regions of \mathcal{D} are β -fat unit regions forming an arrangement with ply Δ , and bounded region complexity, the thickness between two points is at most $\Delta(2\beta + 1)^2$ times their resilience.*

This change in the upper bound of the thickness in terms of the resilience implies similar changes in Lemmas 2, 3, 4 and 5. The following lemmas summarize these changes; they are proved in the same way as their counterparts for disks, thus we only sketch the differences with the original proofs (if any).

Lemma 8. *When the regions of \mathcal{D} are β -fat unit regions forming an arrangement with ply Δ , and bounded region complexity, the minimum resilience path between p and q cannot traverse cells whose thickness to p or q is larger than $(1 + \Delta(2\beta + 1)^2)\frac{t}{2}$.*

Proof. We use the same reasoning as in the proof of Lemma 2. On the one hand there is the minimum thickness path between p and q , whose thickness is t . On the other hand, we also have the minimum resilience path ρ between the same points, whose thickness is at most $\Delta(2\beta + 1)^2t$ by Corollary 2. Assume now that any cell C traversed by ρ has thickness $k\Delta(2\beta + 1)^2t$ from p , for some $0 < k < 1$. The alternative path goes from p to C , via q , and its thickness is at most $(1 - k)\Delta(2\beta + 1)^2t + t$. The bound we need is obtained for the value of k that makes both expressions equal, which is $k = \frac{1}{2} + \frac{1}{2\Delta(2\beta+1)^2}$, leading to the claimed value. \square

Thus, for β -fat objects our domain R now becomes be the union of the cells of the arrangement that have thickness from p at most $(1 + \Delta(2\beta + 1)^2)\frac{t}{2}$.

Lemma 9. *There exists a point q' on the outer boundary of R and a tree that spans p , q , and q' that has total thickness at most $(3 + \Delta(2\beta + 1)^2)\frac{t}{2}$.*

As before, we use π to denote the path from q' to q' that traverses the tree from the previous lemma.

Lemma 10. *For any problem instance \mathcal{D} , there are at most $2^{O(\Delta^2\beta^4r + \Delta\beta^2r\log(\Delta\beta r))}$ crossing patterns between π and ρ .*

Proof. Let $\mu = \Delta(2\beta + 1)^2$ and $v = \frac{3 + \Delta(2\beta + 1)^2}{2}$. We proceed as in the proof of Lemma 4. Recall that previously we had $2^{2.5t} \times 2r \cdot (5t \times 5t \times 9)^{2r}$ crossing patterns, but now we must use the bounds that depend on β instead. What before was $2r$ now becomes μr , and the $2.5t$ terms now become vt . Making these changes in the previous expression, we obtain that the number of crossings is bounded by

$$2^{vt} \times \mu r \times (2vt \times 2vt \times 9)^{\mu r}.$$

Since $t \leq \mu r$ (and by simplifying the expression), this is upper bounded by

$$2^{v\mu r} \times \mu r \times (6v\mu r)^{2\mu r} = 2^{v\mu r + \log(\mu r) + 2\mu r \log(6v\mu r)}.$$

Finally, we apply that both $\mu, v \in O(\Delta\beta^2)$, and obtain the desired bound. \square

Lemma 11. *Any crossing pattern has at most $O(\Delta^2\beta^4r^2)$ forbidden pairs.*

Proof. As in the unit disc case, each crossing between π and ρ creates an additional vertex in the boundary (i.e., a potential vertex of H). Further note that π and ρ can cross at most $2\mu r$ times (since they traverse through at most that many cells of $\mathcal{A}(\mathcal{D})$). A bound on the number of vertices of H immediately implies a quadratic bound on the number of edges in H as well. Thus, we obtain that the number of forbidden pairs is at most $O((2\mu r + 1)^2) = O(\Delta^2\beta^4r^2)$ as claimed. \square

With these results in place, the rest of the algorithm remains unchanged: the only additional property of unit disks that we use is the fact that they are connected, to be able to phrase the problem as a vertex cut in the region intersection graph.

Theorem 2. *Let \mathcal{D} be a collection of n connected β -fat unit regions of bounded region complexity in \mathbb{R}^2 forming an arrangement of ply Δ , and let p and q be two points. Let r be a parameter. There exists an algorithm to test whether $r(p, q) \leq r$, and if so, to compute a path with that resilience, in $O(2^{f(\Delta, \beta, r)}n^3)$ time, where $f(\Delta, \beta, r) \in O(\Delta^2\beta^4r^2 \log(\Delta\beta r))$.*

Proof. As before, the running time is bounded by the product of the number of crossing patterns and the time needed to solve a single instance of the vertex multicut problem. By Lemmas 10 and 11, these bounds now become $O(2^{O(\Delta^2\beta^4r + \Delta\beta^2r\log(\Delta\beta r))})$ and $O(2^{O(\Delta^2\beta^4r^2 \log(\Delta\beta r))}n^3)$, respectively. The product of both is dominated by the second term, hence the theorem is shown. \square

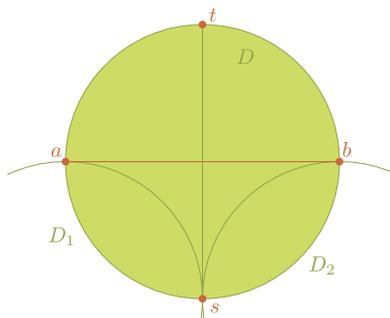


Fig. 10. Barring symmetric configurations, this is the only way of making two disks that cross the segment st (dashed line) and avoids the segment ab (solid line). However, in this case the disks D_1 and D_2 do not separate s and t .

4. $(1 + \epsilon)$ -approximation

In this section we present an efficient polynomial-time approximation scheme (EPTAS) for computing the resilience of an arrangement of disks of bounded ply Δ . The general idea of the algorithm is very simple: first, we compute all pairs of regions that can be reached by removing at most k disks, for $k = \lceil (16\Delta - 12)/\epsilon^2 \rceil$. Then, we compute a shortest path in the dual graph of the arrangement of regions, augmented with some extra edges. We prove that the length of the resulting path is a $(1 + \epsilon)$ -approximation of the resilience.

As in the previous section, we first consider the case in which \mathcal{D} is a set of n unit disks in \mathbb{R}^2 (note that this time we have the additional constraint that no point is covered in more than Δ disks). Let $\mathcal{A}(\mathcal{D})$ be the arrangement induced by the regions of \mathcal{D} , and let $G_{\mathcal{A}(\mathcal{D})}$ be the dual graph of $\mathcal{A}(\mathcal{D})$. Recall that $G_{\mathcal{A}(\mathcal{D})}$ has a vertex for every cell of $\mathcal{A}(\mathcal{D})$, and a directed edge between all pairs of adjacent cells of cost 1 when entering a disk, and cost 0 when leaving a disk. For any given k , let G_k be the graph obtained from $G_{\mathcal{A}(\mathcal{D})}$ by adding, for each pair of cells $A, B \in \mathcal{A}(\mathcal{D})$ with resilience at most k , a shortcut edge \overrightarrow{AB} of cost $r(A, B)$.

For a pair of cells of $\mathcal{A}(\mathcal{D})$, we can test whether $r(A, B)$ is smaller than k , and if so, compute it, in $O(2^{f(k)}n^3)$ time (where $f(k) = k^2 \log k + o(k^2 \log k)$) by applying Theorem 1 to a point $p \in A$ and a point $q \in B$. Since the number of pairs of cells of the arrangement is also bounded by a polynomial in n , we overall get a EPTAS since k is a constant that depends only on ϵ and Δ . Again, we emphasize that the bounds presented in this section are not tight, but our objective is to show the existence of an EPTAS for this problem.

4.1. Analysis

Lemma 12. *Let $D \in \mathcal{D}$, where $\mathcal{A}(\mathcal{D})$ has ply Δ , and let s, t be any two points inside D . Then the resilience between s and t in \mathcal{D} is at most $4\Delta - 3$.*

Proof. Let c be the number of disks that contain s or t (or both). Clearly these disks must be removed. Also notice that $c \leq 2\Delta - 1$, since D contains both points and no point is contained in more than Δ disks. Now we analyze how many other disks may need to be removed too.

Consider a minimum resilience path between s and t among those that stay inside D . For each disk D_1 (not containing neither s nor t) that needs to be removed in an optimal solution, there must be another disk D_2 that intersects D_1 , so that D_1 and D_2 together separate s and t inside D . We call such a pair of disks a *separating pair*. Thus if the resilience is $(c + c')$, there must be at least c' disjoint⁴ separating pairs intersecting D . Let a and b be the diametral pair on D that is orthogonal to segment st . We claim that one of the disks of any separating pair must cover either a or b . Indeed, assume on the contrary that there exists two unit disks D_1 and D_2 that separate s and t but do not contain neither a nor b (nor s or t). Without loss of generality, we may assume that both s and t lie on the boundary of D . Observe that in order to separate s and t , the union of D_1 and D_2 must cross segment st and cannot cross segment ab (otherwise it would contain s or t , since D, D_1 and D_2 are unit disks). However, the only possible way of doing so is if D_1 and D_2 are tangent to a, b and either s or t (see Fig. 10). However, in this case s and t are not separated, a contradiction.

That is, for each separating pair we have a unique disk that covers either a or b . Since no point is contained in more than Δ disks (and D contains both a and b) we conclude that there cannot be more than $2(\Delta - 1)$ separating pairs, completing the proof of the lemma. \square

⁴ By *disjoint* we refer to the identities of the disks, not to the regions they occupy.

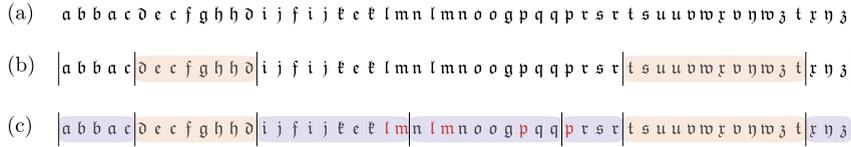


Fig. 11. (a) A string of 52 symbols, each appearing twice. (b) First, we identify a maximal set of *long* segments bounded by equal symbols, and longer than $\lambda = 4$ (shown in red). (c) Then, we segment the remaining pieces into *short* segments, with length $k = 10$ (shown in blue). Red symbols are double-counted.

The previous lemma implies that in an optimal resilience path, if a disk appears twice, the two entry points have resilience at most $4\Delta - 3$ apart (when counting the cells traversed by the path between the two occurrences of the disk). Note that a lower bound of Δ is also easy to construct, so the result is (asymptotically speaking) tight.

To prove the result in this section it will be convenient to focus on the sequence of disks encountered by a path when going from p to q . It turns out that such problem is essentially a string problem, where each symbol represents a disk encountered by the path. In that context, the thickness will be equivalent to the number of symbols of the string (recall that we assume that p is not contained in any disk), and the resilience to the number of distinct symbols.

Let $S = \langle s_1 \dots s_n \rangle$ be a string of n symbols from some alphabet \mathcal{A} , such that no symbol appears more than twice. Let T be a substring of S . We define $\ell(T)$ to be the length of T , and $d(T)$ to be the number of distinct symbols in T . Clearly, $\frac{1}{2}\ell(T) \leq d(T) \leq \ell(T)$. Let σ and k be two fixed integers such that $\sigma < k$. We define the *cost* of a substring T of S to be:

$$\psi(T) = \begin{cases} \sigma & \text{if } T = \langle a\tau a \rangle \text{ for some } a \in \mathcal{A}, \text{ string } \tau \text{ s.t. } a \notin \tau, \text{ and } \ell(T) > \sigma, \\ d(T) & \text{if } \ell(T) \leq k \text{ (and the first condition fails),} \\ \ell(T) & \text{otherwise.} \end{cases}$$

Note that, in the string context, d acts as the resilience, ℓ as the thickness, and ψ is the approximation we compute. Intuitively, if T is short (i.e., length at most k) we can compute the exact value $d(T)$. If T has a symbol whose two appearances are far away we will use a “shortcut” and pay σ (i.e., for unit disk regions, by Lemma 12, we have $\sigma = 4\Delta - 3$). Otherwise, we will approximate d by ℓ .

Given a long string, we wish to subdivide S into a *segmentation* \mathcal{T} , composed of m disjoint segments (i.e., substrings of S) T_1, \dots, T_m , that minimize the total cost $\psi(\mathcal{T}) = \sum_i \psi(T_i)$. Clearly, $\psi(\mathcal{T}) \leq \ell(S)$.

Lemma 13. *Let S be a string. There exists a segmentation \mathcal{T} such that $\psi(\mathcal{T}) \leq (1 + \varepsilon)d(S)$, where $\varepsilon = 2\sqrt{\sigma/k}$.*

Proof. Let λ be an integer such that $\sigma < \lambda < k$, of exact value to be specified later. First, we consider all pairs of equal symbols in S that are more than λ apart. We would like to take all of these pairs as separate segments; however, we cannot take segments that are not disjoint. So, we greedily take the leftmost symbol s whose partner is more than λ further to the right, and mark this as a segment. We recurse on the substring remaining to the right of the rightmost s .⁵ Finally, we segment the remaining pieces greedily into pieces of length k . Fig. 11 illustrates the resulting segmentation. We will refer to the segments of the first type, with length more than λ , as *long* segments. The remaining segments, of length at most k , will be called *short*.

Now, we prove that the resulting segmentation has a cost of at most $(1 + \varepsilon)d(S)$. First, consider a symbol to be *counted* if it appears in only one short (blue) segment, and to be *double-counted* if it appears in two different short segments. Suppose s is double-counted. Then the distance between its two occurrences must be smaller than λ , otherwise it would have formed a long (red) segment. Therefore, it must appear in two adjacent short segments. The leftmost of these two segments has length exactly k , but only λ of these can have a partner in the next segment. So, at most a fraction λ/k symbols are double-counted.

Second, we need to analyze the cost of the long (red) segments. In the worst case, all symbols in the segment also appear in another place, where they were already counted. In this case, the true cost would be 0, and we pay σ too much. However, we can assign this cost to the at least λ symbols in the segment; since each symbol appears only twice they can be charged at most once. So, we charge at most σ/λ to each symbol. The total cost is then bounded by $(1 + \lambda/k + \sigma/\lambda)d(S)$. To optimize the approximation factor, we choose λ such that $\lambda/k = \sigma/\lambda$; more precisely we take $\lambda = \lceil \sqrt{k\sigma} \rceil$. \square

Recall that for our resilience approximation we have $\sigma = 4\Delta - 3$ (Lemma 12). Thus, the actual value of k is obtained by solving $\varepsilon = 2\sqrt{\sigma/k}$ for k , which leads to $k = \lceil (16\Delta - 12)/\varepsilon^2 \rceil$.

⁵ In fact, we could choose any disjoint collection such that after their removal there are no more segments of this type longer than λ .

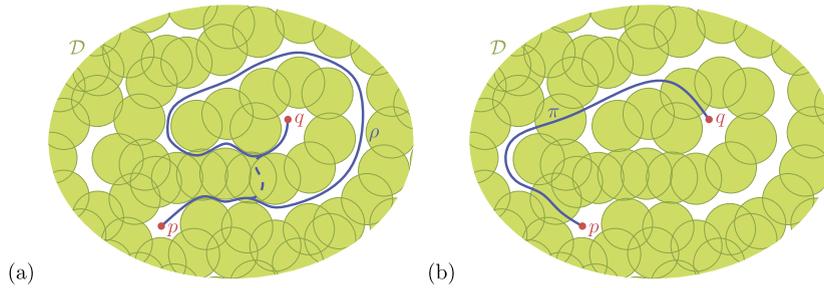


Fig. 12. (a) The optimal path ρ , achieving a resilience of 2. There is a segmentation of ρ of cost 3, using the dashed shortcut. (b) A minimum cost path π found by the algorithm. In this example, the resilience of π is 3.

4.2. Application to resilience approximation

We now show that the shortest path between any p, q in G_k is a $(1 + \epsilon)$ -approximation of their resilience. Let π be a path from p to q in \mathbb{R}^2 , and let $S(\pi)$ be the sequence that records every disk of \mathcal{D} we enter along π , plus the disks that contain the start point of π , added at the beginning of the sequence, in any order. Then we have $|S(\pi)| = t(\pi)$.

Lemma 14. For every path π from p to q and every segmentation \mathcal{T} of $S(\pi)$, there exists a path from p to q in G_k of cost at most $\psi(\mathcal{T})$.

Proof. We describe how to construct a path in G_k based on \mathcal{T} . For every segment T of \mathcal{T} , we create a piece of path whose length in G_k is at most the cost of the segment $\psi(T)$.

There are three types of segments. The first type are segments that start and end with the same symbol a , which corresponds to a disk $D \in \mathcal{D}$. For those, we make a shortcut path that stays inside D , as per Lemma 12. The second type are segments whose length is at most k . For those, by definition, G_k contains a shortcut edge whose cost is exactly the resilience between the corresponding cells of $\mathcal{A}(\mathcal{D})$. The third type are the remaining segments. For those, we simply use the piece of π that corresponds to T . \square

Lemma 15. For any $p, q \in \mathbb{R}^2$, it holds $\text{cost}_{G_k}(\mathfrak{P}_{G_k}(p, q)) \leq (1 + \epsilon)r(p, q)$.

Proof. Let ρ be a path from p to q of optimal resilience $r^* = r(\rho) = r(p, q)$. Then, consider the sequence $S(\rho)$, that is, the sequence of disks that ρ enters. Now, by Lemma 13, there exists a segmentation \mathcal{T} of $S(\rho)$ of cost at most $(1 + \epsilon)d(S(\rho)) = (1 + \epsilon)r^*$. By Lemma 14, there exists a path in G_k of equal or smaller cost. Fig. 12 illustrates this.

Now, consider the path π that our algorithm produces. The resilience of π is smaller than the cost of π in G_k , which is smaller than the cost of ρ in G_k , which is smaller than $1 + \epsilon$ times the resilience of ρ . That is: $r(\pi) \leq \text{cost}_{G_k}(\pi) \leq \text{cost}_{G_k}(\rho) \leq (1 + \epsilon)r(\rho) = (1 + \epsilon)r^*$. \square

Theorem 3. Let \mathcal{D} be a set of unit disks of ply Δ in \mathbb{R}^2 . We can compute a path π between any two given points $p, q \in \mathbb{R}^2$ whose resilience is at most $(1 + \epsilon)r(p, q)$ in $O(2^{f(\Delta, \epsilon)}n^5)$ time, where $f(\Delta, \epsilon) = O\left(\frac{\Delta^2 \log(\Delta/\epsilon)}{\epsilon^4}\right)$.

Proof. The running time of the algorithm is dominated by the preprocessing stage: determining if the resilience between every pair of vertices of $G_{\mathcal{A}(\mathcal{D})}$ is at most $\lceil (16\Delta - 12)/\epsilon^2 \rceil$. Since $G_{\mathcal{A}(\mathcal{D})}$ is an arrangement of disks with ply at most Δ , it has $O(\Delta n)$ cells.⁶ We execute the algorithm of Theorem 1 for every pair of cells (thus, $O(\Delta^2 n^2)$ times), and we obtain the desired bound. \square

4.3. Extension to fat regions

As in Section 3.2, we now generalize the result to arbitrary β -fat unit regions. We again assume that our collection of regions has bounded ply Δ , and that the region boundaries have $O(1)$ pairwise intersections. As in Section 3.2, for simplicity in the notation our analysis assumes that the region boundaries have at most two pairwise intersections, implying that the intersection between any two overlapping regions has one connected component. However, our results generalize to $k = O(1)$ pairwise intersections between region boundaries.

Lemma 16. Let $D \in \mathcal{D}$, where $\mathcal{A}(\mathcal{D})$ has ply Δ , and let p, q be any two points inside D . Then the resilience between p and q in \mathcal{D} is at most $(2\beta + 1)^2 \Delta$.

⁶ We thank the anonymous referee that pointed this to us and allowed the dependency in n to be lowered.

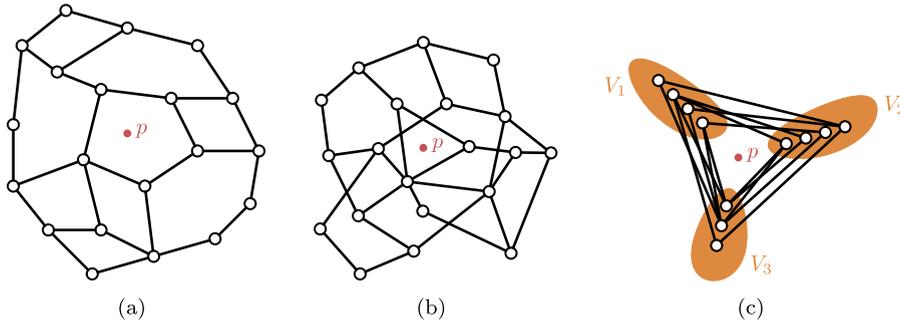


Fig. 13. (a) A planar odd embedded graph. (b) A non-planar one. (c) A tripartite graph, oddly embedded around p .

Proof. The resilience between p and q is upper-bounded by the number of regions that intersect D . We can give an upper bound using a simple packing argument. Since p and q belong to a β -(unit)fat region D , they are both inside a circle C with center c and radius β . Any other β -fat region D' that interferes with the path from p to q must intersect C . Such an intersecting region, being also β -fat, must contain a unit-disk whose center cannot be more than 2β away from c . Therefore all regions intersecting C have their unit-disks centered at distance at most β from c . Moreover, such disks are totally contained in a disk of radius $2\beta + 1$ centered at c . As in the proof of Lemma 7, we can show that at most $(2\beta + 1)^2$ disjoint unit-disks fit into a disk of radius $(2\beta + 1)$. Since the ply is at most Δ , the maximum number of unit-disks inside a disk of radius β in \mathcal{D} is $(2\beta + 1)^2\Delta$. \square

As before, the rest of the arguments do not rely on the geometry of the regions anymore, and we can proceed as in the disk case. The only difference is that the value σ of doing a shortcut has increased to $(2\beta + 1)^2\Delta$.

Theorem 4. Let \mathcal{D} be a set of β -fat regions of ply Δ in \mathbb{R}^2 . We can compute a path π between any two points $p, q \in \mathbb{R}^2$ whose resilience is at most $(1 + \varepsilon)r(p, q)$ in $O(2^{f(\Delta, \beta, \varepsilon)}n^5)$ time, where $f(\Delta, \beta, \varepsilon) = O\left(\frac{\Delta^4\beta^8}{\varepsilon^4} \log(\beta\Delta/\varepsilon)\right)$.

5. NP-hardness

In this section we show that computing the resilience of certain types of fat regions is NP-hard. We recall that several NP-hardness results for other shapes are already known, but most of them are for skinny objects. For example, hardness for the case in which regions are line segments in \mathbb{R}^2 was shown in [2,23] and [26, Section 5.1]. Our contribution is to show that hardness holds for the case in which ranges have bounded fatness (i.e., ranges are not skinny). The only hardness proof that we know for objects of positive area is by Tseng [22], who shows that if the regions are rotations and translations of a fixed square or ellipse the problem is NP-hard.

In addition to showing that the problem is difficult for other shapes, our construction is of independent interest, since it is completely different from those given in [2], [22], [23], and [26, Section 5.1]. Moreover, our proof has the advantage of being easy to extend to other shapes. We also note that the construction of Tseng uses several rotations of a fixed shape (i.e., 3 for a square, 4 for an ellipse), whereas our construction only needs two different rotations of the same shape.

First we show NP-hardness for general connected regions, and later we extend it to axis-aligned rectangles of aspect ratio $1 : 1 + \varepsilon$ and $1 + \varepsilon : 1$. We start the section establishing some useful graph-theoretical results.

Let G be a graph, and let p be a point in the plane. Let Γ be an embedding of G into the plane, which behaves properly (vertices go to distinct points, edges are curves that do not meet vertices other than their endpoints and do not triple cross), and such that p is not on a vertex or edge of the embedding. We say Γ is an *odd embedding* around p if it has the following property: every cycle of G has odd length if and only if the winding number of the corresponding closed curve in the plane in Γ around p is odd. We say a graph G is *oddly embeddable* if there exists an odd embedding Γ for it (Fig. 13 shows some examples). We claim that vertex cover is NP-hard for this constrained class of graphs. The proof of this statement is based on two observations.

Observation 1. Every tripartite graph is oddly embeddable.

Proof. The vertices of a tripartite graph G can be divided into three groups V_1, V_2, V_3 such that there are no internal edges in any of these groups. Now, consider a triangle Δ around p . We create an embedding Γ where all vertices in V_1 are close to one corner of Δ , the vertices in V_2 are close to a second corner, and the vertices in V_3 are close to the remaining corner. All edges are straight line segments. See Fig. 13(c).

Consider the graph H obtained from G by contracting all vertices in V_i to a single vertex v_i ; H is a triangle (or a subgraph of a triangle). Now consider any cycle in G , and project it to H . Since there were no edges in G connecting

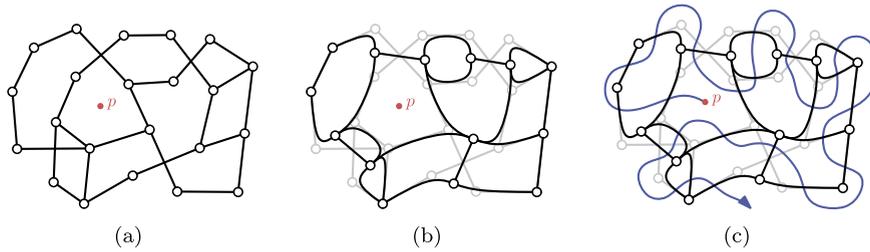


Fig. 14. (a) An oddly embedded graph with four crossings. (b) The crossings are flattened according to the parity of their vertices. (c) An odd Euler path from p to the outer face.

vertices within a group V_i , this does not change the length of the cycle, nor does it change the winding number around p . Any two consecutive edges from v_i to v_j , and back from v_j to v_i , do not influence the parity of the length of the cycle, nor the winding number around p , so we can remove them from the cycle. We are left with a cycle of length $3w$ and winding number w or $-w$, for some integer w . Clearly, $3w$ is odd if and only if w is odd. Therefore, Γ is an odd embedding of G , as required. \square

The *maximum independent set* problem in a graph asks for the largest set of vertices in the graph such that no two vertices in the set are connected by an edge. This problem is well-known to be NP-hard on general graphs. In fact, it remains NP-hard for tripartite graphs. A simple proof is included for completeness, and because we need the argument later. Note that a minimum vertex cover is the complement of a maximum independent set, hence by proving the NP-hardness of maximum independent set, we are also proving that minimum vertex cover is NP hard.

Observation 2. Let $G = (V, E)$ be a graph, and G' be the tripartite graph obtained from G by subdividing every edge $e \in E$ into an odd number of pieces m_e . Let $m = \sum_e m_e$ be the total number of vertices added. Let I and I' be maximum independent sets of G and G' , respectively. Then, it holds that $|I'| = |I| + m/2$.

Proof. For every independent set $I \subseteq V$ in G , there is a corresponding independent set I' in G' with $|I'| = |I| + m/2$: for every pair of extra vertices on an edge, we can always add one of the two to an independent set. Conversely, for every independent set I' in G' , there is a corresponding independent set I in G with $|I| = |I'| - m/2$: I' cannot use both extra vertices on an edge, so if we simply remove all extra vertices we remove at most $|E|$ elements from I' (clearly, if we remove less than $m/2$ vertices from I' this way, we can remove more vertices until I has the desired cardinality). \square

From the above observations, it follows that maximum independent set is also NP-hard on tripartite graphs, and hence, also on oddly embeddable graphs. Since our construction does not increase the maximum vertex degree, and vertex cover is known to be NP-hard for graphs with maximum degree three, we obtain the following.

Corollary 3. Minimum vertex cover on oddly embeddable graphs of maximum degree 3 is NP-hard.

Given an embedded graph Γ , we say that a curve in the plane is an *odd Euler path* if it does not go through any vertex of Γ and it crosses every edge of Γ an odd number of times.

Lemma 17. Let p be a point in the plane, and Γ an oddly embedded graph around p . Then there exists an odd Euler path for Γ that starts at p and ends in the outer face. Moreover, such path can be computed in polynomial time.

Proof. First, we insert an even number of extra vertices on every edge of Γ such that in the resulting embedded graph Γ' , every edge crosses at most one other edge. Now we construct an Euler path that crosses every edge of Γ' exactly once; note that this path will therefore cross every edge of Γ an odd number of times. Consider a pair of crossing edges and the four vertices concerned. For each pair of consecutive vertices (vertices that are not endpoints of the same edge), find a path in the graph that does not go around p (when seen as a cycle, after adding the crossing).

The parity of the length of this path does not depend on which path we take: if there would be an even-length path and an odd-length path between the same two vertices, both of which do not go around p , then there would be an odd cycle that does not contain p , which contradicts the oddly embeddedness of Γ' . Now, if the path has even length, we identify these two vertices. Note that of the four pairs of vertices involved in a crossing (i.e., ignoring the two pairs forming edges in Γ), exactly two pairs will have odd length connecting paths, so effectively we “flatten” the crossing. We do this for all crossings, and call the resulting multigraph Γ'' . See Fig. 14(b). (If the two crossing edges belong to different connected components of Γ , there are no paths connecting their vertices; in this case we make an arbitrary choice of which vertices to identify.)

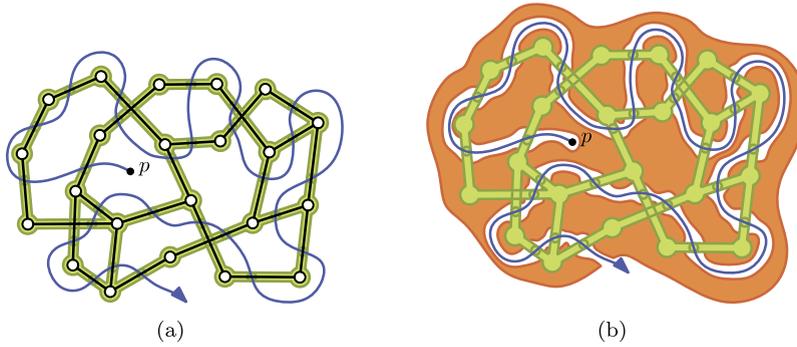


Fig. 15. Creating regions to follow Γ and T .

Now Γ'' is planar. Furthermore, by construction, all faces of Γ'' have even length, except the one containing p and the outer face. Therefore, the dual multigraph of Γ'' has only two vertices of odd degree, and hence has an Euler path between these vertices. Furthermore, this Euler path crosses every edge of Γ' exactly once, and therefore every edge of Γ an odd number of times. Note that the proof is constructive. Moreover, both the transformations and the Euler path can be done in polynomial time, hence such path can also be obtained in polynomial time. \square

Lemma 18. *Let p be a given point in the plane, and Γ an oddly embedded graph (not necessarily planar) around p . Furthermore, let T be a curve that forms an odd Euler path from p to the outer face. Then we can construct a set \mathcal{D} of connected regions such that a minimum set of regions from \mathcal{D} to remove corresponds exactly to a minimum vertex cover in Γ .*

Proof. If T is self-intersecting, then we can rearrange the pieces between self-intersections to remove all self-intersections. Thus we assume that T is a simple path.

If T crosses any edge of Γ more than once, we insert an even number of extra vertices on that edge such that afterwards, every edge is crossed exactly once. Let Γ' be the resulting graph. Since we inserted an even number of vertices on every edge, finding a minimum vertex cover in Γ' will give us a minimum vertex cover in Γ .

Now, for each vertex v in Γ' , we create one region D_v in \mathcal{D} . This region consists of the point where v is embedded, and the pieces of the edges adjacent to v up to the point where they cross T . Fig. 15(a) shows an example (the regions have been dilated by a small amount for visibility; if the embedding Γ has enough room this does not interfere with the construction). Note that all regions are simply connected.

Finally, we create one more special region W in \mathcal{D} that forms a corridor for T . Then W is duplicated at least n times to ensure that crossing this "wall" will always be more expensive than any other solution. Fig. 15(b) shows this.

Now, in order to escape, anyone starting at p must roughly follow T in order to not cross the wall. This means that for every edge of Γ' that T passes, one of the regions blocking the path (one of the vertices incident to the edge) must be disabled. The smallest number of regions to disable to achieve this corresponds to a minimum vertex cover in Γ' . \square

Combining this result with Corollary 3, we obtain our first hardness result for the barrier resilience problem.

Theorem 5. *The barrier resilience problem for a collection of connected regions is NP-hard.*

5.1. Extension to fat regions

We now adapt the previous approach to also work for a much more restricted class of regions: axis-aligned rectangles of sizes $1 \times (1 + \varepsilon)$ and $(1 + \varepsilon) \times 1$ for any $\varepsilon > 0$ (as long as ε depends polynomially on n). For simplicity, we limit Γ to have maximum degree 3. Maximum independent set is still known to be NP-hard in that case [10], and making them tripartite does not change the maximum degree.

The idea of the reduction is the following. We start from a sufficiently spacious (but polynomial) embedding of Γ , as illustrated in Fig. 16(a). On each edge we add a large even number of extra vertices. Each new vertex will be replaced by a rectangle, so every edge in Γ will become a chain of overlapping rectangles, like the green rectangles in Fig. 16(b). Therefore the first phase consists in replacing the embedding of Γ by an equivalent embedding of rectangles. We call these rectangles *graph rectangles* (green in the figures). Some care must be taken in the placement of graph rectangles around degree-3 vertices and in crossings, so that the rest of the construction can be made to work. Next, we place *wall rectangles* (orange in the figures; these consist of many copies of the same rectangle) across each graph rectangle. The gaps between adjacent wall rectangles should cover the overlapping part of two adjacent graph rectangles, so that a path can pass through them only whenever one of the two graph rectangles is removed. Then, we find a curve T from p that goes through every gap exactly once (note that T exists, by Lemma 17). Fig. 16(c) illustrates this phase of the construction. Finally, we add

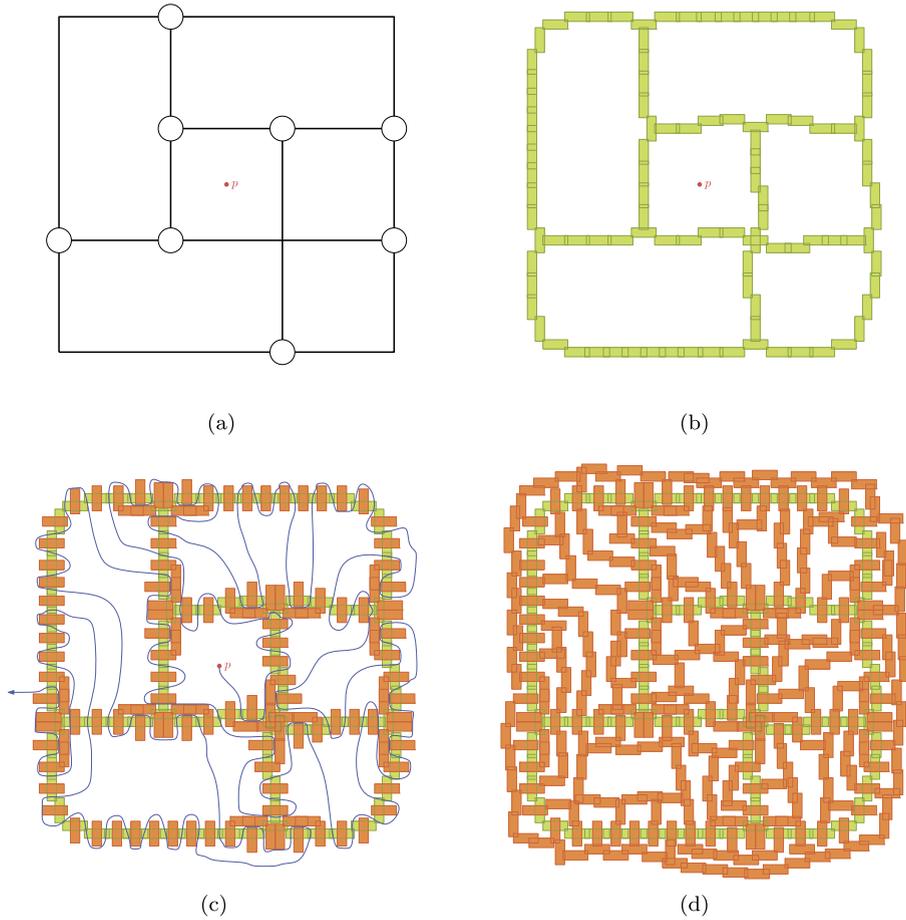


Fig. 16. (a) A non-planar oddly embedded cubic graph, embedded on a grid. (b) A set of rectangles, containing exactly one rectangle for each input vertex, and an even number of rectangles for each input edge. Note that crossings can be embedded if sufficiently far apart. (c) Local walls are added to make “tunnels”, each tunnel contains the overlapping part of two adjacent yellow rectangles. To go through a tunnel, one of the two yellow rectangles has to be removed. Then we choose an Euler path from p to the outside, that goes through each tunnel exactly once. (d) The final set of rectangles, designed to force any path from p to the outside to be homotopically equal to the one we drew.

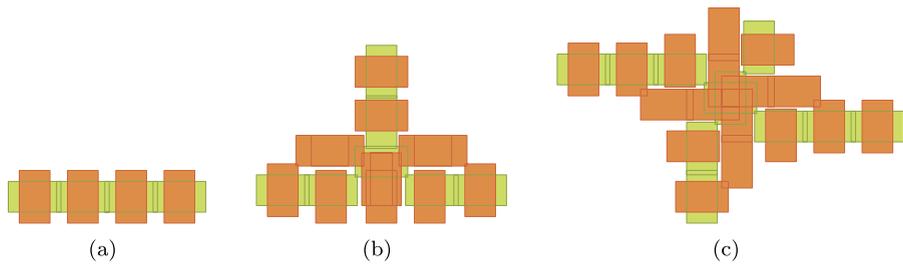


Fig. 17. Local details of the construction. Note that we use rectangles with a large aspect ratio for visibility, but the same constructions can be made with aspect ratio arbitrarily close to 1. (a) Overlapping rectangles to create edges (with an even number of extra vertices). (b) A vertex of degree at most 3, which is just a single rectangle. (c) A crossing between two chains.

more wall rectangles around T , to force any potential minimum resilience path from p that does not go through the wall rectangles to be homotopic to T . Fig. 16(d) shows the final set of rectangles. Now, computing an optimal resilience path among this set of rectangles would correspond to a maximal independent set in Γ .

For the construction to work, there needs to be enough space to place the wall rectangles. It is clear that this is possible far away from the graph rectangles, but close to the graph rectangles we proceed as follows: first, Fig. 17(a) shows the placement of rectangles along an edge of Γ . Fig. 17(b) shows how to place the rectangles at degree-3 vertices. Crossings are handled as shown in Fig. 17(c). These gadgets force some of the gaps in the chain to join each other. But this is no problem if every edge has enough rectangles. Also, note that at the center of the construction in Fig. 17(c) there are two overlapping

green rectangles, which belong to the two crossing chains. This is the only place where we vitally use the fact that the regions are not pseudodisks.

Lemma 19. *Let p be a given point in the plane, and Γ an oddly embedded graph with maximum vertex degree 3 (not necessarily planar) around p . Furthermore, let T be a curve that forms an odd Euler path from p to infinity. Then we can construct a set \mathcal{D} of axis-aligned rectangles of aspect ratio $1 : (1 + \varepsilon)$ such that a minimum set of regions from \mathcal{D} to remove corresponds exactly to a minimum vertex cover in Γ .*

Proof. We first add groups of extra vertices on every edge of Γ so that we have room to place the rectangles, in an even number per edge. Then replace edges by chains as of rectangles as in Fig. 17, and connect the orange (wall) rectangles to force the only optimal path from p to the outer face to be along the Euler path T . The path may have to be rerouted locally close to the crossings, but since there is a sufficiently large number of crossings with every edge anyway, this is always possible. Orange rectangles have to be duplicated sufficiently many times again, to make sure that no optimal path will ever cross them. \square

Theorem 6. *The barrier resilience problem for regions that are axis-aligned rectangles of aspect ratio $1 : (1 + \varepsilon)$ is NP-hard.*

A similar approach can likely be used to show NP-hardness of other classes of regions as well. However, it seems that a necessary property for our approach is that the regions are able to completely cross each other: in other words, the regions in \mathcal{D} cannot be pseudodisks.⁷

Acknowledgements

The authors would like to thank some anonymous referees for their thorough check of a previous version of this document. M.K was partially supported by KAKENHI projects number 17K12635, 15H02665, and 24106007. M.L. was supported by the Netherlands Organisation for Scientific Research (NWO) under grant 639.021.123. R.I.S. was partially supported by projects MINECO MTM2015-63791-R/FEDER and Gen. Cat. 2017SGR1640, and by MINECO through the Ramón y Cajal program.

References

- [1] P.K. Agarwal, J. Pach, M. Sharir, State of the union (of geometric objects), in: *Surveys on Discrete and Computational Geometry: Twenty Years Later*, in: *Contemp. Math.*, vol. 453, AMS, 2008.
- [2] H. Alt, S. Cabello, P. Giannopoulos, C. Knauer, On some connection problems in straight-line segment arrangements, in: *Proc. EuroCG*, 2011, pp. 27–30, also available as CoRR, arXiv:1104.4618.
- [3] S. Bereg, D.G. Kirkpatrick, Approximating barrier resilience in wireless sensor networks, in: *Proc. ALGOSENSORS*, 2009, pp. 29–40.
- [4] S. Cabello, P. Giannopoulos, The complexity of separating points in the plane, *Algorithmica* 74 (2) (2016) 643–663.
- [5] D.Y.C. Chan, D.G. Kirkpatrick, Multi-path algorithms for minimum-colour path problems with applications to approximating barrier resilience, *Theor. Comput. Sci.* 553 (2014) 74–90.
- [6] C.-Y. Chang, C.-Y. Hsiao, C.-T. Chang, The k-barrier coverage mechanism in wireless visual sensor networks, in: *Proc. IEEE WCNC*, 2012, pp. 2318–2322.
- [7] D.Z. Chen, Y. Gu, J. Li, H. Wang, Algorithms on minimizing the maximum sensor movement for barrier coverage of a linear domain, *Discrete Comput. Geom.* 50 (2) (2013) 374–408.
- [8] M. de Berg, Improved bounds on the union complexity of fat objects, *Discrete Comput. Geom.* 40 (1) (July 2008) 127–140.
- [9] A. Efrat, The complexity of the union of (α, β) -covered objects, *SIAM J. Comput.* 34 (2005) 775–787.
- [10] M. Garey, D. Johnson, L. Stockmeyer, Some simplified NP-complete graph problems, *Theor. Comput. Sci.* 1 (3) (1976) 237–267.
- [11] M. Gibson, G. Kanade, K. Varadarajan, On isolating points using disks, in: *Proc. ESA*, 2011, pp. 61–69.
- [12] S. He, J. Chen, X. Li, X. Shen, Y. Sun, Cost-effective barrier coverage by mobile sensor networks, in: *Proc. INFOCOM*, 2012, pp. 819–827.
- [13] T.C. Hu, Multi-commodity network flows, *Oper. Res.* 11 (3) (1963) 344–360.
- [14] M. Korman, M. Löffler, R.I. Silveira, D. Strash, On the complexity of barrier resilience for fat regions, in: *Proc. ALGOSENSORS*, 2013, pp. 201–216.
- [15] S. Kumar, T.-H. Lai, A. Arora, Barrier coverage with wireless sensors, in: *Proc. MOBICOM*, 2005, pp. 284–298.
- [16] S. Kumar, T.-H. Lai, A. Arora, Barrier coverage with wireless sensors, *Wirel. Netw.* 13 (6) (2007) 817–834.
- [17] D. Marx, Parameterized graph separation problems, *Theor. Comput. Sci.* 351 (3) (2006) 394–406.
- [18] J. Matousek, *Lectures on Discrete Geometry*, Springer-Verlag, New York, Inc., Secaucus, NJ, USA, 2002.
- [19] G.L. Miller, S. Teng, W.P. Thurston, S.A. Vavasis, Separators for sphere-packings and nearest neighbor graphs, *J. ACM* 44 (1) (1997) 1–29.
- [20] J. Pach, M. Sharir, On the boundary of the union of planar convex sets, *Discrete Comput. Geom.* 21 (3) (1999) 321–328.
- [21] R. Penninger, I. Vigan, Point set isolation using unit disks is NP-complete, CoRR, arXiv:1303.2779, 2013.
- [22] K.-C.R. Tseng, Resilience of Wireless Sensor Networks, Master's thesis, University of British Columbia, 2011.
- [23] K.-C.R. Tseng, D. Kirkpatrick, On barrier resilience of sensor networks, in: *Proc. ALGOSENSORS*, 2011, pp. 130–144.
- [24] S. Whitesides, R. Zhao, K-Admissible Collections of Jordan Curves and Offsets of Circular Arc Figures, Technical Report SOCS 90.08, School of Computer Science, McGill University, 1990.
- [25] M. Xiao, Simple and improved parameterized algorithms for multiterminal cuts, *Theory Comput. Syst.* 46 (4) (2010) 723–736.
- [26] S. Yang, Some Path Planning Algorithms in Computational Geometry and Air Traffic Management, PhD thesis, State University of New York, Stony Brook, 2012.

⁷ A similar fact was also observed in [23].