

Minimum Bisection is NP-hard on Unit Disk Graphs

Josep Díaz¹ and George B. Mertzios^{2*}

¹ Departament de Llenguatges i Sistemes Informàtics,
Universitat Politècnica de Catalunya, Spain.

² School of Engineering and Computing Sciences, Durham University, UK.
diaz@lsi.upc.edu, george.mertzios@durham.ac.uk

Abstract. In this paper we prove that the MIN-BISECTION problem is NP-hard on *unit disk graphs*, thus solving a longstanding open question.

Keywords: Minimum bisection problem, unit disk graphs, planar graphs, NP-hardness.

1 Introduction

The problem of appropriately partitioning the vertices of a given graph into subsets, such that certain conditions are fulfilled, is a fundamental algorithmic problem. Apart from their evident theoretical interest, graph partitioning problems have great practical relevance in a wide spectrum of applications, such as in computer vision, image processing, and VLSI layout design, among others, as they appear in many divide-and-conquer algorithms (for an overview see [2]). In particular, the problem of partitioning a graph into equal sized components, while minimizing the number of edges among the components turns out to be very important in parallel computing. For instance, to parallelize applications we usually need to evenly distribute the computational load to processors, while minimizing the communication between processors.

Given a simple graph $G = (V, E)$ and $k \geq 2$, a *balanced k -partition* of $G = (V, E)$ is a partition of V into k vertex sets V_1, V_2, \dots, V_k such that $|V_i| \leq \lceil \frac{|V|}{k} \rceil$ for every $i = 1, 2, \dots, k$. The *cut size* (or simply, the *size*) of a balanced k -partition is the number of edges of G with one endpoint in a set V_i and the other endpoint in a set V_j , where $i \neq j$. In particular, for $k = 2$, a balanced 2-partition of G is also termed a *bisection* of G . The *minimum bisection* problem (or simply, MIN-BISECTION) is the problem, given a graph G , to compute a bisection of G with the minimum possible size, also known as the *bisection width* of G .

Due to the practical importance of MIN-BISECTION, several heuristics and exact algorithms have been developed, which are quite efficient in practice [2], from the first ones in the 70's [16] up to the very efficient one described in [7]. However, from the theoretical viewpoint, MIN-BISECTION has been one of the

*Partially supported by the EPSRC Grant EP/K022660/1.

most intriguing problems in algorithmic graph theory so far. This problem is well known to be NP-hard for general graphs [11], while it remains NP-hard when restricted to the class of everywhere dense graphs [18] (i.e. graphs with minimum degree $\Omega(n)$), to the class of bounded maximum degree graphs [18], or to the class of d -regular graphs [5]. On the positive side, very recently it has been proved that MIN-BISECTION is fixed parameter tractable [6], while the currently best known approximation ratio is $O(\log n)$ [20]. Furthermore, it is known that MIN-BISECTION can be solved in polynomial time on trees and hypercubes [9, 18], on graphs with bounded treewidth [13], as well as on grid graphs with a constant number of holes [10, 19].

In spite of this, the complexity status of MIN-BISECTION on planar graphs, on grid graphs with an arbitrary number of holes, and on unit disk graphs have remained longstanding open problems so far [8, 10, 14, 15]. The first two of these problems are equivalent, as there exists a polynomial time reduction from planar graphs to grid graphs with holes [19]. Furthermore, there exists a polynomial time reduction from planar graphs with maximum degree 4 to unit disk graphs [8]. Therefore, since grid graphs with holes are planar graphs of maximum degree 4, there exists a polynomial reduction of MIN-BISECTION from planar graphs to unit disk graphs. Another motivation for studying MIN-BISECTION on unit disk graphs comes from the area of wireless communication networks [1, 3], as the bisection width determines the communication bandwidth of the network [12].

Our contribution. In this paper we resolve the complexity of MIN-BISECTION on unit disk graphs. In particular, we prove that this problem is NP-hard by providing a polynomial reduction from a variant of the maximum satisfiability problem, namely from the monotone Max-XOR(3) problem (also known as the monotone Max-2-XOR(3) problem). Consider a monotone XOR-boolean formula ϕ with variables x_1, x_2, \dots, x_n , i.e. a boolean formula that is the conjunction of XOR-clauses of the form $(x_i \oplus x_k)$, where no variable is negated. If, in addition, every variable x_i appears in exactly k XOR-clauses in ϕ , then ϕ is called a *monotone XOR(k)* formula. The *monotone Max-XOR(k)* problem is, given a monotone XOR(k) formula ϕ , to compute a truth assignment of the variables x_1, x_2, \dots, x_n that XOR-satisfies the largest possible number of clauses of ϕ . Recall here that the clause $(x_i \oplus x_k)$ is XOR-satisfied by a truth assignment τ if and only if $x_i \neq x_k$ in τ . Given a monotone XOR(k) formula ϕ , we construct a unit disk graph H_ϕ such that the truth assignments that XOR-satisfy the maximum number of clauses in ϕ correspond bijectively to the minimum bisections in H_ϕ , thus proving that MIN-BISECTION is NP-hard on unit disk graphs.

Organization of the paper. Necessary definitions and notation are given in Section 2. In Section 3, given a monotone XOR(3)-formula ϕ with n variables, we construct an auxiliary unit disk graph G_n , which depends only on the size n of ϕ (and not on ϕ itself). In Section 4 we present our reduction from the monotone Max-XOR(3) problem to MIN-BISECTION on unit disk graphs, by modifying the graph G_n to a unit disk graph H_ϕ which also depends on the formula ϕ

itself. Finally we discuss the presented results and remaining open problems in Section 5.

2 Preliminaries and Notation

We consider in this article simple undirected graphs with no loops or multiple edges. In an undirected graph $G = (V, E)$, the edge between vertices u and v is denoted by uv , and in this case u and v are said to be *adjacent* in G . For every vertex $u \in V$ the *neighborhood* of u is the set $N(u) = \{v \in V \mid uv \in E\}$ of its adjacent vertices and its *closed neighborhood* is $N[u] = N(u) \cup \{u\}$. The subgraph of G that is *induced* by the vertex subset $S \subseteq V$ is denoted $G[S]$. Furthermore a vertex subset $S \subseteq V$ induces a *clique* in G if $uv \in E$ for every pair $u, v \in S$.

A graph $G = (V, E)$ with n vertices is the *intersection graph* of a family $F = \{S_1, \dots, S_n\}$ of subsets of a set S if there exists a bijection $\mu : V \rightarrow F$ such that for any two distinct vertices $u, v \in V$, $uv \in E$ if and only if $\mu(u) \cap \mu(v) \neq \emptyset$. Then, F is called an *intersection model* of G . A graph G is a *disk graph* if G is the intersection graph of a set of disks (i.e. circles together with their internal area) in the plane. A disk graph G is a *unit disk graph* if there exists a disk intersection model for G where all disks have equal radius (without loss of generality, all their radii are equal to 1). Given a disk (resp. unit disk) graph G , an intersection model of G with disks (resp. unit disks) in the plane is called a *disk* (resp. *unit disk*) *representation* of G . Alternatively, unit disk graphs can be defined as the graphs that can be represented by a set of points on the plane (where every point corresponds to a vertex) such that two vertices intersect if and only if the corresponding points lie at a distance at most some fixed constant c (for example $c = 1$). Although these two definitions of unit disk graphs are equivalent, in this paper we use the representation with the unit disks instead of the representation with the points.

Note that any unit disk representation R of a unit disk graph $G = (V, E)$ can be completely described by specifying the centers c_v of the unit disks D_v , where $v \in V$, while for any disk representation we also need to specify the radius r_v of every disk D_v , $v \in V$. Given a graph G , it is NP-hard to decide whether G is a disk (resp. unit disk) graph [4, 17]. Given a unit disk representation R of a unit disk graph G , in the remainder of the paper we may not distinguish for simplicity between a vertex of G and the corresponding unit disk in R , whenever it is clear from the context. It is well known that the Max-XOR problem is NP-hard. Furthermore, it remains NP-hard even if the given formula ϕ is restricted to be a monotone XOR(3) formula. For the sake of completeness we provide in the next lemma a proof of this fact.

Lemma 1. *Monotone Max-XOR(3) is NP-hard.*

3 Construction of the unit disk graph G_n

In this section we present the construction of the auxiliary unit disk graph G_n , given a monotone XOR(3)-formula ϕ with n variables. Note that G_n depends only on the size of the formula ϕ and not on ϕ itself. Using this auxiliary graph G_n we will then construct in Section 4 the unit disk graph H_ϕ , which depends also on ϕ itself, completing thus the NP-hardness reduction from monotone Max-XOR(3) to the minimum bisection problem on unit disk graphs.

We define G_n by providing a unit disk representation R_n for it. For simplicity of the presentation of this construction, we first define a set of halflines on the plane, on which all centers of the disks are located in the representation R_n .

3.1 The half-lines containing the disk centers

Denote the variables of the formula ϕ by $\{x_1, x_2, \dots, x_n\}$. Define for simplicity the values $d_1 = 5.6$ and $d_2 = 7.2$. For every variable x_i , where $i \in \{1, 2, \dots, n\}$, we define the following four points in the plane:

- $p_{i,0} = (2i \cdot d_1, 2(i-1) \cdot d_2)$ and $p_{i,1} = ((2i-1) \cdot d_1, (2i-1) \cdot d_2)$, which are called the *bend points* for variable x_i , and
- $q_{i,0} = ((2i-1) \cdot d_1, 2(i-1) \cdot d_2)$ and $r_{i,0} = (2i \cdot d_1, 2i \cdot d_2)$, which is called the *auxiliary points* for variable x_i .

Then, starting from point $p_{i,j}$, where $i \in \{1, 2, \dots, n\}$ and $j \in \{0, 1\}$, we draw in the plane one halfline parallel to the x -axis pointing to the left and one halfline parallel to the y -axis pointing upwards. The union of these two halflines on the plane is called the *track* $T_{i,j}$ of point $p_{i,j}$. Note that, by definition of the points $p_{i,j}$, the tracks $T_{i,0}$ and $T_{i,1}$ do not have any common point, and that, whenever $i \neq k$, the tracks $T_{i,j}$ and $T_{k,\ell}$ have exactly one common point. Furthermore note that, for every $i \in \{1, 2, \dots, n\}$, both auxiliary points $q_{i,0}$ and $r_{i,0}$ belong to the track $T_{i,0}$.

We will construct the unit disk representation R_n of the graph G_n in such a way that the union of all tracks $T_{i,j}$ will contain the centers of all disks in R_n . The construction of R_n is done by repeatedly placing on the tracks $T_{i,j}$ multiple copies of three particular unit disk representations $Q_1(p)$, $Q_2(p)$, and $Q_3(p)$ (each of them including $2n^6 + 2$ unit disks), which we use as gadgets in our construction. Before we define these gadgets we need to define first the notion of a (t, p) -crowd.

Definition 1. *Let $\varepsilon > 0$ be infinitesimally small. Let $t \geq 1$ and $p = (x_p, y_p)$ be a point in the plane. Then, the horizontal (t, p) -crowd (resp. the vertical (t, p) -crowd) is a set of t unit disks whose centers are equally distributed between the points $(x_p - \varepsilon, y_p)$ and $(x_p + \varepsilon, y_p)$ (resp. between the points $(x_p, y_p - \varepsilon)$ and $(x_p, y_p + \varepsilon)$).*

Note that, by Definition 1, both the horizontal and the vertical (t, p) -crowds represent a clique of t vertices. Furthermore note that both the horizontal and

the vertical $(1, p)$ -crowds consist of a single unit disk centered at point p . For simplicity of the presentation, we will graphically depict in the following a (t, p) -crowd just by a disk with a *dashed contour* centered at point p , and having the number t written next to it. Furthermore, whenever the point p lies on the horizontal (resp. vertical) halfline of a track $T_{i,j}$, then any (t, p) -crowd will be meant to be a horizontal (resp. vertical) (t, p) -crowd.

3.2 Three useful gadgets

Let $p = (p_x, p_y)$ be a point on a track $T_{i,j}$. Whenever p lies on the horizontal halfline of $T_{i,j}$, we define for any $\delta > 0$ (with a slight abuse of notation) the points $p - \delta = (p_x - \delta, p_y)$ and $p + \delta = (p_x + \delta, p_y)$. Similarly, whenever p lies on the vertical halfline of $T_{i,j}$, we define for any $\delta > 0$ the points $p - \delta = (p_x, p_y - \delta)$ and $p + \delta = (p_x, p_y + \delta)$. Assume first that p lies on the *horizontal* halfline of $T_{i,j}$. Then we define the unit disk representation $Q_1(p)$ as follows:

- $Q_1(p)$ consists of the horizontal $(n^3, p + 0.9)$ -crowd, the horizontal $(2n^6 - 2n^3 + 2, p + 2.8)$ -crowd, and the horizontal $(n^3, p + 4.7)$ -crowd, as it is illustrated in Figure 1(a).

Assume now that p lies on the *vertical* halfline of $T_{i,j}$, we define the unit disk representations $Q_2(p)$ and $Q_3(p)$ as follows:

- $Q_2(p)$ consists of a single unit disk centered at point p , the vertical $(n^6, p + 1.8)$ -crowd, a single unit disk centered at point $p + 3.6$, and the vertical $(n^6, p + 5.4)$ -crowd, as it is illustrated in Figure 1(b).
- $Q_3(p)$ consists of a single unit disk centered at point p , the vertical $(n^6, p + 1.7)$ -crowd, a single unit disk centered at point $p + 3.6$, and the vertical $(n^6, p + 5.4)$ -crowd, as it is illustrated in Figure 1(c).

In the above definition of the unit disk representation $Q_k(p)$, where $k \in \{1, 2, 3\}$, the point p is called the *origin* of $Q_k(p)$. Note that the origin p of the representation $Q_2(p)$ (resp. $Q_3(p)$) is a center of a unit disk in $Q_2(p)$ (resp. $Q_3(p)$). In contrast, the origin p of the representation $Q_1(p)$ is not a center of any unit disk of $Q_1(p)$, however p lies in $Q_1(p)$ within the area of each of the n^3 unit disks of the horizontal $(n^3, p + 0.9)$ -crowd of $Q_1(p)$. For every point p , each of $Q_1(p)$, $Q_2(p)$, and $Q_3(p)$ has in total $2n^6 + 2$ unit disks (cf. Figure 1).

Furthermore, for any $i \in \{1, 2, 3\}$ and any two points p and p' in the plane, the unit disk representation $Q_i(p')$ is an isomorphic copy of the representation $Q_i(p)$, which is placed at the origin p' instead of the origin p . Moreover, for any point p in the vertical halfline of a track $T_{i,j}$, the unit disk representations $Q_2(p)$ and $Q_3(p)$ are almost identical: their only difference is that the vertical $(n^6, p + 1.8)$ -crowd in $Q_2(p)$ is replaced by the vertical $(n^6, p + 1.7)$ -crowd in $Q_3(p)$, i.e. this whole crowd is just moved downwards by 0.1 in $Q_3(p)$.

Observation 1 *Let $k \in \{1, 2, 3\}$ and $p \in T_{i,j}$, where $i \in \{1, 2, \dots, n\}$ and $j \in \{0, 1\}$. For every two adjacent vertices u, v in the unit disk graph defined by $Q_k(p)$, u and v belong to a clique of size at least $n^6 + 1$.*

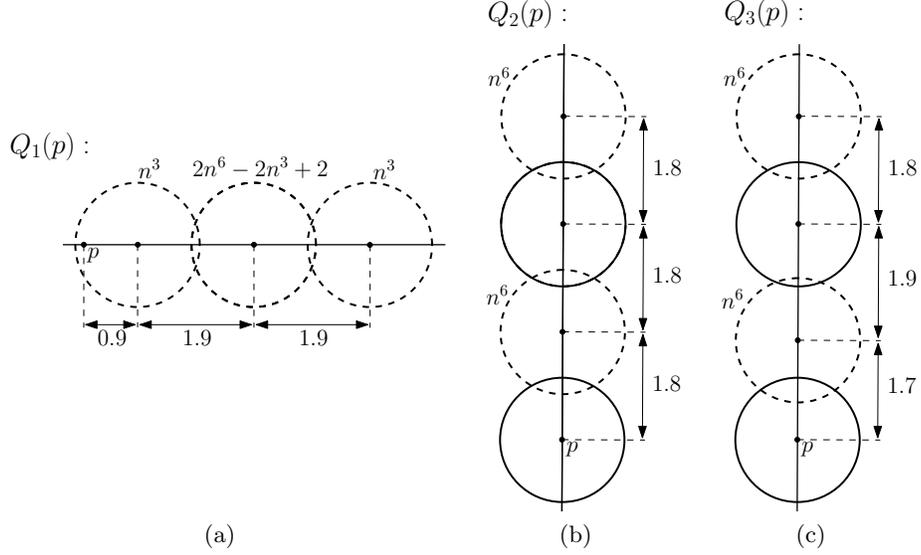


Fig. 1. The unit disk representations $Q_1(p)$, $Q_2(p)$, and $Q_3(p)$, where p is a point on one of the tracks $T_{i,j}$, where $1 \leq i \leq n$ and $j \in \{0, 1\}$.

3.3 The unit disk representation R_n of G_n

We are now ready to iteratively construct the unit disk representation R_n of the graph G_n , using the above gadgets $Q_1(p)$, $Q_2(p)$, and $Q_3(p)$, as follows:

- (a) for every $i \in \{1, 2, \dots, n\}$ and for every $j \in \{0, 1\}$, add to R_n :
 - the gadget $Q_1(p)$, with its origin at the point $p = (0, (2(i-1) + j) \cdot d_2)$,
- (b) for every $i \in \{1, 2, \dots, n\}$, add to R_n :
 - the gadgets $Q_1(q_{i,0})$, $Q_2(r_{i,0})$, $Q_3(p_{i,0})$, and $Q_3(p_{i,1})$,
 - the gadgets $Q_1(p)$ and $Q_1(p')$, with their origin at the points $p = (-d_1, (2i-1) \cdot d_2)$ and $p' = (-2d_1, (2i-1) \cdot d_2)$ of the track $T_{i,1}$, respectively,
- (c) for every $i, k \in \{1, 2, \dots, n\}$ and for every $j, \ell \in \{0, 1\}$, where $i \neq k$, add to R_n :
 - the gadgets $Q_1(p)$ and $Q_2(p)$, with their origin at the (unique) point p that lies on the intersection of the tracks $T_{i,j}$ and $T_{k,\ell}$.

This completes the construction of the unit disk representation R_n of the graph $G_n = (V_n, E_n)$, in which the centers of all unit disks lie on some track $T_{i,j}$, where $i \in \{1, 2, \dots, n\}$ and $j \in \{0, 1\}$.

Definition 2. Let $i \in \{1, 2, \dots, n\}$ and $j \in \{0, 1\}$. The vertex set $S_{i,j} \subseteq V_n$ consists of all vertices of those copies of the gadgets $Q_1(p)$, $Q_2(p)$, and $Q_3(p)$, whose origin p belongs to the track $T_{i,j}$.

For every $v \in V_n$ let c_v be the center of its unit disk in the representation R_n . Note that, by Definition 2, the unique vertex $v \in V_n$, for which $c_v \in T_{i,j} \cap T_{k,\ell}$, where $i < k$ (i.e. c_v lies on the intersection of the vertical halfline of $T_{i,j}$ with the horizontal halfline of $T_{k,\ell}$), we have that $v \in S_{i,j}$. Furthermore note that $\{S_{i,j} : 1 \leq i \leq n, j \in \{0,1\}\}$ is a partition of the vertex set V_n of G_n . In the next lemma we show that this is also a balanced $2n$ -partition of G_n , i.e. $|S_{i,j}| = |S_{k,\ell}|$ for every $i, k \in \{1, 2, \dots, n\}$ and $j, \ell \in \{0, 1\}$.

Lemma 2. *For every $i \in \{1, 2, \dots, n\}$ and $j \in \{0, 1\}$, we have that $|S_{i,j}| = 4(n+1)(n^6+1)$.*

Consider the intersection point p of two tracks $T_{i,j}$ and $T_{k,\ell}$, where $i \neq k$. Assume without loss of generality that $i < k$, i.e. p belongs to the vertical halfline of $T_{i,j}$ and on the horizontal halfline of $T_{k,\ell}$, cf. Figure 2(a). Then p is the origin of the gadget $Q_2(p)$ in the representation R_n (cf. part (c) of the construction of R_n). Therefore p is the center of a unit disk in R_n , i.e. $p = c_v$ for some $v \in S_{i,j} \subseteq V_n$. All unit disks of R_n that intersect with the disk centered at point p is shown in Figure 2(a). Furthermore, the induced subgraph $G_n[\{v\} \cup N(v)]$ on the vertices of G_n , which correspond to these disks of Figure 2(a), is shown in Figure 2(c). In Figure 2(c) we denote by K_{n^6} and K_{n^3} the cliques with n^6 and with n^3 vertices, respectively, and the thick edge connecting the two K_{n^3} 's depicts the fact that all vertices of the two K_{n^3} 's are adjacent to each other.

Now consider a bend point $p_{i,j}$ of a variable x_i , where $j \in \{0, 1\}$. Then $p_{i,j}$ is the origin of the gadget $Q_3(p_{i,j})$ in the representation R_n (cf. the first bullet of part (b) of the construction of R_n). Therefore $p_{i,j}$ is the center of a unit disk in R_n , i.e. $p = c_v$ for some $v \in S_{i,j} \subseteq V_n$. All unit disks of R_n that intersect with the disk centered at point $p_{i,j}$ are shown in Figure 2(b). Furthermore, the induced subgraph $G_n[\{v\} \cup N(v)]$ of G_n that corresponds to the disks of Figure 2(b), is shown in Figure 2(d). In both Figures 2(a) and 2(b), the area of the intersection of two crowds (i.e. disks with dashed contour) is shaded gray for better visibility.

Lemma 3. *Consider an arbitrary bisection \mathcal{B} of G_n with size strictly less than n^6 . Then for every set $S_{i,j}$, $i \in \{1, 2, \dots, n\}$ and $j \in \{0, 1\}$, all vertices of $S_{i,j}$ belong to the same color class of \mathcal{B} .*

4 Minimum bisection on unit disk graphs

In this section we provide our polynomial-time reduction from the monotone Max-XOR(3) problem to the minimum bisection problem on unit disk graphs. To this end, given a monotone XOR(3) formula ϕ with n variables and $m = \frac{3n}{2}$ clauses, we appropriately modify the auxiliary unit disk graph G_n of Section 3 to obtain the unit disk graph H_ϕ . Then we prove that the truth assignments that satisfy the maximum number of clauses in ϕ correspond bijectively to the minimum bisections in H_ϕ .

We construct the unit disk graph $H_\phi = (V_\phi, E_\phi)$ from $G_n = (V_n, E_n)$ as follows. Let $(x_i \oplus x_k)$ be a clause of ϕ , where $i < k$. Let p_0 (resp. p_1) be the

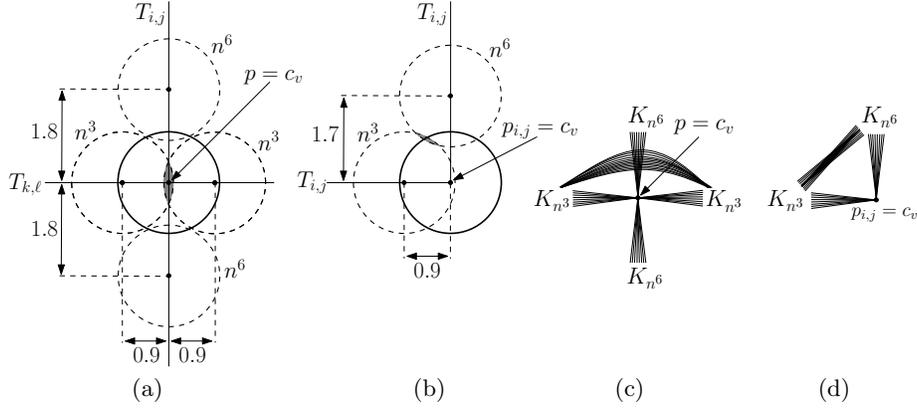


Fig. 2. The disks in R_n (a) around the intersection point $p = c_v$ of two tracks $T_{i,j}$ and $T_{k,\ell}$, where $i < k$, and (b) around the bend point $p_{i,j} = c_v$ of a variable x_i , where $j \in \{0, 1\}$. (c) The induced subgraph of G_n on the vertices of part (a), and (d) the induced subgraph of G_n for part (b).

unique point in the unit disk representation R_n that lies on the intersection of the tracks $T_{i,0}$ and $T_{k,1}$ (resp. on the intersection of the tracks $T_{i,1}$ and $T_{k,0}$). For every point $p \in \{p_0, p_1\}$, where we denote $p = (p_x, p_y)$, we modify the gadgets $Q_1(p)$ and $Q_2(p)$ in the representation R_n as follows:

- (a) replace the horizontal $(n^3, p + 0.9)$ -crowd of $Q_1(p)$ by the horizontal $(n^3 - 1, p + 0.9)$ -crowd and a single unit disk centered at $(p_x + 0.9, p_y + 0.02)$,
- (b) replace the vertical $(n^6, p + 1.8)$ -crowd of $Q_2(p)$ by the vertical $(n^6 - 1, p + 1.8)$ -crowd and a single unit disk centered at $(p_x + 0.02, p_y + 1.8)$.

That is, for every point $p \in \{p_0, p_1\}$, we first move one (arbitrary) unit disk of the horizontal $(n^3, p + 0.9)$ -crowd of $Q_1(p)$ upwards by 0.02, and then we move one (arbitrary) unit disk of the vertical $(n^6, p + 1.8)$ -crowd of $Q_2(p)$ to the right by 0.02. In the resulting unit disk representation these two unit disks intersect, whereas they do not intersect in the representation R_n . Furthermore it is easy to check that for any other pair of unit disks, these disks intersect in the resulting representation if and only if they intersect in R_n .

Denote by R_ϕ the unit disk representation that is obtained from R_n by performing the above modifications for all clauses of the formula ϕ . Then H_ϕ is the unit disk graph induced by R_ϕ . Note that, by construction, the graphs H_ϕ and G_n have exactly the same vertex set, i.e. $V_\phi = V_n$, and that $E_n \subset E_\phi$. In particular, note that the sets $S_{i,j}$ (cf. Definition 2) induce the same subgraphs in both H_ϕ and G_n , and thus the next corollary follows directly by Lemma 3.

Corollary 1. *Consider an arbitrary bisection \mathcal{B} of H_ϕ with size strictly less than n^6 . Then for every set $S_{i,j}$, $i \in \{1, 2, \dots, n\}$ and $j \in \{0, 1\}$, all vertices of $S_{i,j}$ belong to the same color class of \mathcal{B} .*

Theorem 1. *There exists a truth assignment τ of the formula ϕ that satisfies at least k clauses if and only if the unit disk graph H_ϕ has a bisection with value at most $2n^4(n-1) + 3n - 2k$.*

Proof (sketch). The (\Rightarrow) part of the proof is omitted due to lack of space.

(\Leftarrow) Assume that H_ϕ has a minimum bisection \mathcal{B} with value at most $2n^4(n-1) + 3n - 2k$. Denote the two color classes of \mathcal{B} by blue and red, respectively. Since the size of \mathcal{B} is strictly less than n^6 , Corollary 1 implies that for every $i \in \{1, 2, \dots, n\}$ and $j \in \{0, 1\}$, all vertices of the set $S_{i,j}$ belong to the same color class of \mathcal{B} . Therefore, all cut edges of \mathcal{B} have one endpoint in a set $S_{i,j}$ and the other endpoint in a set $S_{k,\ell}$, where $(i, j) \neq (k, \ell)$. Furthermore, since \mathcal{B} is a bisection of H_ϕ , Lemma 2 implies that exactly n of the sets $\{S_{i,j} : 1 \leq i \leq n, j \in \{0, 1\}\}$ are colored blue and the other n ones are colored red in \mathcal{B} .

First we will prove that, for every $i \in \{1, 2, \dots, n\}$, the sets $S_{i,0}$ and $S_{i,1}$ belong to different color classes in \mathcal{B} . To this end, let $t \geq 0$ be the number of variables x_i , $1 \leq i \leq n$, for which both sets $S_{i,0}$ and $S_{i,1}$ are colored blue (such variables x_i are called *blue*). Then, since \mathcal{B} is a bisection of H_ϕ , there must be also t variables x_i , $1 \leq i \leq n$, for which both sets $S_{i,0}$ and $S_{i,1}$ are colored red (such variables x_i are called *red*), whereas $n - 2t$ variables x_i , for which one of the sets $\{S_{i,0}, S_{i,1}\}$ is colored blue and the other one red (such variables x_i are called *balanced*). Using the minimality of the bisection \mathcal{B} , we will prove that $t = 0$.

Every cut edge of \mathcal{B} occurs at the intersection of the tracks of two variables x_i, x_k , where either both x_i, x_k are balanced variables, or one of them is a balanced and the other one is a blue or red variable, or one of them is a blue and the other one is a red variable. Furthermore recall by the construction of the graph H_ϕ from the graph G_n that every clause $(x_i \oplus x_k)$ of the formula ϕ corresponds to an intersection of the tracks of the variables x_i and x_k . Among the m clauses of ϕ , let m_1 of them correspond to intersections of tracks of two balanced variables, m_2 of them correspond to intersections of tracks of a balanced variable and a blue or red variable, and m_3 of them correspond to intersections of tracks of a blue variable and a red variable. Note that $m_1 + m_2 + m_3 \leq m$.

Let $1 \leq i < k \leq n$. In the following we distinguish the three cases of the variables x_i, x_k that can cause a cut edge in the bisection \mathcal{B} .

- **x_i and x_k are both balanced variables:** in total there are $\frac{(n-2t)(n-2t-1)}{2}$ such pairs of variables, where exactly m_1 of them correspond to a clause $(x_i \oplus x_k)$ of the formula ϕ . It is easy to check that, for every such pair x_i, x_k that does not correspond to a clause of ϕ , the intersection of the tracks of x_i and x_k contributes exactly $2n^3 + 2n^3 = 4n^3$ edges to the value of \mathcal{B} . Furthermore, for each of the m_1 other pairs x_i, x_k that correspond to a clause of ϕ , the intersection of the tracks of x_i and x_k contributes either $4n^3$ or $4n^3 + 2$ edges to the value of \mathcal{B} . In particular, if the vertices of the sets $S_{i,0}$ and $S_{k,1}$ have the same color in \mathcal{B} then the pair x_i, x_k contributes $4n^3$ edges to the value of \mathcal{B} , otherwise it contributes $4n^3 + 2$ edges. Among these m_1 clauses, let m_1^* of them contribute $4n^3$ edges each and the remaining $m_1 - m_1^*$ of them contribute $4n^3 + 2$ edges each.

- **one of x_i, x_k is a balanced variable and the other one is a blue or red variable:** in total there are $(n - 2t)2t$ such pairs of variables, where exactly m_2 of them correspond to a clause $(x_i \oplus x_k)$ of the formula ϕ . It is easy to check that, for every such pair x_i, x_k that does not correspond to a clause of ϕ , the intersection of the tracks of x_i and x_k contributes exactly $2n^3 + 2n^3 = 4n^3$ edges to the value of \mathcal{B} . Furthermore, for each of the m_2 other pairs x_i, x_k that correspond to a clause of ϕ , the intersection of the tracks of x_i and x_k contributes $4n^3 + 1$ edges to the value of \mathcal{B} .
- **one of x_i, x_k is a blue variable and the other one is a red variable:** in total there are t^2 such pairs of variables, where exactly m_3 of them correspond to a clause $(x_i \oplus x_k)$ of the formula ϕ . It is easy to check that, for every such pair x_i, x_k that does not correspond to a clause of ϕ , the intersection of the tracks of x_i and x_k contributes exactly $4 \cdot 2n^3 = 8n^3$ edges to the value of \mathcal{B} . Furthermore, for each of the m_3 other pairs x_i, x_k that correspond to a clause of ϕ , the intersection of the tracks of x_i and x_k contributes $8n^3 + 2$ edges to the value of \mathcal{B} .

Therefore, the value of \mathcal{B} can be computed (the exact details are omitted due to lack of space) as $2n^4(n - 1) + 4n^3t + 2(m_1 - m_1^*) + m_2 + 2m_3$. Note now that $0 \leq 2(m_1 - m_1^*) + m_2 + 2m_3 \leq 2m = 3n < 4n^3$. Therefore, since the value of the bisection \mathcal{B} is minimum by assumption, it follows that $t = 0$. Thus for every $i \in \{1, 2, \dots, n\}$ the variable x_i of ϕ is balanced in the bisection \mathcal{B} , i.e. the sets $S_{i,0}$ and $S_{i,1}$ belong to different color classes in \mathcal{B} . That is, $m_1 = m$ and $m_2 = m_3 = 0$, and thus the value of \mathcal{B} is equal to $2n^4(n - 1) + 2(m - m_1^*)$. On the other hand, since the value of \mathcal{B} is at most $2n^4(n - 1) + 3n - 2k$ by assumption, it follows that $2(m - m_1^*) \leq 3n - 2k$. Therefore, since $m = \frac{3n}{2}$, it follows that $m_1^* \geq k$.

We define now from \mathcal{B} the truth assignment τ of ϕ as follows. For every $i \in \{1, 2, \dots, n\}$, if the vertices of the set $S_{i,0}$ are blue and the vertices of the set $S_{i,1}$ are red in \mathcal{B} , then we set $x_i = 0$ in τ . Otherwise, if the vertices of the set $S_{i,0}$ are red and the vertices of the set $S_{i,1}$ are blue in \mathcal{B} , then we set $x_i = 1$ in τ . Recall that m_1^* is the number of clauses of ϕ that contribute $4n^3$ edges each to the value of \mathcal{B} , while the remaining $m - m_1^*$ clauses of ϕ contribute $4n^3 + 2$ edges each to the value of \mathcal{B} . Thus, by the construction of H_ϕ from G_n , for every clause $(x_i \oplus x_k)$ of ϕ that contributes $4n^3$ (resp. $4n^3 + 2$) to the value of \mathcal{B} , the vertices of the sets $S_{i,0}$ and $S_{k,1}$ have the same color (resp. $S_{i,0}$ and $S_{k,1}$ have different colors) in \mathcal{B} . Therefore, by definition of the truth assignment τ , there are exactly m_1^* clauses $(x_i \oplus x_k)$ of ϕ where $x_i \neq x_k$ in τ , and there are exactly $m - m_1^*$ clauses $(x_i \oplus x_k)$ of ϕ where $x_i = x_k$ in τ . That is, τ satisfies exactly $m_1^* \geq k$ of the m clauses of ϕ . This completes the proof of the theorem. \square

We can now state our main result, which follows by Theorem 1 and Lemma 1.

Theorem 2. MIN-BISECTION is NP-hard on unit disk graphs.

5 Concluding Remarks

In this paper we proved that MIN-BISECTION is NP-hard on unit disk graphs by providing a polynomial time reduction from the monotone Max-XOR(3) problem, thus solving a longstanding open question. As pointed out in the Introduction, our results indicate that MIN-BISECTION is probably also NP-hard on planar graphs, or equivalently on grid graphs with an arbitrary number of holes, which remains yet to be proved.

References

1. I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. Wireless sensor networks: A survey. *Computer Networks*, 38:393–422, 2002.
2. C.-E. Bichot and P. Siarry, editors. *Graph Partitioning*. Wiley, 2011.
3. M. Bradonjic, R. Elsässer, T. Friedrich, T. Sauerwald, and A. Stauffer. Efficient broadcast on random geometric graphs. In *Proceedings of the 21st annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1412–1421, 2010.
4. H. Brey and D. G. Kirkpatrick. Unit disk graph recognition is NP-hard. *Computational Geometry*, 9(1-2):3–24, 1998.
5. T. Bui, S. Chaudhuri, T. Leighton, and M. Sipser. Graph bisection algorithms with good average case behavior. *Combinatorica*, 7:1987, 171–191.
6. M. Cygan, D. Lokshtanov, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. Minimum bisection is fixed parameter tractable. In *Proceedings of the 46th Annual Symposium on the Theory of Computing (STOC)*, 2014. To appear.
7. D. Delling, A. V. Goldberg, I. Razenshteyn, and R. F. Werneck. Exact combinatorial branch-and-bound for graph bisection. In *Proceedings of the 14th Meeting on Algorithm Engineering & Experiments (ALENEX)*, pages 30–44, 2012.
8. J. Díaz, M. D. Penrose, J. Petit, and M. J. Serna. Approximating layout problems on random geometric graphs. *Journal of Algorithms*, 39(1):78–116, 2001.
9. J. Díaz, J. Petit, and M. Serna. A survey on graph layout problems. *ACM Computing Surveys*, 34:313–356, 2002.
10. A. E. Feldmann and P. Widmayer. An $O(n^4)$ time algorithm to compute the bisection width of solid grid graphs. In *Proceedings of the 19th annual European Symposium on Algorithms (ESA)*, pages 143–154, 2011.
11. M. R. Garey and D. S. Johnson. *Computers and intractability: A guide to the theory of NP-completeness*. W. H. Freeman & Co., 1979.
12. J. Hromkovic, R. Klasing, A. Pelc, P. Ruzicka, and W. Unger. *Dissemination of Information in Communication Networks - Broadcasting, Gossiping, Leader Election, and Fault-Tolerance*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2005.
13. K. Jansen, M. Karpinski, A. Lingas, and E. Seidel. Polynomial time approximation schemes for Max-Bisection on planar and geometric graphs. *SIAM Journal on Computing*, 35(1):110–119, 2005.
14. S. Kahruman-Anderoglu. *Optimization in geometric graphs: Complexity and approximation*. PhD thesis, Texas A & M University, 2009.
15. M. Karpinski. Approximability of the minimum bisection problem: An algorithmic challenge. In *Proceedings of the 27th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, pages 59–67, 2002.

16. B. Kernighan and S. Lin. An efficient heuristic procedure for partitioning graphs. *Bell System Technical Journal*, 49(2):291–307, 1970.
17. J. Kratochvíl. Intersection graphs of noncrossing arc-connected sets in the plane. In *Proceedings of the 4th Int. Symp. on Graph Drawing (GD)*, pages 257–270, 1996.
18. R. MacGregor. *On partitioning a graph: A theoretical and empirical study*. PhD thesis, University of California, Berkeley, 1978.
19. C. H. Papadimitriou and M. Sideri. The bisection width of grid graphs. *Mathematical Systems Theory*, 29(2):97–110, 1996.
20. H. Räcke. Optimal hierarchical decompositions for congestion minimization in networks. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 255–264, 2008.