Linear Time Approximation for Dominating Sets and Independent Dominating Sets in Unit Disk Graphs

Guilherme D. da Fonseca[1], Celina M. H. de Figueiredo[2],
Vinícius G. P. de Sá[3], and Raphael Machado[4]

[1] Dept. de Informática Aplicada, UniRio, Brazil
[2] PESC/COPPE, UFRJ, Brazil
[3] Dept. de Ciência da Computação, UFRJ, Brazil
[4] Instituto Nacional de Metrologia, Qualidade e Tecnologia, Brazil

Abstract. A unit disk graph is the intersection graph of n congruent disks in the plane. Dominating sets in unit disk graphs are widely studied due to their application in wireless ad-hoc networks. Since the minimum dominating set problem for unit disk graphs is NP-hard, several approximation algorithms with different merits have been proposed in the literature. On one extreme, there is a linear time 5-approximation algorithm. On another extreme, there are two PTAS whose running times are polynomials of very high degree. We introduce a linear time approximation algorithm that takes the usual adjacency representation of the graph as input and attains a 44/9 approximation factor. This approximation factor is also attained by a second algorithm we present, which takes the geometric representation of the graph as input and runs in $O(n \log n)$ time, regardless of the number of edges. The analysis of the approximation factor of the algorithms, both of which are based on local improvements, exploits an assortment of results from discrete geometry to prove that certain graphs cannot be unit disk graphs. It is noteworthy that the dominating sets obtained by our algorithms are also independent sets.

1 Introduction

A *unit disk graph* G is a graph whose n vertices can be mapped to points in the plane and whose m edges are defined by pairs of points within Euclidean distance at most 1 from one another. Alternatively, one can regard the vertices of G as mapped to coplanar disks of unit diameter, so that two vertices are adjacent whenever the corresponding disks intersect.

A *dominating set* D is a subset of the vertices of a graph such that every vertex not in D is adjacent to some vertex in D. An *independent dominating set* is a dominating set which is also an independent set. Note that any maximal independent set is an independent dominating set.

Dominating sets in unit disk graphs are widely studied due to their application in wireless ad-hoc networks [12]. Since it is NP-hard to compute the minimum dominating set of a unit disk graph [2], a number of approximation algorithms have been proposed [3, 5, 9, 10, 12, 15, 19]. Such algorithms are of two main types. *Graph-based algorithms* receive as input the adjacency representation of the graph and assume no knowledge of the point coordinates, whereas *geometric algorithms* work in the Real RAM model of computation and receive solely the vertex coordinates as input[5]. Thus far these two types of algorithms have been tackled separately in the literature for the dominating set problem in unit disk graphs. In this paper, we introduce approximation algorithms of both types, benefiting from the same approximation factor analysis.

Previous algorithms. A graph-based 5-approximation algorithm that runs in $O(n + m)$ time was presented in [12]. The algorithm computes a maximal independent set, which turns out to be a 5-approximation because unit disk graphs contain no $K_{1,6}$ as induced subgraphs.

Polynomial-time approximation schemes (PTAS) were first presented as geometric algorithms [10] and later as graph-based algorithms [15]. Also, a graph-based PTAS for the more general disk graphs is proposed in [9]. Unfortunately, the complexities of the existing PTAS are high degree polynomials. For example, the PTAS presented in [15] takes $O(n^{225})$ time to obtain a 5-approximation (using the analysis from [3]). Although its analysis is not tight, the running time is too high even for moderately large graphs. The reason is that these PTAS invoke a subroutine that verifies (by brute force) whether a graph admits a dominating set with k vertices. Such subroutine is applied to several subgraphs, and the value of k grows as the approximation error decreases. A similar strategy is used to obtain a PTAS for the minimum independent dominating set [11].

The lack of fast algorithms with approximation factors less than 5 was noticed in [3], where geometric algorithms with approximation factors of 3 and 4 and running times respectively $O(n^{18})$ and $O(n^9)$ were presented. While a significant step towards approximating large instances, those algorithms require the geometric representation of the graph, and the running times are polynomials of rather high degrees. Linear and near-linear time approximation algorithms constitute an active topic of research, even for problems that can be solved exactly in polynomial time, such as maximum flow and maximum matching [1, 18].

It is useful to contrast the minimum dominating set problem with the maximum independent set problem. While a maximal independent set is a 5-approximation to both problems, it is easy to obtain a geometric 3-approximation to the maximum independent set problem in $O(n \log n)$ time [14]. In the graph-based version, a related strategy takes roughly $O(n^5)$ time, though. No similar results are known for the minimum dominating set problem.

Packing problems usually arise in situations where one wants to enclose non-overlapping objects as densely as possible into recipients of given shape. Unit disk graphs are subject to a number of packing constraints that limit the size of independent sets (which correspond to disjoint disks) as a function of the distance

[5] The Real RAM model is a technical necessity, otherwise storing the coordinates of the vertices would require an exponential number of bits [13].

between vertices. The existing PTAS for dominating sets in unit disk graphs are based on some of these packing constraints, such as the *bounded growth property*: the size of an independent set formed by vertices within distance r of a given vertex is at most $(1 + 2r)^2$. Note, however, that the bounded growth property is not tight. For example, for $r = 1$, it gives an upper bound of 9 vertices where the actual maximum size is 5. Since the bounded growth property is strongly connected to the problem of packing circles in a circle [7], obtaining exact values for all r seems unlikely.

Our contribution. Our main result consists of two approximation algorithms: a graph-based algorithm, which runs in linear $O(n + m)$ time, and its geometric counterpart, which runs in $O(n \log n)$ time in the Real RAM model, regardless of the number of edges.

The approximation factor of our algorithms is 44/9. The strategy for both algorithms is to construct a 5-approximate solution using the algorithm from [12] and to perform subsequent local improvements to that initial dominating set. Our main lemma (Lemma 7) uses forbidden subgraphs to show that a solution that admits no local improvement is a 44/9-approximation. Since the dominating sets produced by our algorithms are independent sets, the same approximation factor holds for the independent dominating set problem.

Proving that a certain graph is *not* a unit disk graph (and is therefore a forbidden induced subgraph) is no easy feat[6]. We make use of an assortment of results from discrete geometry in order to prove properties of unit disk graphs that are interesting *per se*. For example, we use universal covers and disk packings to show that the neighborhood of a clique in a unit disk graph contains at most 12 independent vertices. These properties, along with a tighter version of the bounded growth property, allow us to show that certain graphs are not unit disk graphs. Consequently, our algorithms employ a broader set of forbidden subgraphs, including, but not being limited to, the $K_{1,6}$.

2 Forbidden Subgraphs

In this section, we introduce some lemmas about the structure of unit disk graphs. These lemmas will be applied to prove our approximation factor in Section 3. We start by stating three previous results from the area of discrete geometry. The first lemma comes from the study of universal covers (for a recent survey see [8]).

Lemma 1 (Pál [16]). *If a set of points P has diameter 1, then P can be enclosed by a circle of radius $1/\sqrt{3}$.*

Packing congruent disks in a circle is a well studied problem. Exact bounds on the radius of the smallest circle packing k congruent disks are known for some

[6] The fastest known algorithm to decide whether a given graph is a unit disk graph is doubly exponential [17].

small values of k, namely $k \leq 13$ and $k = 19$ [7]. The bound for $k = 13$ will be useful to us.

Lemma 2 (Fodor [7]). *The radius of the smallest circle enclosing* 13 *points with mutual distances at least* 1 *is* $(1+\sqrt{5})/2$.

The *density* of a packing is the ratio between the covered area and the total area. The following upper bound is useful when no exact bound is known.

Lemma 3 (Fejes Tóth [6]). *Every packing of two or more congruent disks in a convex region has density at most* $\pi/\sqrt{12}$.

Given a graph $G = (V, E)$ and a vertex $v \in V$, let $N(v)$ denote the *open neighborhood* of v and let $N[v] = N(v) \cup \{v\}$ denote the *closed neighborhood* of v. More generally, the *open r-neighborhood* of a vertex v is the set of vertices w such that the distance between v and w in G is exactly r, while the *closed r-neighborhood* of a vertex v is the set of vertices w such that the distance between v and w in G is at most r. For a set $S \subseteq V$, we let $N_S(v) = N(v) \cap S$ and $N_S[v] = N[v] \cap S$. Finally, given a subgraph G' of G, the *closed neighborhood* of G' is the set of vertices that belong to the closed neighborhood of some vertex of G'. The following two lemmas concern neighborhoods in unit disk graphs.

Lemma 4. *The closed neighborhood of a clique in a unit disk graph contains at most* 12 *independent vertices.*

Proof. By Lemma 1, the points which define a clique in a unit disk graph are contained inside a circle of radius $1/\sqrt{3}$. Therefore, the points corresponding to the closed neighborhood of such clique are contained inside a circle of radius $1 + (1/\sqrt{3})$. By Lemma 2, we have that a circle enclosing 13 points with mutual distances at least 1 has radius at least $(1+\sqrt{5})/2$. Since $(1+\sqrt{5})/2 > 1+(1/\sqrt{3})$, the lemma follows. □

Lemma 5. *Given an integer $r \geq 1$, the closed r-neighborhood of a vertex in a unit disk graph contains at most* $\lfloor \pi(2r+1)^2/\sqrt{12} \rfloor$ *independent vertices.*

Proof. All n disks of diameter 1 corresponding to the closed r-neighborhood of a vertex v must be enclosed by a circle W of radius $(2r+1)/2$ centered on v. Each disk of diameter 1 has area $\pi/4$ and W has area $(2r+1)^2\pi/4$. Using Lemma 3, we have $(n\,\pi/4)/((2r+1)^2\pi/4) \leq \pi/\sqrt{12}$, and the lemma follows. □

We say that a graph G is *(k, ℓ)-pendant* if there is a vertex v in G with k vertices of degree 1 in the open neighborhood of v and ℓ vertices of degree 1 in the open 2-neighborhood of v. We refer to v as a *generator* of the (k, ℓ)-pendant graph. The following lemma bounds the value of the parameter ℓ for a $(4, \ell)$-pendant unit disk graphs.

Lemma 6. *If G is a $(4, \ell)$-pendant unit disk graph, then $\ell \leq 8$.*

Proof. Let v be a generator of G. Since $K_{1,6}$ is a forbidden induced subgraph [12] and v has 4 neighbors of degree 1, we have that the remaining neighbors of v together with v itself form a clique. By Lemma 4, we have that $4 + \ell \leq 12$. □

3 Approximation Algorithms

In this section, we present our approximation algorithms. The key property to analyze the approximation factor is presented in Lemma 7, while the running time analyses are presented in Sections 3.1 and 3.2.

Hereafter, let $G = (V, E)$ be a unit disk graph, and let $D \subseteq V$ be an independent dominating set of G. If $v \in D$ and $uv \in E$, we say that v *dominates* u and, conversely, that u *is dominated* by v.

As already mentioned, unit disk graphs are free of induced $K_{1,6}$. Therefore, at most 5 vertices of D may belong to the closed neighborhood of any given vertex $v \in V$. A *corona* is a set $C \subseteq D$ consisting of exactly 5 neighbors of some vertex $c \in V \setminus D$. Such a vertex c is called a *core* of the corona C, and it is not necessarily unique. Notice that the subgraph induced by a corona C and a corresponding core c is a star, i.e. a graph formed by an independent set and a universal vertex.

A corona C is said to be *reducible* if there is a core c of C such that $D \cup \{c\} \setminus C$ is a dominating set. If no such core exists, C is dubbed *irreducible*. Given a reducible corona C and a corresponding core c, we refer to the operation that converts D into the smaller dominating set $D \cup \{c\} \setminus C$ as a *reduction*.

Lemma 7. *Let $G = (V, E)$ be a unit disk graph, D an independent dominating set in G, and D^* a minimum dominating set of G. If D contains no reducible coronas, then $\rho = |D|/|D^*| \leq 44/9$.*

Proof. We use a charging argument to bound the ratio between the cardinalities of D and D^*. Consider that each vertex $u \in D$ splits a *unit charge* evenly among the vertices in the closed neighborhood $N_{D^*}[u]$. The function $f : D^* \to (0, 5]$ below corresponds to the total charges assigned to each vertex $v^* \in D^*$, accumulating the (fractional) charges that v^* received from the vertices in $N_D[v^*]$:

$$f(v^*) = \sum_{u \in N_D[v^*]} \frac{1}{|N_{D^*}[u]|}. \tag{1}$$

Note that, since D and D^* are dominating sets, neither $N_{D^*}[u]$ nor $N_D[v^*]$ are ever empty, and $f(v^*) \leq N_D[v^*]$. Such function f allows us to write the cardinality of D as

$$|D| = \sum_{v^* \in D^*} f(v^*).$$

Since

$$\rho = \frac{|D|}{|D^*|} = \frac{\sum_{v^* \in D^*} f(v^*)}{|D^*|}$$

is precisely the average value of $f(\cdot)$ over the elements of D^*, we obtain the desired bound $\rho \leq 44/9$ by showing that the existence of vertices c^* in D^* with $f(c^*) > 44/9$ is counterbalanced by a sufficiently large number of vertices x^* in D^* with $f(x^*) \leq 4$.

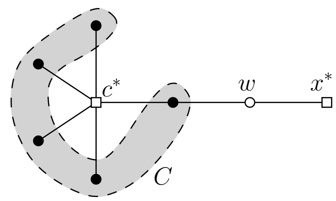

Fig. 1. Figure for the proof of Lemma 7. A proper, induced subgraph, where squares were used for a subset of D^*, solid circles for a subset of D (the corona C) and hollow circles for vertices not in $D \cup D^*$. Vertices w and x^* are respectively witness and reliever of core c^*.

Before we continue, we observe that $f(c^*) > 44/9$ means exactly $f(c^*) = 5$, because the sum in (1) has at most 5 terms, all of which are of the form $1/i$ for integer $i \geq 1$.

Thus, let c^* be a vertex in D^* with $f(c^*) = 5$. Clearly, $c^* \notin D$, otherwise $f(c^*) \leq |N_D[c^*]| = 1$, because D is an independent set. Moreover, c^* must have exactly 5 neighbors in D, since a greater number of neighbors in D would imply the existence of an induced $K_{1,6}$ in G, which is not possible, and a lesser number would imply $f(c^*) \leq |N_D[c^*]| \leq 4$, a contradiction. Vertex c^* is therefore a core.

Now let $C \subset D$ be the corona of which c^* is a core. Since C is irreducible (by the hypothesis of the lemma), there must be a vertex $w \in V \setminus (C \cup \{c^*\})$, such that:

(i) w is only dominated, in D, by vertices that belong to C; and
(ii) w is not adjacent to c^*.

We call w a *witness* of c^* (meaning the corona having c^* as a core fails to be reducible due to w). Now, for all $u \in C$, it holds that the only vertex in $N_{D^*}[u]$ must be the very core c^*, otherwise the contribution of u in (1) would be at most $1/2$, and $f(c^*)$ would be at most $9/2 < 5$, a contradiction. In particular, the witness w cannot belong to D^*. But D^* is a dominating set, so there must exist a vertex $x^* \in D^*$ that is adjacent to w. We call x^* a *reliever* of c^*. Figure 1 illustrates this situation.

We now show that $|N_D[x^*]| \leq 4$. For the sake of contradiction, assume $|N_D[x^*]| > 4$. Because G is free of induced $K_{1,6}$, such number must be exactly 5, so that x^* is the core of a corona $C' \subset D$. However, due to (i) above, $N_{C'}(w) = \emptyset$, hence $C' \cup \{w\}$ is an independent set of G, constituting, along with the core x^*, an induced $K_{1,6}$ in G, a contradiction. Since $f(x^*) \leq |N_D[x^*]|$, we have $f(x^*) \leq 4$.

We have just shown that the existence of a vertex c^* in D^* with $f(c^*) = 5$ implies the existence of a vertex $x^* \in D^*$ such that $f(x^*) \leq 4$. Were this correspondence one-to-one, we would be able to state that the average of $f(\cdot)$ over the elements of D^* was no greater than 4.5. Unfortunately, this correspondence is not necessarily one-to-one, as illustrated in Figure 2.

Still, the lemmas in Section 2 allow us to bound the ratio between the number of vertices c^* with $f(c^*) = 5$ and the number of vertices x^* for which the values

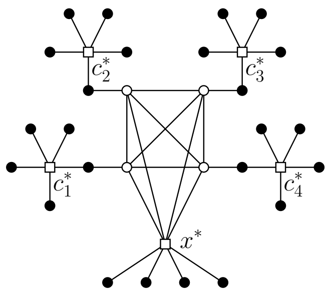

Fig. 2. A unit disk graph where 4 distinct cores $c_1^*, \ldots, c_4^*$ share the same reliever x^*.

of f are significantly lower. Let $x^* \in D^* \setminus D$ be a reliever. In order to obtain the claimed bound, we consider two cases according to the size of $N_D[x^*]$:

(i) $|\boldsymbol{N_D[x^*]}| \leq \mathbf{3}$. By Lemma 5, the closed 4-neighborhood of x^* contains at most 73 independent vertices. Since each corona contains 5 independent vertices (only adjacent to their cores), at most $\lfloor 73/5 \rfloor = 14$ coronas may share a common reliever[7]. Let $c_1^*, \ldots, c_{14}^*$ denote the cores of such coronas. If $|N_D[x^*]| \leq 3$, then the average value of $f(\cdot)$ among $x^*, c_1^*, \ldots, c_{14}^*$ is at most

$$\frac{3 + 14 \cdot 5}{15} < 4.867.$$

(ii) $|\boldsymbol{N_D[x^*]}| = \mathbf{4}$. By Lemma 6, if $|N_D[x^*]| = 4$, then at most 8 cores $c_1^*, \ldots, c_8^*$ may have x^* as their common reliever, for otherwise we obtain a $(4, 9)$-pendant graph, which cannot be a unit disk graph, . Thus, the average value of $f(\cdot)$ among $x^*, c_1^*, \ldots, c_8^*$ is at most

$$\frac{4 + 8 \cdot 5}{9} = 44/9 = 4.888\ldots.$$

The worst case is therefore the one in which $|N_D[x^*]| = 4$, for an average $\rho = 44/9$, and the lemma follows. $\square$

3.1 Graph-based Algorithm

By Lemma 7, an independent dominating set with no reducible coronas is a 44/9-approximation to the minimum dominating set. In this section, we describe how to obtain such set in linear time given the adjacency list representation of the graph.

We can easily compute a maximal independent set D, which is a 5-approximation to the minimum dominating set [12], in $O(n + m)$ time. An independent dominating set with no reducible coronas can then be obtained by iteratively performing reductions. However, naively performing such reductions leads to a running time of $O(n^2 m)$, since (i) there are $O(n)$ candidates to being the core

[7] We would like to thank an anonymous referee for this simplified argument.

of a reducible corona, (ii) detecting whether a vertex v is in fact the core of a reducible corona by inspecting the 3-neighborhood of v takes $O(m)$ time, and (iii) we may need to reduce a total of $O(n)$ coronas. Fortunately, the following algorithm modifies the set D and returns an independent dominating set with no reducible coronas in $O(n+m)$ time.

(1) For each vertex $v \in V \setminus D$, compute $N_D(v)$.
(2) For each vertex $v \in V \setminus D$, if $|N_D(v)| = 5$, add $N_D(v)$ to the list of coronas $\mathcal{C}$ (unless it is already there).
(3) Let $B \leftarrow \emptyset$. For each corona $C \in \mathcal{C}$, if there is a vertex c such that $D \cup \{c\} \setminus C$ is a dominating set, then add c to the set B.
(4) Choose a maximal subset B' of B such that the pairwise distance of the vertices in B' is at least 5.
(5) For each vertex $c \in B'$, perform a reduction $D \leftarrow D \cup \{c\} \setminus N_D(c)$.
(6) Repeat all the steps above until $B' = \emptyset$.

The algorithm is correct since all changes made to D along its execution preserve the property that D is an independent dominating set. Notice that, in step (4), we only reduce coronas that are sufficiently far from each other, in order to guarantee that we do not reduce a corona that may have ceased to be reducible due to a previous reduction. Moreover, the algorithm always terminates because the size of D decreases at every iteration, except for the last one. Next, we show that the running time is $O(n+m)$.

Step (1) can be easily implemented to run in $O(n+m)$ time. To execute step (2) in $O(n+m)$ time, we must determine in constant time whether a corona is already in the list $\mathcal{C}$. This can be achieved by indexing each corona C by an arbitrary vertex $v \in C$ (say, the one with the lowest index), and by storing with v a list of coronas that are in $\mathcal{C}$ and whose index is v. Note that, because of the packing constraints inherent to unit disk graphs, the number of coronas that contain a given vertex is $O(1)$.

Step (3) can be implemented as follows (for each corona $C \in \mathcal{C}$):

(3a) Let S_1 be the union of the open neighborhoods of the 5 vertices in C.
(3b) Let S_2 be the subset of S_1 containing only the vertices v with $N_D(v) \subseteq C$.
(3c) Let S_3 be the intersection of the closed neighborhoods $N[v]$ of all $v \in S_2 \cup C$.
(3d) If $S_3 \neq \emptyset$, then add an arbitrary vertex of S_3 to the set B.

The steps above take $O(n+m)$ total time when executed for all coronas $C \in \mathcal{C}$, because the number of coronas that contain or are adjacent to a given vertex is also $O(1)$ by packing constraints.

It is easy to perform steps (4) and (5) in linear time. It remains to show that the whole process is only repeated for a constant number of iterations. Let $\mathcal{C}_1, \ldots, \mathcal{C}_k$ denote the set of reducible coronas at each iteration of the algorithm with $\mathcal{C}_k = \emptyset$. Note that the reductions performed in step (5) never create a new reducible corona. Therefore $\mathcal{C}_1 \supset \cdots \supset \mathcal{C}_k$. Let C denote a corona that was reduced in the last iteration k. If C was not reduced during a previous iteration $i < k$, then another corona within constant distance from C was reduced at that very iteration i. Since, again by packing constraints, the maximum number of coronas within constant distance from C is itself a constant, we have $k = O(1)$.

The following theorem summarizes the result from this section.

Theorem 1. *Given the adjacency list representation of a unit disk graph with n vertices and m edges, we can find a 44/9-approximation to the minimum dominating set in $O(n + m)$ time.*

3.2 Geometric Algorithm

In this section, we describe how to obtain an independent dominating set with no reducible corona in $O(n \log n)$ time given the geometric representation of the graph. The input for our algorithm is a set P of n points. Without loss of generality, we assume that the corresponding unit disk graph is connected (otherwise, we can compute the connected components in $O(n \log n)$ time using a Delaunay triangulation [4]). We use terms related to vertices of the graph and to the corresponding points interchangeably. For example, we say a set of points is independent if all pairwise distances are greater than 1.

We want the points of P to be structured in suitable fashion. Thus, as a preliminary step, we sort the points by x-coordinates and by y-coordinates separately (such orderings will also be useful later on), and we partition the points of P according to an infinite grid with unitary square cells by performing two sweeps on the sorted points. Without loss of generality, we assume that no point lies on the boundary of a grid cell. Given $p \in P$, let $\sigma(p)$ denote the grid cell that contains p. We refer to the set of at most 8 non-empty grid cells surrounding a cell Q as the *open vicinity* of Q, denoted $N(Q)$, and to the union of Q and its open vicinity as the *closed vicinity* of Q, denoted $N[Q]$. Note that a point p can only be adjacent to points in the closed vicinity of $\sigma(p)$, that is, $N[p] \subset N[\sigma(p)]$. Each point $p \in P$ stores a pointer to its containing cell $\sigma(p)$. Also, each cell stores the list of points it contains and pointers to the cells in its open vicinity. Since the diameter of the point set is at most n due to the graph connectivity, this whole step can be done in $O(n \log n)$ time.

We are now able to show how to compute a maximal independent set D efficiently. We begin by making a copy P' of P, and by letting $D \leftarrow \emptyset$. Then we repeat the two following steps while set P' is non-empty. (i) Choose an arbitrary point $p \in P'$ and add it to set D. (ii) For each point p' in the closed vicinity of $\sigma(p)$, remove p' from P' if $\|pp'\| \leq 1$. When P' becomes empty, D is an independent dominating set. This process takes $O(n)$ time due to the two following facts. First, a cell belongs to the closed vicinity of a constant number of cells. Second, the number of points inside a cell with pairwise distances greater than 1 is at most a constant.

We now have that D is a maximal independent set, and therefore a 5-approximation to the minimum dominating set. Next, we show how to modify D in order to produce an independent dominating set with no reducible corona, therefore a 44/9-approximation to the minimum dominating set. The algorithm follows a close parallel to the one in Section 3.1, but each step takes no more than $O(n \log n)$ time using the geometric representation of the graph.

Since D is an independent set and a grid cell Q has side 1, a simple packing argument shows that $|D \cap Q| \leq 4$. We store the set $D \cap Q$ in the corresponding cell Q. In order to compute $N_D(p)$, it suffices to inspect at most the 36 points in $D \cap Q$ for $Q \in N[\sigma(p)]$. We can then build a list of coronas in $O(n)$ time (steps (1) and (2) of Section 3.1).

To perform step (3), we need to find out whether there is a vertex c such that $D \cup \{c\} \setminus C$ is a dominating set, for each corona $C = \{p_1, \ldots, p_5\}$. First, we make S_1 the union of $N_D(p_i)$ for $1 \leq i \leq 5$. Then, we make S_2 the subset of S_1 containing only the points p with $N_D(p) \subseteq C$. These first two steps are similar to steps (3a) and (3b) in Section 3.1. The remaining sub-steps of step (3) are significantly different, though.

We proceed by making $S_3 = S_2 \cup C$. We need to determine whether there is a point $p \in S_3$ that is adjacent to all points in S_3. For each $p \in S_3$, let $\beta(p)$ denote the disk of radius 1 centered at p. Let R denote the convex region defined by the intersection of $\beta(p)$ for all $p \in S_3$. A point p is adjacent to all points in S_3 if and only if $p \in R$. We can compute the region R in $O(|S_3| \log |S_3|)$ time using divide-and-conquer in a manner analogous to half-plane intersection [4]. We can then test whether each point $p \in S_3$ belongs to the region R in logarithmic time using binary search (remember the points were previously sorted). If there is at least one point $p \in S_3 \cap R$, then we add p to the set B. Therefore, the whole step (3) takes $O(n \log n)$ time.

In step (4) of the geometric algorithm, we choose an alternative set $B' \subset B$ which can be computed in $O(n)$ time as follows. For each $p \in B$, we add p to B' and then remove from B all points that are contained in the cells within Euclidean distance at most 4 of $\sigma(p)$. Since by packing constraints there are $O(1)$ points in the intersection of D and the closed vicinity of a cell, we can easily perform step (5) in $O(n)$ time.

We summarize the result from this section in the following theorem.

Theorem 2. *Given a set of n points representing a unit disk graph, we can find a $44/9$-approximation to the minimum dominating set in $O(n \log n)$ time in the Real RAM model of computation.*

4 Conclusion and Open Problems

We introduced novel linear and near-linear time algorithms for approximating the minimum dominating set and minimum independent dominating set in a unit disk graph, proving an upper bound of 44/9 to the approximation factor of our algorithms. Nevertheless, the best lower bound we are aware of is 4.8, which is attained by the unit disk graph in Figure 2. Closing this gap would likely require the development of new tools to prove that certain graphs are not unit disk graphs. Computer generated proofs may be useful towards this goal.

References

1. P. Christiano, J. A. Kelner, A. Madry, D. A. Spielman, and S.-H. Teng. Electrical flows, Laplacian systems, and faster approximation of maximum flow in undirected graphs. In *Proc. 43rd annual ACM Symp. on Theory of Computing (STOC)*, pages 273–282, 2011.
2. B. N. Clark, C. J. Colbourn, and D. S. Johnson. Unit disk graphs. *Discrete Mathematics*, 86(1–3):165–177, 1990.
3. M. De, G. Das, and S. Nandy. Approximation algorithms for the discrete piercing set problem for unit disk. In *Proc. 23rd Canadian Conference on Computational Geometry (CCCG)*, pages 375–380, 2011.
4. M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer, 2010.
5. T. Erlebach and M. Mihalák. A $(4+\epsilon)$-approximation for the minimum-weight dominating set problem in unit disk graphs. In *Proc. 7th Workshop on Approximation and Online Algorithms*, volume 5893 of *Lecture Notes in Computer Science*, pages 135–146, 2010.
6. L. Fejes Tóth. *Lagerungen in der Ebene, auf der Kugel und im Raum*. Springer-Verlag, 1953.
7. F. Fodor. The densest packing of 13 congruent circles in a circle. *Contributions to Algebra and Geometry*, 44(2):431–440, 2003.
8. S. R. Frinch. *Mathematical Constants*. Number 94 in Encyclopedia of Mathematics and its Applications. Cambridge, 2003.
9. M. Gibson and I. Pirwani. Algorithms for dominating set in disk graphs: Breaking the $\log n$ barrier. In *Proc. 18th Annual European Symposium on Algorithms (ESA)*, volume 6346 of *Lecture Notes in Computer Science*, pages 243–254, 2010.
10. H. B. Hunt III, M. V. Marathe, V. Radhakrishnan, S. Ravi, D. J. Rosenkrantz, and R. E. Stearns. NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs. *Journal of Algorithms*, 26:238–274, 1998.
11. J. L. Hurink and T. Nieberg. Approximating minimum independent dominating sets in wireless networks. *Information Processing Letters*, 109(2):155–160, 2008.
12. M. V. Marathe, H. Breu, H. B. Hunt III, S. S. Ravi, and D. J. Rosenkrantz. Simple heuristics for unit disk graphs. *Networks*, 25(2):59–68, 1995.
13. C. McDiarmid and T. Muller. Integer realizations of disk and segment graphs. *preprint*, arXiv:1111.2931, 2011.
14. T. Nieberg. *Independent and Dominating Sets in Wireless Communication Graphs*. PhD thesis, University of Twente, 2006.
15. T. Nieberg, J. Hurink, and W. Kern. Approximation schemes for wireless networks. *ACM Transactions on Algorithms*, 4(4):49:1–49:17, 2008.
16. J. Pál. Ein minimumprobleme für ovale. *Math. Annalen*, 83:311–319, 1921.
17. J. Spinrad. *Efficient Graph Representations*. Fields Inst. monographs. AMS, 2003.
18. D. E. D. Vinkemeier and S. Hougardy. A linear-time approximation algorithm for weighted matchings in graphs. *ACM Transactions on Algorithms*, 1:107–122, 2005.
19. F. Zou, Y. Wang, X.-H. Xu, X. Li, H. Du, P. Wan, and W. Wu. New approximations for minimum-weighted dominating sets and minimum-weighted connected dominating sets on unit disk graphs. *Theoretical Computer Science*, 412(3):198–208, 2011.