
Department of Applied Mathematics
Faculty of EEMCS



University of Twente
The Netherlands

P.O. Box 217
7500 AE Enschede
The Netherlands

Phone: +31-53-4893400

Fax: +31-53-4893114

Email: memo@math.utwente.nl
www.math.utwente.nl/publications

Memorandum No. 1732

**A PTAS for the minimum
dominating set problem
in unit disk graphs**

T. NIEBERG AND J.L. HURINK

August, 2004

ISSN 0169-2690

A PTAS for the Minimum Dominating Set Problem in Unit Disk Graphs

Tim Nieberg* and Johann Hurink

University of Twente
Faculty of Electrical Engineering, Mathematics & Computer Science
Postbus 217, NL-7500 AE Enschede
{T.Nieberg,J.L.Hurink}@utwente.nl

Abstract. We present a polynomial-time approximation scheme (PTAS) for the minimum dominating set problem in unit disk graphs. In contrast to previously known approximation schemes for the minimum dominating set problem on unit disk graphs, our approach does not assume a geometric representation of the vertices (specifying the positions of the disks in the plane) to be given as part of the input. The algorithm accepts any undirected graph as input, and is robust in the sense that for instances not reflecting unit disk graphs, it either returns a $(1 + \varepsilon)$ -approximate minimum dominating set, or a certificate showing that the input graph is no unit disk graph. The given PTAS can easily be adapted to other classes of related geometric intersection graphs.

Keywords: Minimum Dominating Set, Geometric Intersection Graph, PTAS

AMS-Classification: 05R69, 05R62, 68R10

1 Introduction

In this paper, we consider the *minimum dominating set* (MDS) problem of finding a dominating set of minimum cardinality in a unit disk graph for the case that no geometric representation of the graph is available. A graph is a *unit disk graph* (UDG) if its vertices can be drawn as circular disks of equal radius in the plane in such a way that there is an edge between two vertices if and only if the two disks have a non-empty intersection. Such a drawing, i.e. a list of center points of the vertices/disks, is referred to as geometric representation of the graph. A subset of vertices in such an undirected graph is called *dominating set* if every vertex in the graph either is contained in the subset, or adjacent to a vertex in the set. The MDS problem is NP-hard, even on unit disk graphs when a geometric representation is given [4].

We present a polynomial-time approximation scheme (PTAS) for the MDS problem, that is, given any $\varepsilon > 0$, the algorithm gives in polynomial-time an approximation with a performance guarantee of $(1 + \varepsilon)$.

* This work is partially supported by the European research project EYES (IST-2001-34734).

Unit disk graphs are widely used to model the communication in wireless ad-hoc networks (e.g. [5]). In such a network, structures like dominating sets play an important role, e.g. in global flooding to alleviate the so-called broadcast storm problem. A message broadcast only in the dominating set is an efficient way to ensure that it is received by all transmitters in the network, both in terms of energy and interference. In wireless ad-hoc networks, the transmitters mostly do not know their own location, as obtaining this information is hard or costly, making algorithms that do not rely on positional information favorable.

Most of the work concerning approximation schemes in unit disk graphs assume a given geometric representation, which allows for separation of the graph along a grid ([1],[6]). Approximation schemes for the MDS in unit disk graphs (with representation), and other related problems, are given in [7]. In [3], a PTAS for the minimum connected dominating set is presented, also using separation of the graph based on a grid.

However, the case when no geometric representation is present is significantly different: Computing a possible geometric representation for a given unit disk graph is NP-hard. Indeed, any polynomial-time algorithm computing a geometric representation for UDGs can be used in a straightforward way to solve the corresponding recognition problem, i.e. determining whether a given graph is a UDG, a problem known to be NP-hard [2]. For the case that a representation is not given, several approximation algorithms are presented in [9], including a 5-approximation for the MDS problem.

Besides the independence from a geometric representation, an additional advantage of the presented PTAS lies in the fact that we can extend the algorithm towards a *robust* approximation [11]. The algorithm may then be applied to an arbitrary undirected graph, and the output is either a $(1 + \varepsilon)$ -approximation for the MDS problem in this graph, or a certificate which allows us to prove in polynomial-time that the input graph is no unit disk graph. In other words, we have a polynomial-time algorithm which either approximates the MDS problem, or solves the recognition problem. In case the input graph is a UDG, the algorithm always returns a dominating set of desired quality.

The remainder of the paper is organized as follows. In the following section, we present some basic definitions needed for the description of the approximation scheme. Section 3 introduces the concept of a 2-separated collection of subsets, a structure that is used to efficiently separate a graph into smaller subgraphs for which the problem of computing a dominating set is easier to tackle. The PTAS itself is then presented in Section 4, where we also show the polynomial running-time of the method. In Section 5, we discuss the ro-

bustness of the algorithm, and present some extensions to other intersection graphs of geometric objects.

2 Definitions and preliminaries

A graph $G = (V, E)$ is a unit disk graph (UDG) if it results from the intersection graph of disks of unit radius in the euclidean plane. In other words, G is a UDG if there exists a map $f : V \rightarrow \mathbb{R}^2$ satisfying

$$(u, v) \in E \iff \|f(u) - f(v)\| \leq 2,$$

where $\|\cdot\|$ denotes the euclidean norm. In this context, f is called a geometric representation of G and is not unique for a given graph.

A subset $D \subset V$ is a *dominating set* (for V) if for every vertex $v \in V$, either $v \in D$ holds or there exists an edge $(u, v) \in E$ such that $u \in D$. The minimum dominating set problem (MDS) seeks to find a dominating set of minimum cardinality for a given graph.

In this paper, the goal is to give a polynomial-time approximation scheme (PTAS) for the minimum dominating set problem on unit disk graphs for the case that a geometric representation is not given. That is, we seek for an algorithm which, given a UDG $G = (V, E)$ and a parameter $\varepsilon > 0$, computes a dominating set of cardinality no more than $(1 + \varepsilon)$ the size of a minimum dominating set in G . The running time of the algorithm is allowed to depend on the parameter ε , but should be polynomial with respect to the input instance, i.e. polynomial in $n = |V|$ for fixed $\varepsilon > 0$.

We now present some further definitions needed for the description and discussion of the algorithm and the underlying concepts. Without loss of generality, we may assume the graph G to be connected. If this is not the case, we may consider each connected component separately.

Let $W \subset V$ denote a set of vertices in $G = (V, E)$. In the following, we simultaneously use W to also denote the resulting induced subgraph $G[W] := (W, E \cap (W \times W))$. Obviously, the graph $G[W]$ is a unit disk graph if the original graph is one.

Furthermore, we denote by $N(v)$ the closed neighborhood of a vertex $v \in V$, i.e. $N(v) := \{u \in V \mid (u, v) \in E\} \cup \{v\}$. Analogously, for $W \subset V$, let $N(W) := \bigcup_{w \in W} N(w)$ define the neighborhood of W . In this context, we set $N(\emptyset) := \emptyset$. For $r \in \mathbb{N}$, we denote by $N^r(v) := N(N^{r-1}(v))$ the recursively defined r -th neighborhood of $v \in V$, where $N^1(v) := N(v)$.

For two vertices $u, v \in V$, let $d(u, v)$ denote the (hop-)distance between u and v , that is the number of edges on a shortest path between these two




Fig. 1. Example of a UDG with and without geometric representation

vertices. Thus, alternatively, the r -th neighborhood of $v \in V$ is characterized by $N^r(v) = \{u \in V \mid d(u, v) \leq r\}$.

Denote by $\mathcal{P}(V)$ the set of all subsets of vertices. We then define $D : \mathcal{P}(V) \rightarrow \mathcal{P}(V)$ to be an operation returning a dominating set of minimum cardinality for the subset of vertices given as argument to it. Since for a subset $W \subset V$, the set $D(W)$ dominates W , i.e. for every $w \in W$, either $w \in D(W)$ holds, or there is an edge $(u, w) \in E$ such that $u \in D(W)$, from the above definitions, we can state that $W \subset N(D(W))$ holds. Furthermore, it is easy to see that $D(W) \subset N(W)$ holds. In the following, we are interested in an efficient, i.e. polynomial-time, approximation of $D(V)$ within a factor of $(1 + \varepsilon)$ for any given $\varepsilon > 0$.

Figure 1 illustrates some of the given notations. In the left part, a graph and its geometric representation are given, whereas in the right part only the graph and some neighborhoods of a node v are presented. Furthermore, the circled vertices in the right part give a minimum dominating set for $N^3(v)$, i.e. $D(N^3(v))$. As can be seen from the example, $D(W) \subset W$ need not hold for a subset $W \subsetneq V$: Using the circled vertex in $N^4(v)$, we obtain a dominating set consisting of three vertices, whereas restricting the dominating set only to vertices from $N^3(v)$ yields dominating sets of cardinality 4 or higher.

3 Local dominating sets

In this section, we introduce the concept of a 2-separated collection of subsets. The subgraphs induced by the subsets of such a collection divide the original graph into smaller parts for which it becomes easier to tackle the MDS problem. For a collection of local dominating sets resulting from a separation of




Fig. 2. Example for a 2-separated collection $\mathcal{S} = \{S_1, \dots, S_6\}$

the graph into smaller subgraphs, we show several properties that allow for bounds on the cardinalities with respect to an optimal solution. Throughout this section, we do not assume the graph to be a UDG, the following concepts are valid for all undirected graphs.

For a graph $G = (V, E)$, let $\mathcal{S} := \{S_1, \dots, S_k\}$ be a collection of subsets of vertices $S_i \subset V$, $i = 1, \dots, k$, with the following property:

(P) for any two vertices $s \in S_i$ and $\bar{s} \in S_j$ with $i \neq j$, it is $d(s, \bar{s}) > 2$.

We refer to \mathcal{S} as a 2-separated collection of subsets. An example of such a 2-separated collection is presented in Figure 2. The grey areas mark the different subsets that make up the collection, vertices which are not part of the collection, and thus separate the subsets are white.

The following lemma shows that the sum of the cardinalities of minimum dominating sets $D(S_i)$ for the subsets $S_i \in \mathcal{S}$ of a 2-separated collection forms a lower bound on the cardinality $|D(V)|$ of a minimum dominating set in G .

Lemma 1. *For a 2-separated collection $\mathcal{S} = \{S_1, \dots, S_k\}$ in a graph $G = (V, E)$, we have*

$$|D(V)| \geq \sum_{i=1}^k |D(S_i)|.$$

Proof. For each subset $S_i \in \mathcal{S}$, consider the neighborhood $N(S_i)$. As a direct result of property (P), these neighborhoods are pairwise disjoint. Furthermore, any vertex outside $N(S_i)$ has distance more than one to all vertices in S_i . Thus,

$D(V) \cap N(S_i)$ has to dominate all vertices in S_i , since $D(V)$ dominates the entire vertex set V .

On the other hand, also $D(S_i) \subset N(S_i)$ dominates S_i using a minimum number of vertices in G . Therefore, we get

$$|D(V) \cap N(S_i)| \geq |D(S_i)|.$$

Combining this for all subsets of the 2-separated collection, we get

$$|D(V)| \geq \sum_{i=1}^k |D(V) \cap N(S_i)| \geq \sum_{i=1}^k |D(S_i)|,$$

as claimed. \square

Lemma 1 states that a 2-separated collection \mathcal{S} leads to a lower bound on the cardinality of a MDS. Additionally, such a collection may help in getting an approximation of this cardinality. If we are able to enlarge the subsets S_i to subsets T_i in such a way that the dominating sets of the expansions are locally bounded and the unions of these forms a dominating set for V , we get a global approximation for the MDS in G .

Corollary 1. *Let $\mathcal{S} = \{S_1, \dots, S_k\}$ be a 2-separated collection in $G = (V, E)$, and let T_1, \dots, T_k be subsets of V with $S_i \subset T_i$ for all $i = 1, \dots, k$. If there exists a bound $\rho \geq 1$ such that*

$$|D(T_i)| \leq \rho \cdot |D(S_i)|$$

holds for all $i = 1, \dots, k$, and if $\bigcup_{i=1}^k D(T_i)$ forms a dominating set in G , the set $\bigcup_{i=1}^k D(T_i)$ is a ρ -approximation of an MDS in G .

Proof. $|\bigcup_{i=1}^k D(T_i)| \leq \sum_{i=1}^k |D(T_i)| \leq \rho \cdot \sum_{i=1}^k |D(S_i)| \leq \rho \cdot |D(V)|.$ \square

In the following section, we focus on the construction of suitable subsets $T_i \subset V$, which contain a 2-separated collection $S_i \subset T_i$, in a way that a local $(1 + \varepsilon)$ -approximation can be guaranteed. Furthermore, we create these subsets in such a way that the union of the respective local dominating sets also dominates the entire set of vertices, resulting in a global $(1 + \varepsilon)$ -approximation for the MDS.

4 Efficient construction of suitable subsets

From the previous discussion, we recall that if we have a 2-separated collection $\mathcal{S} := \{S_1, \dots, S_k\}$, corresponding sets $T_i \supset S_i$ together with a bound of $(1 + \varepsilon)$

for the local dominating sets $D(S_i)$ and $D(T_i)$, then the union of the $D(T_i)$ satisfies the approximation bound required for a PTAS for the MDS problem. In this section, we show how to construct suitable subsets, for which the union of the local dominating sets also forms a dominating set for V . Furthermore, we prove that this can be achieved in polynomial running-time with respect to the size of the input instance for fixed $\varepsilon > 0$ if the input graph is a UDG. For ease of notation, let $\rho := (1 + \varepsilon)$ denote the desired approximation guarantee of the algorithm.

The basic idea of the construction is simple: we compute a local dominating set for a neighborhood of a vertex, and expand this neighborhood until we have formed sets S and $T \supset S$ which satisfy a desired bound. Then, we eliminate the current neighborhood and continue the same steps for the remaining graph.

In more detail, the algorithm works as follows. We start with an arbitrary vertex $v \in V$ and consider for $r = 0, 1, 2, \dots$, the r -th neighborhoods $N^r(v)$. Starting with $N^0(v) = v$, we compute dominating sets of minimum cardinality for these neighborhoods as long as

$$|D(N^{r+2}(v))| > \rho \cdot |D(N^r(v))| \quad (1)$$

holds. Denote by \hat{r}_1 the smallest r for which (1) is violated.

We go on iteratively with this procedure for the graph induced by $V_{i+1} := V_i \setminus N^{\hat{r}_i+2}(v_i)$, where $V_1 := V$. The vertex $v_i \in V_i$ is chosen as an arbitrary central vertex of the neighborhoods. In further iterations, we thus consider for $r = 0, 1, 2, \dots$ the neighborhoods $N^r(v_i)$ with respect to V_i , i.e. we have $N^r(v_i) \subset V_i$. Note that the dominating sets $D(\cdot)$ are always computed with respect to the entire input graph G .

This process is then repeated until V_{i+1} contains no more vertices. Let $k \in \mathbb{N}$ be the total number of iterations. Obviously we have $k < n$. In the following, let $N_i, i = 1, \dots, k$, denote the respective neighborhoods when the stopping criterion (1) is violated, i.e. $N_i := N^{\hat{r}_i+2}(v_i)$.

Looking at the dominating sets for these neighborhoods, $D(N_i)$, we have the following lemma which shows that a dominating set for the entire graph is given by the union of the sets $D(N_i)$.

Lemma 2. *For the collection of neighborhoods $\{N_1, \dots, N_k\}$ created by the above algorithm, the union $D := \bigcup_{i=1}^k D(N_i)$ forms a dominating set for the input graph G .*

Proof. It is $V_{i+1} = V_i \setminus N_i$ and $N_i \subset V_i$, thus we have $V_i = V_{i+1} \cup N_i$. We stop the algorithm at $V_{k+1} = \emptyset$, which implies $V_k = N_k$. Therefore $\bigcup_{i=1}^k N_i = V$ by induction, and the claim follows. \square

Next, we show that the solution set $D := \bigcup_{i=1}^k D(N_i)$ returned by the algorithm satisfies the $(1 + \varepsilon)$ -bound on the approximation. We show that $\mathcal{N} := \{N^{\hat{r}_1}(v_1), \dots, N^{\hat{r}_k}(v_k)\}$ is a 2-separated collection in G , and then apply Corollary 1 to the respective local dominating sets $D(N_i)$.

Lemma 3. *The subsets $N^{\hat{r}_i}(v_i), i = 1, \dots, k$, created by the algorithm form a 2-separated collection $\mathcal{N} := \{N^{\hat{r}_1}(v_1), \dots, N^{\hat{r}_k}(v_k)\}$ in G .*

Proof. For ease of notation, let \overline{N}_i denote the neighborhood $N^{\hat{r}_i}(v_i)$ for iteration $i \in \{1, \dots, k\}$ of the algorithm. Recall that a 2-separated collection is characterized by property (P), i.e. vertices of two different subsets of the collection have distance more than 2 from one another.

Clearly, $\{\overline{N}_1, V_2\}$ is a 2-separated collection in G , since $V_2 = V \setminus N(N(\overline{N}_1))$. For induction, suppose that $\{\overline{N}_1, \dots, \overline{N}_{i-1}, V_i\}$ is a 2-separated collection in G . Any vertex in V_i has distance more than 2 from any other vertex in $\overline{N}_1, \dots, \overline{N}_{i-1}$. Considering $V_{i+1} = V_i \setminus N(N(\overline{N}_i))$, we see that both V_{i+1} and \overline{N}_i satisfy (P). Therefore, $\{\overline{N}_1, \dots, \overline{N}_i, V_{i+1}\}$ again is a 2-separated collection. \square

Additionally, the criterion (1) for stopping to expand the neighborhood guarantees that each pair of local dominating sets satisfies

$$|D(N_i)| \leq \rho \cdot |D(N^{\hat{r}_i}(v_i))| \quad (i = 1, \dots, k). \quad (2)$$

Using Corollary 1 and Lemma 2, we now obtain the following result for the approximation.

Corollary 2. *The above algorithm returns a dominating set $\bigcup_{i=1}^k D(N_i)$ of cardinality no more than $(1 + \varepsilon)$ the size of a minimum dominating set in $G = (V, E)$. \square*

At this point, it is noteworthy to remind that this Corollary 2 is valid for any undirected graph G , even if it is not a unit disk graph.

It remains to show that the $(1 + \varepsilon)$ -approximation algorithm has polynomial running-time. In contrast to Corollary 2, the polynomial running-time relies on the fact that the input graph G is a unit disk graph. So, for the further discussion in this section, we assume G to be a unit disk graph.

The number k of iterations is bounded by $n = |V|$. We may thus limit the further discussion to one iteration only. Since any V_i during the execution of the algorithm again induces a unit disk graph, we focus w.l.o.g. on the graph $G = (V, E)$ in the first iteration. We show two things:

- (1) we can compute the minimal dominating set $D(N^r(v))$ in polynomial time if the value of r is a constant or polynomially bounded; and

- (2) there exists a constant bound for \hat{r}_1 , i.e. the diameter of the largest neighborhood we need to consider until the stopping criterion (1) is violated.

Before showing that $D(N^r(v))$ can be computed efficiently, we need to introduce the notion of an independent set, and briefly state a key result for independent sets in UDGs.

Let $W \subset V$. A set $I \subset W$ is called an *independent set* if for every two vertices $u, v \in I$, there does not exist an edge $(u, v) \in E$. An independent set is called *maximal* in W if we cannot add any other vertex from W to I without violating the independence property (of no two vertices being adjacent). Clearly, any maximal independent set in W also dominates W .

For a UDG, the following result of [10] bounds the size of an independent set in the neighborhood $N^r(v)$. We give the short proof, since we rely on it in the next section.

Lemma 4. *Let $G = (V, E)$ be a UDG. Any independent set $I^r \subset N^r(v), v \in V$, satisfies*

$$|I^r| \leq (2r + 1)^2 = O(r^2).$$

Proof. Let $f : V \rightarrow \mathbb{R}^2$ be a geometric representation of G . From the definition of a unit disk graph, we conclude that any $w \in N^r(v)$ satisfies

$$\|f(v) - f(w)\| \leq 2r.$$

Thus, I^r consists of pairwise disjoint disks of unit radius inside a disk of radius $2r + 1$ around $f(v)$, and therefore $|I^r| \leq \pi(2r + 1)^2/\pi$. \square

As a consequence of Lemma 4, any independent set in $N^r(v)$ is polynomially bounded in r , including maximal independent sets. The cardinality of a minimum dominating set in $N^r(v)$ is bounded from above by the cardinality of a maximal independent set in $N^r(v)$, and, therefore, we get

Corollary 3. $|D(N^r(v))| \leq (2r + 1)^2 = O(r^2)$. \square

Assuming r to be fixed or polynomially bounded, a minimum dominating set $D(N^r(v))$ can then be computed in polynomial time, e.g. by complete enumeration in time $O(n^\vartheta)$, with $\vartheta = O(r^2)$.

Next, we show that, for a UDG, there exists such a bound on \hat{r}_1 , the first value of r which violates (1). This bound only depends on the approximation ratio ρ , and not on the size of the unit disk graph $G = (V, E)$ given as input.

Lemma 5. *There exists a constant $c = c(\rho)$ such that $\hat{r}_1 \leq c$, that is, the largest neighborhood to be considered during the iteration of the algorithm is bounded by a constant.*

Proof. It is $|D(N^0(v))| = |D(N^1(v))| = 1$, as the central vertex v dominates itself and all its neighbors.

Consider an arbitrary value of $r < \hat{r}_1$. First, if r is an even number, due to the stopping criterion (1) we have

$$(2r + 1)^2 \geq |D(N^{r+2}(v))| > \rho |D(N^r(v))| > \cdots > \rho^{\frac{r}{2}} |D(N^0(v))| = (\sqrt{\rho})^r.$$

Second, if r is uneven, we get

$$(2r + 1)^2 \geq |D(N^{r+2}(v))| > \rho |D(N^r(v))| > \cdots > \rho^{\frac{r}{2}} |D(N^1(v))| = (\sqrt{\rho})^r.$$

Since $\rho > 1$, and thus $\sqrt{\rho} > 1$, in both cases the above inequations have to be violated eventually. The bound on \hat{r}_1 when these inequations are violated the first time only depends on ρ and not on the size of the overall graph G . The claim follows directly. \square

Setting $c := \frac{1}{\varepsilon^2} \ln \frac{1}{\varepsilon}$ yields $(2c + 1)^2 < (1 + \varepsilon)^c$, thus the constant for c is bounded by $O(\frac{1}{\varepsilon^2} \log \frac{1}{\varepsilon})$.¹

Summarizing, if the input graph is a UDG, each iteration has polynomial running time, and therefore the presented algorithm is a polynomial-time approximation scheme for the MDS problem. Note that the computation of $D(N^r(v))$ for the largest neighborhood, dominates the running-time of the algorithm. Therefore, the overall time complexity of the approximation is $O(n^{c^2})$ with $c = O(\frac{1}{\varepsilon^2} \log \frac{1}{\varepsilon})$.

5 Discussion

Unit disk graphs are a special subclass of undirected graphs. As we have shown in the previous part, the presented algorithm accepts an arbitrary undirected graph as input, and returns a dominating set of desired quality for this graph. However, the polynomial running-time relies on the UDG characterization. This raises the question of robustness for algorithms designed for a restricted domain [11]:

An algorithm \mathcal{A} , defined on a set \mathcal{G} of instances, is robust on a restricted class $\mathcal{U} \subset \mathcal{G}$ if it solves the problem for all instances in \mathcal{U} , and for instances not in \mathcal{U} , the algorithm either solves the problem, or provides a certificate that the input does not belong to \mathcal{U} . Of course, the notion of a robust algorithm is especially interesting when \mathcal{A} has polynomial running-time with respect to the size of the input instance, and the decision whether an instance belongs

¹ For sufficiently small $\varepsilon \leq \frac{1}{2}$, we can use the inequality $\ln(1 + \varepsilon) + \varepsilon^2 \geq \varepsilon$ to obtain this result.

to the subclass $\mathcal{U} \subset \mathcal{G}$ is not as easy to decide. In our situation, \mathcal{G} is the set of undirected graphs, \mathcal{A} computes a $(1 + \varepsilon)$ -approximation of the cardinality of an MDS, and \mathcal{U} is the subclass of UDGs.

In case the input graph is a unit disk graph, the algorithm always returns a $(1 + \varepsilon)$ -approximate dominating set in polynomial running-time. Also, when the input is any undirected graph, such an approximation is returned. However, the polynomial running-time in this case cannot be guaranteed. In the following, we consider the case that the input is no UDG.

The time complexity of the algorithm is a direct result of the possibility to bound the cardinality of a minimum dominating set in a neighborhood of bounded diameter. This bound results from the fact that a maximal independent set I^r in such a neighborhood is bounded, i.e. for the r -th neighborhood of a vertex $v \in V$, we have

$$|D(N^r(v))| \leq |I^r| \leq (2r + 1)^2.$$

If, during the execution of the algorithm, we now find a neighborhood $N^r(v)$ for which a minimum dominating set of size less than or equal to $(2r + 1)^2$ cannot be found, we terminate the algorithm, and output the neighborhood $N^r(v)$ as a certificate to show that the input is no UDG. For this neighborhood, we can then construct a maximal independent set, e.g. by a simple greedy strategy in time $O(n^2)$, which has to violate Lemma 4. This immediately shows that the input graph cannot be a unit disk graph.

Note that for robustness, we do not need to explicitly consider the bound $r \leq c$ (Lemma 5) on the diameter of the neighborhoods $N^r(v)$ constructed, as this bound follows from the polynomial bound on the cardinality of the dominating sets in the neighborhoods.

The PTAS presented in this paper can be extended in a straightforward way to intersection graphs of other, related geometric objects, e.g. the unit disk graph may be defined using other geometric norms. From the discussion on the complexity in the previous section, it can be seen that a sufficient condition for the existence of a PTAS for the MDS problem in a geometric intersection graph is given when there is a polynomial bound on the ratio of maximum geometric diameter divided by minimum volume of the objects that make up the intersection graph (see Lemma 4). Thus, the objects in consideration do not necessarily need to be of equal size or shape, e.g., the unit disks may be replaced by disks with fixed lower and upper bounds on the radius. This condition includes Quasi Unit Disk Graphs which are used to give a more realistic model of a wireless, ad-hoc network [8]. An extension to a (fixed) dimension $d > 2$ is also immediately possible.

6 Conclusion

In this paper, we present a new polynomial-time approximation scheme for the minimum dominating set problem in unit disk graphs. The algorithm does not need a geometric representation of the graph to compute the $(1 + \varepsilon)$ -approximate dominating set. In fact, it accepts any undirected graph as input and returns either a dominating set which satisfies the desired bound, or a certificate to show that the input graph is no UDG. Of course, if the input graph satisfies the characterization of a UDG, a dominating set is always returned.

The approximation algorithm that results in the PTAS works by exploiting the fact that inside a neighborhood of bounded diameter, a locally optimal minimum dominating set is polynomially bounded in size and can thus be efficiently computed. For UDGs, the largest neighborhood which has to be considered in the algorithm is also bounded by a constant that only depends on the desired approximation factor, and not on the size of the input graph. The overall time complexity of the (robust) approximation algorithm is $O(n^{c^2})$ with $c = O(\frac{1}{\varepsilon^2} \log \frac{1}{\varepsilon})$.

References

1. B.S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *Journal of the ACM*, 41(1):153–180, 1994.
2. H. Breu and D.G. Kirkpatrick. Unit disk graph recognition is NP-hard. *Computational Geometry. Theory and Applications*, 9(1-2):3–24, 1998.
3. X. Cheng, X. Huang, D. Li, W. Wu, and D.-Z. Du. A polynomial-time approximation scheme for the minimum-connected dominating set in ad hoc wireless networks. *Networks*, 42:202–208, 2003.
4. B. N. Clark, C. J. Colburn, and D. S. Johnson. Unit disks graphs. *Discrete Mathematics*, 86:165–177, 1990.
5. EYES (Energy-Efficient Sensor Networks). website. <http://eyes.eu.org>.
6. D.S. Hochbaum and W. Maass. Approximation schemes for covering and packing problems. *Journal of the ACM*, 32(1):130–136, 1985.
7. H.B. Hunt III, M.V. Marathe, V. Radhakrishnan, S. S. Ravi, D.J. Rosenkrantz, and R.E. Stearns. NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs. *J. Algorithms*, 26(2):238–274, 1998.
8. F. Kuhn, R. Wattenhofer, and A. Zollinger. Ad-hoc networks beyond unit disk graphs. In *Proceedings of the 1st ACM DIALM-POMC Joint Workshop on Foundations of Mobile Computing*, 2003.
9. M.V. Marathe, H. Breu, H.B. Hunt III, S. S. Ravi, and D.J. Rosenkrantz. Simple heuristics for unit disk graphs. *Networks*, 25:59–68, 1995.
10. T. Nieberg, J. Hurink, and W. Kern. A robust PTAS for maximum independent sets in unit disk graphs. In *Proceedings of the 30th workshop on graph theoretic concepts in computer science*, pages xx–xx. to be published, LNCS, 2004.

11. V. Raghavan and J. Spinrad. Robust algorithms for restricted domains. In *Proceedings of the 12th annual ACM-SIAM symposium on discrete algorithms*, pages 460–467. Society for Industrial and Applied Mathematics, 2001.