J Comb Optim
DOI 10.1007/s10878-009-9216-y

Polynomial time approximation schemes for minimum disk cover problems

Chen Liao · Shiyan Hu

Abstract The following planar *minimum disk cover* problem is considered in this paper: given a set $\mathcal{D}$ of n disks and a set $\mathcal{P}$ of m points in the Euclidean plane, where each disk covers a subset of points in $\mathcal{P}$, to compute a subset of disks with minimum cardinality covering $\mathcal{P}$. This problem is known to be NP-hard and an algorithm which approximates the optimal disk cover within a factor of $(1+\epsilon)$ in $\mathcal{O}(mn^{\mathcal{O}(\frac{1}{\epsilon^2}\log^2\frac{1}{\epsilon})})$ time is proposed in this paper. This work presents the first polynomial time approximation scheme for the minimum disk cover problem where the best known algorithm can approximate the optimal solution with a large constant factor. Further, several variants of the minimum disk cover problem such as the incongruent disk cover problem and the weighted disk cover problem are considered and approximation schemes are designed.

Keywords Minimum disk cover · Polynomial time approximation scheme · Wireless network · Minimum weight disk cover

1 Introduction

Let $\mathcal{P}$ be a set of n points in the plane. A *disk* D covers a point $p \in \mathcal{P}$ if the Euclidean distance between the center of D and p is no greater than the radius of D. The following *minimum disk cover* problem is considered in this paper: given a set $\mathcal{D}$ of n disks and a set $\mathcal{P}$ of m points in the Euclidean plane, where each disk covers a subset

C. Liao · S. Hu (✉)
Department of Electrical and Computer Engineering, Michigan Technological University, Houghton, MI 49931, USA
e-mail: shiyan@mtu.edu

C. Liao
e-mail: cliao@mtu.edu

Published online: 28 February 2009

of points in $\mathcal{P}$, to compute a subset of disks with minimum cardinality covering all the points in $\mathcal{P}$.

The minimum disk cover problem is known to be NP-hard (Masuyama et al. 1981). The best approximation algorithm is due to (Bronnimann and Goodrich 1995), where a constant factor approximation algorithm is proposed. The algorithm runs in $\mathcal{O}(t^2 n \log n \log(n/t))$ time where t denotes the cardinality of the optimal disk cover. A recent work (Ambühl et al. 2006) considers a natural weighted extension of the minimum disk cover problem, called *minimum weight disk cover problem*, where each disk is weighted. A 72-approximation algorithm is proposed in Ambühl et al. (2006).

In addition to its theoretical interest, minimum disk cover problem has important practical implications on wireless network design. In mobile ad hoc wireless networks, each host is equipped with radio-frequency (RF) transceivers to provide reliable transmission inside a circular range, represented by a disk, within some distance. Due to the high mobility of hosts, a host needs to frequently re-synchronize with its neighbors which are the network nodes within its coverage. To avoid the notorious broadcast storm problem (Ni et al. 1999) and save energy, it is highly desirable to reduce the transmission load. Treating each neighbor by a point in $\mathcal{P}$ as in Calinescu et al. (2004), this asks to select minimum number of hosts in transmission such that all network nodes can be covered/synchronized. This is the minimum disk cover problem considered in this paper.

1.1 Related work

As mentioned before, the minimum disk cover problem is shown to be NP-hard in Masuyama et al. (1981) and the best approximation algorithms are due to (Bronnimann and Goodrich 1995; Ambühl et al. 2006) where large constant factor approximation algorithms are proposed.

Despite limited advance in the minimum disk cover problem, a closely related unit disk graph minimum dominating set problem has been much better studied. A unit disk graph is defined as a graph such that one can realize all vertices as unit disks in the plane and there is an edge between two vertices if and only if their corresponding disks intersect (Clark et al. 1990). Define a dominating set on a graph as a vertex set where for any vertex, either the vertex itself or one of its adjacent vertices is in the set. The unit disk graph minimum dominating set problem asks to compute a dominating set of minimum cardinality in a unit disk graph. It is known to be NP-hard (Clark et al. 1990). A polynomial time approximation scheme (PTAS) is given in Hunt et al. (1998) using the shifting strategy (Hochbaum and Maass 1985; Baker 1994). Recently, a new PTAS on the same problem but without geometric representation is proposed in Nieberg et al. (2008). In Cheng et al. (2003), an interesting variant of the unit disk graph minimum dominating set problem with the additional requirement that the induced dominating set needs to be connected is considered and a PTAS is proposed for the variant. Further work on this direction is to generalize the problem to minimum weight connected dominating set problem. The work in Ambühl et al. (2006) presents a constant approximation to this problem and the work in Huang et al. (2008) improves the approximation ratio to $(10 + \epsilon)$ by a two-phase algorithm. Algorithms for computing minimum m-connected k-dominating set, which asks to pick m-connected vertices S from a graph such that each $p \in P \backslash S$ is adjacent to at least k vertices in S, are given in Shang et al. (2007 2008).

The minimum disk cover problem can also be viewed as a generalization of the *minimum forwarding set* problem in wireless ad hoc network design. In minimum forward set problem, a source node, the set of its 1-hop nodes, and the set of its 2-hop nodes are given, the problem asks to compute a subset of 1-hop nodes with minimum cardinality to cover all 2-hop nodes. From theoretical point of view, it is a special case of the minimum disk cover problem where all 2-hop nodes lie outside a disk while all 1-hop nodes lie inside the disk (Calinescu et al. 2004). The complexity of the minimum forwarding set problem is unknown and a 3-approximation $\mathcal{O}(n^2)$ time algorithm is proposed in Calinescu et al. (2004).

Finally, it is worth noting that there is a recent work on an optimal polynomial-time algorithm for the minimum disk cover problem (Sun et al. 2007). However, the definition of their minimum disk cover problem is quite different from ours. Their problem asks to choose a subset of the given disks such that the union of the chosen disks is equal to the union of the given disks.

1.2 Summary of results

In this paper, the minimum disk cover problem is studied and an algorithm which approximates the optimal disk cover within a factor of $(1 + \epsilon)$ and runs in $\mathcal{O}(mn^{\mathcal{O}(\frac{1}{\epsilon^2}\log^2\frac{1}{\epsilon})})$ time is proposed, where n is the number of disks and m is the number of points. As extensions to our algorithm, approximation schemes for some variants of the problem such as the minimum incongruent disk cover problem and the minimum weight disk cover problem are designed. In particular, the latter admits a quasi polynomial time approximation scheme, i.e., a $(1 + \epsilon)$-approximation algorithm running in $\mathcal{O}(mn^{\mathcal{O}(\frac{1}{\epsilon^2}\log^2 n)})$ time, provided that the ratio between the largest and the smallest disk weights is polynomially bounded in n.

Our algorithms are motivated from the local-neighborhood-based scheme proposed in Nieberg et al. (2008). Note that the work (Nieberg et al. 2008) considers the unit disk graph minimum dominating set problem where disks can be arbitrarily placed in the plane while in our case the disks can only be chosen from a given set. This makes our algorithms quite different from the work (Nieberg et al. 2008).

1.3 Paper organization

The rest of the paper is organized as follows: Sect. 2 presents the problem formulations. Section 3 describes the polynomial time approximation schemes to our problems. A summary of work is given in Sect. 4.

2 Preliminaries

2.1 Notations and definitions

A set of n disks $\mathcal{D}$ and a set of m points $\mathcal{P}$ are given in the Euclidean plane. A *disk* D covers a point $p \in \mathcal{P}$ if the Euclidean distance between the center of D and p is no greater than the radius of D. A set of points $\mathcal{P}$ is covered by a set of disks $\mathcal{D}$ if every point $p \in \mathcal{P}$ is covered by at least one disk $D \in \mathcal{D}$. Let $|\cdot|$ denote the cardinality of a set. In the weighted version of the minimum disk cover problem, each disk $D \in \mathcal{D}$ is also associated with a positive real number as its weight, denoted by $w(D)$. Let $w(\mathcal{D})$ denote the weight of a set of disks $\mathcal{D}$, which is equal to $\sum_{D \in \mathcal{D}} w(D)$. Let $d(\cdot, \cdot)$ denote the Euclidean distance function.

An algorithm is said to approximate a minimization problem within a factor ρ if the algorithm is guaranteed to produce a solution whose objective function value is at most ρ times the optimal solution. A polynomial time approximation scheme, or PTAS in short, for a minimization problem is an algorithm which approximates the optimal solution within a factor of $(1+\epsilon)$ for any $\epsilon > 0$ running in time polynomial in the input size n. A quasi polynomial time approximation scheme, or QPTAS in short, for a minimization problem is an algorithm which approximates the optimal solution within a factor of $(1+\epsilon)$ for any $\epsilon > 0$ running in time $\mathcal{O}(n^{polylog(n)})$ in the input size n. A problem admitting a QPTAS strongly indicates that it has a PTAS, since the problem is not APX-hard unless $NP \subseteq DTIME[n^{polylog(n)}]$.

2.2 Problem formulation

We begin with formulating the minimum congruent disk cover problem.
Minimum congruent disk cover problem: Given a set of congruent disks $\mathcal{D}$ of a constant radius r, a set $\mathcal{P}$ of points of integer coordinates in the Euclidean plane, to compute a subset of $\mathcal{D}$ of the minimum cardinality covering $\mathcal{P}$.

This problem has been shown to be NP-hard (Masuyama et al. 1981; Johnson 1982) even when the radius r is a constant (e.g., 2 as in Masuyama et al. (1981)) by reduction from planar 3SAT problem. This problem is also closely related to the geometric connected dominating set problem (Johnson 1982; Garey and Johnson 1979).

In our minimum congruent disk cover problem, disks in $\mathcal{D}$ are assumed to be congruent with the constant radius r. However, our algorithm can be easily generalized to handle incongruent disks as long as the maximum radius is still bounded by a constant.
Minimum incongruent disk cover problem: Given a set of disks $\mathcal{D}$ whose maximum radius is bounded above by a constant, a set $\mathcal{P}$ of integer coordinate points in the Euclidean plane, to compute a subset of $\mathcal{D}$ of the minimum cardinality covering $\mathcal{P}$.

Our algorithm is applied to the following density constrained minimum wireless communication cover problem, where a density constraint says that there are at most k points in any unit area in the plane. This is a practical constraint as the density in a wireless communication network must be limited by an upper bound for quality of service. Note that this problem does not require each point to have integer coordinates.
Density constrained minimum wireless communication cover problem: Given a set of disks $\mathcal{D}$, a set of points $\mathcal{P}$ and the density constraint in the Euclidean plane, to compute a subset of $\mathcal{D}$ of the minimum cardinality covering $\mathcal{P}$.

The weighted version of the minimum disk cover problem is also considered, where each disk is associated with a positive real number as its *weight*. The problem asks to compute a disk of minimum weight to cover all the points where the weight of a set of disks is defined as the sum of weights in the disks. It is certainly NP-hard since it contains the minimum disk cover problem as a special case (with all disks having equal weight).

Minimum weight disk cover problem: Given a set of disks $\mathcal{D}$ of constantly bounded radii and each disk is associated with a weight, a set $\mathcal{P}$ of integer coordinate points in the Euclidean plane, to compute a subset of $\mathcal{D}$ of the minimum total weight covering $\mathcal{P}$.

Finally, a practical wireless communication quality problem is considered. The signal quality received at a network node $p \in \mathcal{P}$ is inversely proportional to its distance to the transceiver on the host which is the center of the disk covering p. Thus, optimizing signal quality can be roughly treated as minimizing node-to-center distances.

Given a pair of disk D and a point p such that D covers p, the communication weight $cw(D, p)$ is defined as $d(c_D, p)$, where c_D is the center of D. Given a set of disks $\mathcal{D}$ and the points $\mathcal{P}$ covered by them, the total communication weight is computed as the sum of communication weights on all the points, i.e., $\sum_{D\in\mathcal{D}} \sum_{p:d(p,c_D)\leq r_D} cw(c_D, p)$ where r_D refers to the radius of the disk D. Note that p may be counted with different weights for multiple times given a disk cover.

Minimum communication weight disk cover problem: Given a set of disks $\mathcal{D}$ of constantly bounded radii and each disk is associated with a weight, a set $\mathcal{P}$ of integer coordinate points in the Euclidean plane, to compute a subset of $\mathcal{D}$ of the minimum total communication weight covering $\mathcal{P}$.

3 The algorithms

Our algorithm is motivated from the local-neighborhood-based scheme proposed in Nieberg et al. (2008). Note that the work (Nieberg et al. 2008) considers the problem where the centers of disks can be arbitrarily chosen in the plane while in our case the disks can only be chosen from a given set. At a high level, our algorithm partitions the point set $\mathcal{P}$ and their disk covers into a few disjoint subsets such that $(1 + \epsilon)$ approximation on each subset can be computed efficiently. The difficulty lies in minimizing the interaction between disjoint sets, i.e., a disk covering part of one disjoint set needs to have minimal impact on the other disjoint sets. This difficulty is tackled by gradually expanding a subset of points until a good disk cover on them, which is slightly (i.e., $1 + \epsilon$) larger than the optimal disk cover, can be efficiently computed. After that, the subset of points will be marked and the above process is repeated for computing the second, third, ... subsets until all the points are marked. Refer to Fig. 1 for a brief illustration. The details of the algorithm are elaborated as follows.

3.1 The main algorithm

Before describing the algorithms, the notations involved in the algorithm are summarized for clarity.

- n: the number of disks in $\mathcal{D}$
- m: the number of points in $\mathcal{P}$

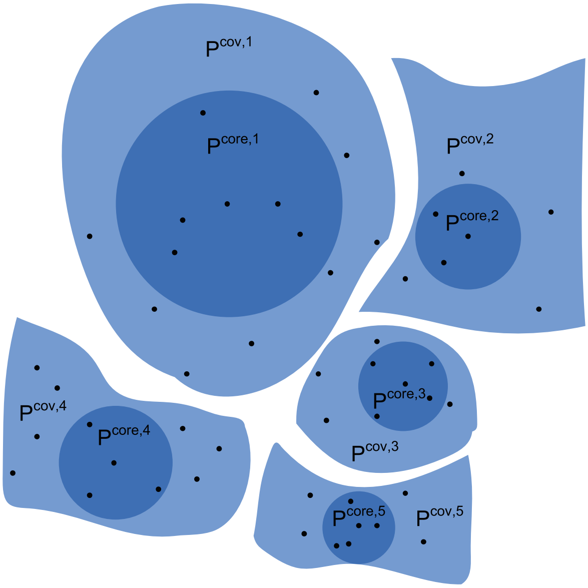

Fig. 1 A brief illustration for our algorithm. It proceeds as partitioning $\mathcal{P}$ into disjoint sets such that a $(1+\epsilon)$ approximation on each disjoint set can be computed efficiently. By our construction, each set can have at most $\mathcal{O}(\frac{1}{\epsilon^2}\log^2\frac{1}{\epsilon})$ points, and $\mathcal{P}^{cov,i}\bigcap\mathcal{P}^{cov,j}=\emptyset, \forall i\neq j$. In addition, the minimum distance between the points in any two cores $\mathcal{P}^{core,i}$ and $\mathcal{P}^{core,j}$, $\forall i\neq j$, is at least $2r$. Thus, the minimum disk cover on $\bigcup_i \mathcal{P}^{core,i}$ forms a lower bound on the minimum disk cover on $\mathcal{P}$. Note that $P^{core,i}$ and $P^{cov,i}$ denote sets of points and their boundaries shown here are only for illustration purpose

- $D_{p,i}$: a disk with p as the center and radius of $\gamma=2ri, \forall i$, where r is the radius of the given disk
- $\mathcal{P}_i$: the points located within $D_{p,i}$
- $\mathcal{D}_i^{cov}$: the minimum disk cover on $\mathcal{P}_i$
- $\mathcal{P}^{cov,j}$: $\mathcal{P}_i$ on the j-th disjoint subset of $\mathcal{P}$ if the disk growing on these points stops at i-th iteration
- $\mathcal{D}^{cov,j}$: $\mathcal{D}_i^{cov}$ on $\mathcal{P}^{cov,j}$ if disk growing stops at i-th iteration
- $\mathcal{D}^{core,j}$: $\mathcal{D}_{i-1}^{cov}$ on $\mathcal{P}^{cov,j}$ if disk growing stops at i-th iteration
- $\mathcal{P}^{core,j}$: $\mathcal{P}_{i-1}$ on $\mathcal{P}^{cov,j}$ if disk growing stops at i-th iteration

We begin with describing how to form a disjoint subset from $\mathcal{P}$ for the minimum congruent disk cover problem. For this, arbitrarily pick a point $p\in\mathcal{P}$. Let $\mathcal{P}_0=\{p\}$. The minimum number of the disks covering p is 1. Arbitrarily pick a disk covering $\mathcal{P}_0$ and denote it by $\mathcal{D}_0^{cov}$. Clearly, $|\mathcal{D}_0^{cov}|=1$.

A disk, denoted by $D_{p,1}$, is built with p as the center and radius of $\gamma=2ri, i=1$, where r is the radius of the congruent disks. It is clear that no disk outside $D_{p,1}$ can cover p and no disk covering p can cover any point outside $D_{p,1}$. Denote by $\mathcal{P}_1$ all the points within $D_{p,1}$. In our minimum congruent disk cover problem, all points in $\mathcal{P}$ are assumed to have integer coordinates. This is of interest due to its practicality in certain wireless network design where hosts may be required to locate on grid points in a uniform lattice. The algorithm on this restricted problem is also the base for the more general problems as described in Sect. 3.2. At most $\lceil\pi\gamma^2\rceil=\lceil 4\pi r^2\rceil$ points may lie inside $D_{p,1}$, i.e., $|\mathcal{P}_1|\leq\lceil 4\pi r^2\rceil$. One then computes the minimum number of disks in $\mathcal{D}$ covering $\mathcal{P}_1$. As a result, at most $\lceil 4\pi r^2\rceil$ disks are needed to cover these points. Thus, the minimum disk cover on $\mathcal{P}_1$ (which are the points within $D_{p,1}$),

Fig. 2 Growing the disk from $D_{p,i-1}$ to $D_{p,i}$

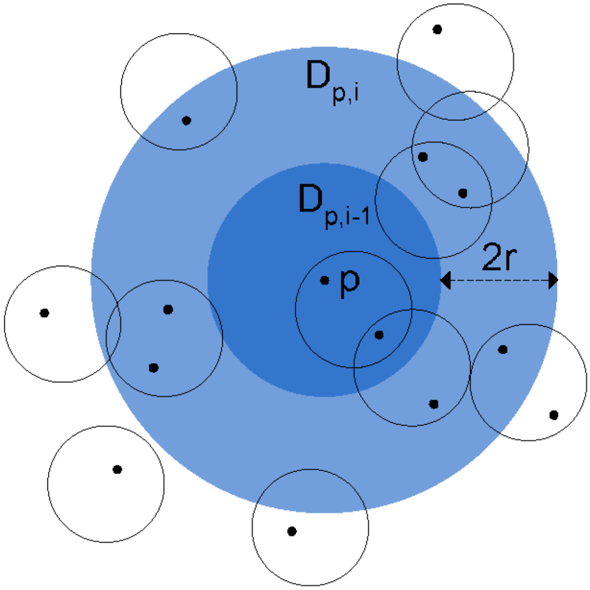

denoted by $\mathcal{D}_1^{cov}$, can be computed in $\mathcal{O}(n^{\lceil \pi\gamma^2 \rceil}) = \mathcal{O}(n^{\lceil 4\pi r^2 \rceil})$ time by enumeration and $|\mathcal{D}_1^{cov}| \leq \lceil 4\pi r^2 \rceil$.

When $|\mathcal{D}_1^{cov}| > (1+\epsilon)|\mathcal{D}_0^{cov}|$, the radius of $D_{p,1}$ is enlarged to $\gamma = 2ri, i = 2$ and denote the new disk by $D_{p,2}$. Denote by $\mathcal{P}_2$ all the points within $D_{p,2}$. At most $\lceil \pi\gamma^2 \rceil = \lceil 16\pi r^2 \rceil$ points may lie inside D_p, i.e., $|\mathcal{P}_2| \leq \lceil 16\pi r^2 \rceil$. The minimum disk cover on $\mathcal{P}_2$, denoted by $\mathcal{D}_2^{cov}$, can be computed in $\mathcal{O}(n^{\lceil \pi\gamma^2 \rceil}) = \mathcal{O}(n^{\lceil 16\pi r^2 \rceil})$ time by enumeration and $|\mathcal{D}_2^{cov}| \leq \lceil 16\pi r^2 \rceil$.

When $|\mathcal{D}_2^{cov}| > (1+\epsilon)|\mathcal{D}_1^{cov}|$, the radius of D_p is enlarged to $\gamma = 2ri, i = 3$. At most $\lceil \pi\gamma^2 \rceil = \lceil 36\pi r^2 \rceil$ points may be inside D_p and thus $|\mathcal{P}_3| \leq \lceil 36\pi r^2 \rceil$. The minimum disk cover, denoted by $\mathcal{D}_3^{cov}$, on them can be computed in $\mathcal{O}(n^{\lceil \pi\gamma^2 \rceil}) = \mathcal{O}(n^{\lceil 36\pi r^2 \rceil})$ time and $|\mathcal{D}_3^{cov}| \leq \lceil 36\pi r^2 \rceil$. Set $\gamma = 2ri, i = 4$ if $|\mathcal{D}_3^{cov}| > (1+\epsilon)|\mathcal{D}_2^{cov}|$.

Refer to Fig. 2. In general, one can see that for any $\gamma = 2ri$, the minimum disk cover $\mathcal{D}_i^{cov}$ with radius γ covering all points within $D_{p,i}$, denoted by $\mathcal{P}_i$, can be computed in $\mathcal{O}(n^{\lceil \pi\gamma^2 \rceil})$ time and $|\mathcal{D}_i^{cov}| \leq \lceil \pi\gamma^2 \rceil$. This process is repeated until $|\mathcal{D}_i^{cov}| \leq (1+\epsilon)|\mathcal{D}_{i-1}^{cov}|$ for some i. This is called *stopping criterion*.

At that moment, stop growing D_p and set $\mathcal{D}^{cov,1} = \mathcal{D}_i^{cov}$, $\mathcal{P}^{cov,1} = \mathcal{P}_i$, $\mathcal{D}^{core,1} = \mathcal{D}_{i-1}^{cov}$, and $\mathcal{P}^{core,1} = \mathcal{P}_{i-1}$. We call $\mathcal{P}^{core,\cdot}$ a *core*. Clearly, all the points in $\mathcal{P}_{i-1}$ cannot be covered by any other disk with center outside $D_{p,i}$ since the minimum distance in between is at least $2r$. Similarly, no disks in $\mathcal{D}_{i-1}^{cov}$ (which covers $\mathcal{P}_{i-1}$) can cover any other point outside $D_{p,i}$.[1] Further, $|\mathcal{D}_i^{cov}| \leq (1+\epsilon)|\mathcal{D}_{i-1}^{core}|$ guarantees that $\mathcal{D}^{cov,1} =$

[1]Note that they are in general not true when disk growing is restricted in some directions later in the algorithm. However, our algorithm only needs that no disk can cover points in different cores, which is the case since the distance between any points in different cores is at least $2r$.

$\mathcal{D}_i^{cov}$ only introduces small number (i.e., with the factor of $(1+\epsilon)$) of additional disks over the optimal disk cover for $\mathcal{P}^{cov,1}$.

The above process is terminated in at most $i \leq \min\{n, m\}$ iterations. According to the stopping criterion, for the algorithm not stopping at the i-th iteration, at least one additional disk needs to be introduced at the i-th iteration, which also means that at least one additional point is introduced compared to the $(i-1)$-th iteration. Thus, $i \leq \min\{n, m\}$.

Actually, as in Nieberg et al. (2008), one can give a much better bound on i which is $\mathcal{O}(\max\{r^2, \frac{1}{\epsilon}\log\frac{1}{\epsilon}\})$. This is due to the fact that $|\mathcal{D}_j^{cov}|$, $j = 1, 2, \ldots, i-1$ needs to be geometrically larger and $|\mathcal{D}_i^{cov}| > (1+\epsilon)^i$ by noting that $|\mathcal{D}_0^{cov}| = 1$. However, there are at most $\lceil \pi\gamma^2 \rceil = \lceil \pi(2ri)^2 \rceil$ points up to i-th iteration which gives an upper bound on the number of used disks $|\mathcal{D}_i^{cov}|$. As stated in Lemma 1, $\lceil \pi(2ri)^2 \rceil \geq |\mathcal{D}_i^{cov}| > (1+\epsilon)^i$ will soon become invalid. Thus, the runtime for computing $\mathcal{D}^{cov,1}$ is bounded by $\mathcal{O}(\sum_i n^{\mathcal{O}(4\pi r^2 i^2)}) = \mathcal{O}(n^{\mathcal{O}(\frac{1}{\epsilon^2}\log^2\frac{1}{\epsilon})})$ by noting that r is a constant.

Lemma 1 *For any integers $i > 0$ and $r > 0$ and any real number $\epsilon > 0$, $\lceil 4\pi r^2 i^2 \rceil > (1+\epsilon)^i$ is valid for i up to $\mathcal{O}(\max\{r^2, \frac{1}{\epsilon}\log\frac{1}{\epsilon}\})$.*

Proof When $\epsilon \geq 1$, for any $i > \max\{\lceil 4\pi r^2 \rceil, 15\}$, $\lceil 4\pi r^2 i^2 \rceil < (1+\epsilon)^i$ since $\lceil 4\pi r^2 i^2 \rceil < i^3 < 2^i$ when $i \geq 15$.

When $0 < \epsilon < 1$, for any $i > \max\{\lceil 4\pi r^2 \rceil, \frac{18}{\epsilon}\log\frac{1}{\epsilon}\}$, $\lceil 4\pi r^2 i^2 \rceil < i^3$. It is then sufficient to show that $(1+\epsilon)^i > i^3$, which can be seen from the following. When $i = \frac{18}{\epsilon}\log\frac{1}{\epsilon}$,

$$(1+\epsilon)^i = (1+\epsilon)^{\frac{18}{\epsilon}\log\frac{1}{\epsilon}} > 2^{\log(\frac{1}{\epsilon})^{18}} = \left(\frac{1}{\epsilon}\right)^{18}. \tag{1}$$

Note that $(\frac{1}{\epsilon})^{18} > (\frac{18}{\epsilon}\log\frac{1}{\epsilon})^3$ since

$$\left(\frac{1}{\epsilon}\right)^{18} > \left(\frac{18}{\epsilon}\log\frac{1}{\epsilon}\right)^3 \tag{2}$$

$$\Longleftrightarrow \quad \left(\frac{1}{\epsilon}\right)^6 > \frac{18}{\epsilon}\log\frac{1}{\epsilon} \tag{3}$$

$$\Longleftrightarrow \quad \left(\frac{1}{\epsilon}\right)^5 > \log\left(\frac{1}{\epsilon}\right)^{18} \tag{4}$$

which is the case since $2^{\frac{1}{\epsilon}^5} > \frac{1}{\epsilon}^{18}$ for $\frac{1}{\epsilon} > 1$. It is easy to see that the above is true for any $i > \frac{18}{\epsilon}\log\frac{1}{\epsilon}$. □

Initially, none of the points in $\mathcal{P}$ are marked. After computing $\mathcal{D}^{cov,1}$, mark all the points in $\mathcal{P}^{cov,1}$ and disallow any future partition to include them. The purpose is to separate $\mathcal{P}^{cov,1}$ from $\mathcal{P}$. In particular, it is to ensure no interaction between the core

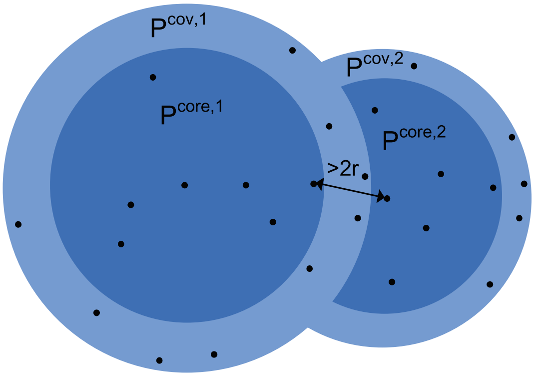

Fig. 3 Growing the disks for $\mathcal{P}^{core,2}$, $\mathcal{P}^{cov,2}$ given the disks for $\mathcal{P}^{core,1}$, $\mathcal{P}^{cov,1}$

$\mathcal{P}^{core,1}$ and any future core $\mathcal{P}^{core,i}$, $\forall i \neq 1$, i.e., no disk can simultaneously cover the points in any two cores. Subsequently, arbitrarily pick a new point $p \in \mathcal{P} \backslash \mathcal{P}^{cov,1}$ to repeat the above procedure where the only difference is that $\mathcal{P}^{cov,2}$ can only take unmarked points. That is, when an enlarged disk $D_{p,i}$ touches a marked point, its disk will not grow in that direction (and the touched point will not be included). Note that $\mathcal{D}_i^{cov}$ may still use some disks which cover marked points. However, no disk can cover points in different cores since the distance between any points in different cores is greater than $2r$ due to our construction. Refer to Fig. 3 for an example where the disks for the second core is constructed given the nearby disks for the first core. It is also clear that $|\mathcal{P}_i|$ is still bounded above by $\lceil \pi \gamma^2 \rceil = \lceil \pi (2ri)^2 \rceil$. Thus, Lemma 1 still holds and the runtime for computing $\mathcal{D}^{cov,2}$ is bounded by $\mathcal{O}(n^{\mathcal{O}(\frac{1}{\epsilon^2} \log^2 \frac{1}{\epsilon})})$. This process is iterated until all the points are marked. Refer to Fig. 1 for a brief illustration of this process. The disk cover $\bigcup_i \mathcal{D}^{cov,i}$ covers every point and is returned as the solution. Each time, at least one point will be removed and this process is iterated up to $|\mathcal{P}| = m$ times. Thus, the total runtime is bounded by $\mathcal{O}(mn^{\mathcal{O}(\frac{1}{\epsilon^2} \log^2 \frac{1}{\epsilon})})$.

The approximation ratio of the above algorithm can be bounded as follows. The minimum disk cover $\bigcup_i \mathcal{D}^{core,i}$ on $\bigcup_i \mathcal{P}^{core,i}$ forms a lower bound for the minimum disk cover on $\mathcal{P}$ since no disk can simultaneously cover points in different cores.

Observation 2 *Let $\mathcal{P}^{core,\cdot}$ be as defined above. $\mathcal{P}^{core,i} \cap \mathcal{P}^{core,j} = \emptyset, \forall i \neq j$.*

Together with the fact that $\mathcal{D}^{core,i}$ is an optimal cover on $\mathcal{P}^{core,i}$, a lower bound on the optimal solution is obtained as the union of the optimal disk cover on all cores. That is,

$$|OPT| \geq \left| \bigcup_i \mathcal{D}^{core,i} \right|, \tag{5}$$

where OPT denotes the optimal disk cover on $\mathcal{P}$. Our solution, denoted by ALG, is $\bigcup_i \mathcal{D}^{cov,i}$ which is a $(1+\epsilon)$ approximation on OPT. This is true since

$$|\mathcal{D}^{cov,i}| < (1+\epsilon)|\mathcal{D}^{core,i}|, \forall i \tag{6}$$

and

$$|ALG| = \left|\bigcup_i \mathcal{D}^{cov,i}\right| < (1+\epsilon)\sum_i |\mathcal{D}^{core,i}| = (1+\epsilon)\left|\bigcup_i \mathcal{D}^{core,i}\right|. \tag{7}$$

Thus, $|ALG| < (1+\epsilon)|OPT|$. We reach the main theorem of this paper.

Theorem 3 *The minimum congruent disk cover can be approximated within a factor of* $(1+\epsilon)$ *in* $\mathcal{O}(mn^{\mathcal{O}(\frac{1}{\epsilon^2}\log^2\frac{1}{\epsilon})})$ *time for any* $\epsilon > 0$, *where* n *is the number of disks and* m *is the number of points.*

3.2 Extensions

The above algorithm can be immediately generalized to the case of incongruent disks where the maximum radius of all disks in $\mathcal{D}$ is assumed to be bounded by a constant r_{max}. The algorithm works the same as above except that each time the radius of the disk $D_{p,i}$ is set to $\gamma = 2r_{max}i$ in contrast to $\gamma = 2ri$ in the congruent disk case. We reach the following theorem.

Theorem 4 *The minimum incongruent disk cover can be approximated within a factor of* $(1+\epsilon)$ *in* $\mathcal{O}(mn^{\mathcal{O}(\frac{1}{\epsilon^2}\log^2\frac{1}{\epsilon})})$ *time for any* $\epsilon > 0$, *where* n *is the number of disks and* m *is the number of points.*

Our algorithm can be directly applied to a practical problem of computing the density constrained minimum wireless communication cover, where the disk radius is bounded above by a constant and the density, i.e., the maximum number of points in any unit square, is bounded above by a constant k. The only impact to the above algorithm is that in disk $D_{p,i}$, previously there are at most $\lceil \pi\gamma^2 \rceil = \lceil 4\pi r^2 i^2 \rceil$ points while there are $\lceil 4\pi r^2 i^2/k \rceil$ points now even if the points are not restricted to have integer-valued coordinates. However, this will not impact the time complexity since one still has $\mathcal{O}(i^2)$ points by noting that k is a constant. Theorem 5 is immediate.

Theorem 5 *The density constrained minimum wireless communication cover can be approximated within a factor of* $(1+\epsilon)$ *in* $\mathcal{O}(mn^{\mathcal{O}(\frac{1}{\epsilon^2}\log^2\frac{1}{\epsilon})})$ *time for any* $\epsilon > 0$, *where* n *is the number of disks and* m *is the number of points.*

The algorithm can also be extended to handle the weighted version of the minimum disk cover problem. In the new problem, each disk in $D \in \mathcal{D}$ is assigned a weight and the problem asks to compute a set of disks with the minimum total weight covering $\mathcal{P}$.

Denote by $\underline{w}$ the smallest disk weight and by $\overline{w}$ the largest disk weight in all disks. Recall that $w(\cdot)$ denote the weight of a set of disks. The algorithm proceeds as before by arbitrarily picking a point and growing the disk centered at it. When growing a disk $D_{p,i}$ centered at p, one still has at most $\lceil 4\pi r^2 i^2 \rceil$ points in $D_{p,i}$ and needs $\mathcal{O}(n^{\lceil 4\pi r^2 i^2 \rceil})$ to compute the minimum weight disk cover by enumeration. Note that the weight of any such disk is bounded by $\lceil 4\pi r^2 i^2 \rceil \overline{w}$ which is the largest possible weights on the points. For any point p, denote by $D_{min}(p)$ the smallest weight disk covering p, thus $w(D_{min}(p)) \geq \underline{w}$. The only modification in the algorithm is the stopping criterion for disk growing. It is changed from $|\mathcal{D}_i^{cov}| \leq (1+\epsilon)|\mathcal{D}_{i-1}^{cov}|$ to $w(\mathcal{D}_i^{cov}) \leq (1+\epsilon) w(\mathcal{D}_{i-1}^{cov})$.

As before, not stopping at i-th iteration implies that $w(\mathcal{D}_i^{cov}) > (1+\epsilon)^i \times w(D_{min}(p))$ due to the fact that $w(\mathcal{D}_j^{cov})$, $j = 1, 2, \ldots, i-1$ is geometrically larger. Since there are at most $\lceil \pi(2ri)^2 \rceil$ points up to i-th iteration, $w(\mathcal{D}_i^{cov}) \leq |\mathcal{D}_i^{cov}| \overline{w} \leq \lceil 4\pi r^2 i^2 \rceil \overline{w}$. Together with $(1+\epsilon)^i w(D_{min}(p)) \geq (1+\epsilon)^i \underline{w}$, one only needs to bound the largest i for $\lceil 4\pi r^2 i^2 \rceil \overline{w} > (1+\epsilon)^i \underline{w}$ to be valid. If the ratio $\frac{\overline{w}}{\underline{w}}$ is bounded by a constant, it is easy to see that $i = \mathcal{O}(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$ by an analysis similar to Lemma 1. Assuming that $\frac{\overline{w}}{\underline{w}}$ is polynomial in n, i.e., $\frac{\overline{w}}{\underline{w}} = \mathcal{O}(n^{\mathcal{O}(1)})$ one can still obtain a quasi polynomial time approximation scheme.

Lemma 6 *For any positive integers $i \leq n$ and r, and any positive real numbers ϵ, $\overline{w}$, and $\underline{w}$ such that $\frac{\overline{w}}{\underline{w}} = \mathcal{O}(n^{\mathcal{O}(1)})$, $\lceil 4\pi r^2 i^2 \rceil \overline{w} > (1+\epsilon)^i \underline{w}$ is valid for i up to $\mathcal{O}(\max\{r^2, \frac{1}{\epsilon} \log n\})$.*

Proof Without loss of generality, assume $\frac{\overline{w}}{\underline{w}} \leq n^c$ for some constant c. When $\epsilon \geq 1$, for any $i > \max\{\lceil 4\pi r^2 \rceil, (c+3)\log n\}$, $\lceil 4\pi r^2 i^2 \rceil n^c < (1+\epsilon)^i$ since $\lceil 4\pi r^2 i^2 \rceil n^c < i^3 n^c \leq n^3 n^c = 2^{\log n^{c+3}}$ since $i \leq n$.

When $0 < \epsilon < 1$, for any $i > \max\{\lceil 4\pi r^2 \rceil, \frac{c+3}{\epsilon} \log n\}$, $\lceil 4\pi r^2 i^2 \rceil n^c < i^3 n^c \leq n^{c+3}$. One can see that $n^{c+3} < (1+\epsilon)^i$ when $i > \frac{c+3}{\epsilon} \log n$ since

$$(1+\epsilon)^i > (1+\epsilon)^{\frac{c+3}{\epsilon} \log n} > 2^{\log n^{c+3}} = n^{c+3}. \tag{8}$$

□

According to Lemma 6, growing a disk takes $\mathcal{O}(\sum_i n^{\mathcal{O}(4\pi r^2 i^2)}) = \mathcal{O}(n^{\frac{1}{\epsilon^2} \log^2 n})$ time. After that, mark all the points involved. Since each time at least one point is marked, the total runtime is bounded by $\mathcal{O}(mn^{\mathcal{O}(\frac{1}{\epsilon^2} \log^2 n)})$. The algorithm follows the one in Sect. 3.1 except for disk growing which has been described as above. At last, the disk cover $\bigcup_i \mathcal{D}^{cov,i}$ covers every point and is returned as the solution. Each time, at least one point will be removed and this process is iterated up to $|\mathcal{P}| = m$ times. Thus, the total runtime is bounded by $\mathcal{O}(mn^{\mathcal{O}(\frac{1}{\epsilon^2} \log^2 \frac{1}{\epsilon})})$.

The approximation ratio can be bounded as follows. The minimum weighted disk cover $\bigcup_i \mathcal{D}^{core,i}$ on $\bigcup_i \mathcal{P}^{core,i}$ forms a lower bound for the minimum disk cover on $\mathcal{P}$ since no disk can simultaneously cover points in different cores. In addition, Observation 2 is still valid. Together with the fact that $\mathcal{D}^{core,i}$ is an optimal cover on $\mathcal{P}^{core,i}$, a lower bound on the optimal solution is obtained as the union of the optimal disk cover on all cores. That is, $w(OPT) \geq w(\bigcup_i \mathcal{D}^{core,i})$. Our solution, denoted by ALG, is $\bigcup_i \mathcal{D}^{cov,i}$ which is a $(1+\epsilon)$ approximation on OPT. This is true since $w(\mathcal{D}^{cov,i}) < (1+\epsilon) w(\mathcal{D}^{core,i})$, $\forall i$ and $w(ALG) = w(\bigcup_i \mathcal{D}^{cov,i}) < (1+\epsilon) \sum_i w(\mathcal{D}^{core,i}) = (1+\epsilon) w(\bigcup_i \mathcal{D}^{core,i})$. Thus, $w(ALG) < (1+\epsilon) w(OPT)$. As before, the algorithm can be generalized to the case of incongruent disks as long as the maximum radius of all disks in $\mathcal{D}$ is bounded by a constant r_{max}. The algorithm works the same as the congruent disk case except that each time the radius of the disk $D_{p,i}$ is set to $\gamma = 2r_{max}i$ in contrast to $\gamma = 2ri$ in the congruent disk case. We reach the following theorem.

Theorem 7 *The minimum weighted disk cover can be approximated within a factor of* $(1+\epsilon)$ *in* $\mathcal{O}(mn^{\mathcal{O}(\frac{1}{\epsilon^2}\log^2 n)})$ *time for any* $\epsilon > 0$ *when the ratio of the largest to smallest disk weight is polynomially bounded by* n, *where* n *is the number of disks and* m *is the number of points. In particular, if the ratio is bounded by a constant, the algorithm runs in* $\mathcal{O}(mn^{\mathcal{O}(\frac{1}{\epsilon^2}\log^2 \frac{1}{\epsilon})})$ *time.*

An application of Theorem 7 is to handle a practical minimum communication weight disk cover problem. Since the signal quality received at a network node $p \in \mathcal{P}$ is inversely proportional to its distance to the transceiver on the host which is the center of the disk covering p. Optimizing signal quality can be roughly treated as minimizing node-to-center distances. As such, for any point p covered by a disk D, the communication weight of p is defined as $d(c_D, p)$, where c_D is the center of D. This is different from the minimum weight disk cover problem since the weight is defined on points but not disks. The minimum communication weight disk cover problem asks to compute a disk cover such that each point in $\mathcal{P}$ is covered and the total communication weight, defined as the sum of communication weights on all the points, is minimized. Although a point may be counted with different weights for multiple times in a disk cover, one can see that minimizing the total communication weight implies minimizing, to some degree, the number of disks in the cover.

The communication weights on points are first transformed to the weights on disks. Define the weight of a disk to be the sum of the weights of all points it covers, i.e., $w(D) = \sum_{p:d(p,C_D)\le r} d(p, c_D)$. Since every disk radius is a constant and every point has integer coordinates, a disk can cover at most $\lceil 4\pi r^2 i^2 \rceil$ points. The minimum weighted disk cover on them can still be found in $\mathcal{O}(n^{\lceil 4\pi r^2 i^2 \rceil})$ time since our algorithm performs enumeration. As before, one needs to bound the largest i for $w(\mathcal{D}_i^{cov}) > (1+\epsilon)^i w(D_{min}(p))$ to be valid. The minimum weight $\underline{w} \ge 0$ and the largest weight $\overline{w} \le \lceil \pi r^2 \rceil r$ by noting that r is the disk radius. The observation is that when $w(\mathcal{D}_0^{cov}) = w(D_{min}(p)) = 0$, if the minimum weight cover for $\mathcal{D}_1^{cov}$ also has zero weight, the disk growing stops since $w(\mathcal{D}_1^{cov}) \le (1+\epsilon)w(\mathcal{D}_0^{cov})$. Otherwise, $w(\mathcal{D}_1^{cov}) \ge 1$. Subsequently, disk growing requires $w(\mathcal{D}_i^{cov}) > (1+\epsilon)^{i-1} \cdot 1$. On the other hand, the largest weight for any disk is always bounded by $\overline{w} = \lceil \pi r^2 \rceil r$. Thus, disk growing implies $\lceil 4\pi r^2 i^2 \rceil \lceil \pi r^2 \rceil r > (1+\epsilon)^{i-1}$ which is valid for i up to $\mathcal{O}(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$ as can be easily seen by an analysis similar to Lemma 1. We reach Theorem 8.

Theorem 8 *The minimum communication weight disk cover can be approximated within a factor of* $(1+\epsilon)$ *in* $\mathcal{O}(mn^{\mathcal{O}(\frac{1}{\epsilon^2}\log^2 \frac{1}{\epsilon})})$ *time for any* $\epsilon > 0$, *where* n *is the number of disks and* m *is the number of points.*

4 Conclusion

This paper studies the planar minimum disk cover problem and proposes an algorithm yielding a polynomial time approximation scheme. The new algorithm approximates the optimal disk cover within a factor of $(1+\epsilon)$ and runs in $\mathcal{O}(mn^{\mathcal{O}(\frac{1}{\epsilon^2}\log^2\frac{1}{\epsilon})})$ time. This work presents the first polynomial time approximation scheme for the minimum disk cover problem where the best known algorithm can approximate the optimal solution with a large constant factor. Our algorithm is also extended to handle incongruent disks. Further, it can be directly applied to the density constrained minimum wireless communication cover problem. A natural generalization is to consider a weighted version of the minimum disk cover problem. For this, our algorithm can be extended to approximate the minimum weight disk cover within a factor of $(1+\epsilon)$ in $\mathcal{O}(mn^{\mathcal{O}(\frac{1}{\epsilon^2}\log^2 n)})$ time provided that the ratio between the largest and smallest disk weights is polynomially bounded. The runtime becomes $\mathcal{O}(mn^{\mathcal{O}(\frac{1}{\epsilon^2}\log^2\frac{1}{\epsilon})})$ when the ratio is constantly bounded. This quasi polynomial time approximation algorithm for the minimum weight disk cover problem indicates the existence of a polynomial time approximation algorithm unless $NP \subseteq DTIME[n^{polylog(n)}]$. An interesting future work would be to design such an algorithm.

References

Ambühl C, Erlebach T, Mihalák M, Nunkesser M (2006) Constant-factor approximation for minimum-weight (connected) dominating sets in unit disk graphs. In: Proceedings of the 9th international workshop on approximation, randomization, and combinatorial optimization, pp 3–14

Baker BS (1994) Approximation algorithms for np-complete problems on planar graphs. J ACM 41(1):153–180

Bronnimann H, Goodrich MT (1995) Almost optimal set covers in finite VC-dimension. Discrete Comput Geom 14(4):463–479

Calinescu G, Mandoiu II, Wan P-J, Zelikovsky AZ (2004) Selecting forwarding neighbors in wireless ad hoc networks. ACM Mob Netw Appl 9:101–111

Cheng X, Huang X, Li D, Wu W, Du D-Z (2003) A polynomial-time approximation scheme for the minimum-connected dominating set in ad hoc wireless networks. Networks 42(4):202–208

Clark BN, Colburn CJ, Johnson DS (1990) Unit disk graphs. Discrete Math 86:165–177

Garey MR, Johnson DS (1979) Computers and intractability: a guide to the theory of NP-completeness

Hochbaum DS, Maass W (1985) Approximation schemes for covering and packing problems in image processing and VLSI. J ACM 32(1):130–136

Huang Y, Gao X, Zhang Z, Wu W (2008) A better constant-factor approximation for weighted dominating set in unit disk graph. J Comb Optim

Hunt III HB, Marathe MV, Radhakrishnan V, Ravi SS, Rosenkrantz DJ, Stearns RE (1998) NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs. J Algorithms 26(2):238–274

Johnson DS (1982) NP-completeness columns: an onging guide. J Algorithms 3(2):182–195

Masuyama S, Ibaraki T, Hasegawa T (1981) The computational complexity of the m-center problems on the plane. IEICE Trans E64:57–64

Ni S-Y, Tseng Y-C, Chen Y-S, Sheu J-P (1999) The broadcast storm problem in a mobile ad hoc network. In: Proceedings of the 5th ACM/IEEE international conference on mobile computing and networking, pp 151–162

Nieberg T, Hurink J, Kern W (2008) Approximation schemes for wireless networks. ACM Trans Algorithms 4(4):49:1–49:17

Shang W, Wan P, Yao F, Hu X (2007) Algorithms for minimum m-connected k-tuple dominating set problem. Theor Comput Sci 381(1–3):241–247
Shang W, Yao F, Wan P, Hu X (2008) On minimum m-connected k-dominating set problem in unit disc graphs. J Comb Optim 16(2):99–106
Sun M-T, Yi C-W, Yang C-K, Lai T-H (2007) An optimal algorithm for the minimum disc cover problem. Algorithmica 50(1):58–71