

Linear Time Approximation Schemes for Geometric Maximum Coverage

Jian Li[†], Haitao Wang[‡], Bowei Zhang[†], and Ningye Zhang[†]

[†] Institute for Interdisciplinary Information Sciences (IIIS),
Tsinghua University, Beijing, China, 100084

[‡] Department of Computer Science, Utah State University, Utah, USA, 84322
{lijian83@mail, bw-zh14@mails, zhangny12@mails}.tsinghua.edu.cn
haitao.wang@usu.edu

Abstract. We study approximation algorithms for the following geometric version of the maximum coverage problem: Let \mathcal{P} be a set of n weighted points in the plane. We want to place m $a \times b$ rectangles such that the sum of the weights of the points in \mathcal{P} covered by these rectangles is maximized. For any fixed $\varepsilon > 0$, we present efficient approximation schemes that can find a $(1 - \varepsilon)$ -approximation to the optimal solution. In particular, for $m = 1$, our algorithm runs in linear time $O(n \log(\frac{1}{\varepsilon}))$, improving over the previous result. For $m > 1$, we present an algorithm that runs in $O(\frac{n}{\varepsilon} \log(\frac{1}{\varepsilon}) + m(\frac{1}{\varepsilon})^{O(\min(\sqrt{m}, \frac{1}{\varepsilon})})$ time.

Keywords: Maximum Coverage, Geometric Set Cover, Polynomial-Time Approximation Scheme

1 Introduction

The maximum coverage problem is a classic problem in theoretical computer science and combinatorial optimization. In this problem, we are given a universe \mathcal{P} of weighted elements, a family of subsets and a number k . The goal is to select at most k of these subsets such that the sum of the weights of the covered elements in \mathcal{P} is maximized. It is well-known that the most natural greedy algorithm achieves an approximation factor of $1 - 1/e$, which is essentially optimal (unless $P=NP$) [22, 27, 18]. However, for several geometric versions of the maximum coverage problem, better approximation ratios can be achieved (we will mention some of such results below). In this paper, we mainly consider the following geometric maximum coverage problem:

Definition 1. ($\text{MaxCov}_R(\mathcal{P}, m)$) Let \mathcal{P} be a set of n points in a 2-dimensional Euclidean plane \mathbb{R}^2 . Each point $p \in \mathcal{P}$ has a given weight $w_p \geq 0$. The goal of our geometric max-coverage problem (denoted as $\text{MaxCov}_R(\mathcal{P}, m)$) is to place m $a \times b$ rectangles such that the sum of the weights of the covered points by these rectangles is maximized. More precisely, let S be the union of m rectangles we placed. Our goal is to maximize

$$\text{Cover}(\mathcal{P}, S) = \sum_{p \in \mathcal{P} \cap S} w_p.$$

We also study the same coverage problem with unit disks, instead of rectangles. We denote the corresponding problem as $\text{MaxCov}_D(\mathcal{P}, m)$.

One natural application of the geometric maximum coverage problem is the facility placement problem. In this problem, we would like to locate a certain number of facilities to serve the maximum number of clients. Each facility can serve a region (depending on whether the metric is L_1 or L_2 , the region is either a square or a disk).

1.1 $m = 1$

Previous Results: We first consider $\text{MaxCov}_R(\mathcal{P}, 1)$. Imai and Asano [23], Nandy and Bhattacharya [26] gave two different exact algorithms for $\text{MaxCov}_R(\mathcal{P}, 1)$, both running in time $O(n \log n)$. It is also known that solving $\text{MaxCov}_R(\mathcal{P}, 1)$ exactly in algebraic decision tree model requires $\Omega(n \log n)$ time [5]. Tao et al. [28] proposed a randomized approximation scheme for $\text{MaxCov}_R(\mathcal{P}, 1)$. With probability $1 - 1/n$, their algorithm returns a $(1 - \varepsilon)$ -approximate answer in $O(n \log(\frac{1}{\varepsilon}) + n \log \log n)$ time. In the same paper, they also studied the problem in the external memory model.

Our Results: For $\text{MaxCov}_R(\mathcal{P}, 1)$ we show that there is an approximation scheme that produces a $(1 - \varepsilon)$ -approximation and runs in $O(n \log(\frac{1}{\varepsilon}))$ time, improving the result by Tao et al. [28].

1.2 General $m > 1$

Previous Results: Both $\text{MaxCov}_R(\mathcal{P}, m)$ and $\text{MaxCov}_D(\mathcal{P}, m)$ are NP-hard if m is part of the input [24]. The most related work is Berg, Cabello and Har-Peled [13]. They mainly focused on using unit disks (i.e., $\text{MaxCov}_D(\mathcal{P}, m)$). They proposed a $(1 - \varepsilon)$ -approximation algorithm for $\text{MaxCov}_D(\mathcal{P}, m)$ with time complexity $O(n \log n + n(m/\varepsilon)^{O(\sqrt{m})})$.

¹ We note that their algorithm can be easily extended to MaxCov_R with the same time complexity.

We are not aware of any explicit result for $\text{MaxCov}_R(\mathcal{P}, m)$ for general $m > 1$. It is known [13] that the problem admits a PTAS via the standard shifting technique [21].²

Our Results: Our main result is an approximation scheme for $\text{MaxCov}_R(\mathcal{P}, m)$ which runs in time

$$O\left(\frac{n}{\varepsilon} \log \frac{1}{\varepsilon} + m \left(\frac{1}{\varepsilon}\right)^\Delta\right),$$

¹ They were mainly interested in the case where m is a constant. So the running time becomes $O(n \log n + n(1/\varepsilon)^{O(\sqrt{m})})$ (which is the bound claimed in their paper) and the exponential dependency on m does not look too bad for $m = O(1)$. Since we consider the more general case, we make the dependency on m explicit.

² Hochbaum and Maass [21] obtained a PTAS for the problem of covering given points with a minimal number of rectangles. Their algorithm can be easily modified into a PTAS for $\text{MaxCov}_R(\mathcal{P}, m)$ with running time $n^{O(1/\varepsilon)}$.

where $\Delta = O(\min(\sqrt{m}, \frac{1}{\varepsilon}))$. Our algorithm can be easily extended to other shapes. In Appendix D, we sketch an extension for approximating $\text{MaxCov}_D(\mathcal{P}, m)$. The running time of our algorithm is

$$O\left(n\left(\frac{1}{\varepsilon}\right)^{O(1)} + m\left(\frac{1}{\varepsilon}\right)^\Delta\right).$$

We want to mention that n is not involved in the second term. Following the convention of approximation algorithms, ε is a fixed constant. Hence, the second term is essentially $O(m)$ and the overall running time is essentially linear $O(n)$. So even if m is a constant, the running time of our algorithm is much better than that in [13]. Our algorithm follows the standard shifting technique [21], which reduces the problem to a smaller problem restricted in a constant size cell. The same technique is also used in Berg et al. [13]. They proceeded by first solving the problem exactly in each cell, and then use dynamic programming to find the optimal allocation for all cells.³ Our improvement comes from another two simple yet useful ideas. First, we notice that we only have to solve the problem in each cell approximately (this is done by taking an ε -approximation again). Second, we solve the dynamic program approximately. In fact, we show that a simple greedy strategy (along with some additional observations) can be used for this purpose, which allows us to save another $O(m)$ term.

1.3 Other Related Work

There are many different variants for this problem. We mention some most related problems here.

Barequet et al. [4], Dickerson and Scharstein [14] studied the max-enclosing polygon problem which aims to find a position of a given polygon to cover maximum number of points. This is the same as $\text{MaxCov}_R(\mathcal{P}, 1)$ if a polygon is a rectangle. Imai et al. [23] gave an optimal algorithm for the max-enclosing rectangle problem with time complexity $O(n \log n)$.

$\text{MaxCov}_D(\mathcal{P}, m)$ was introduced by Drezner [16]. Chazelle and Lee [10] gave an $O(n^2)$ -time exact algorithm for the problem $\text{MaxCov}_D(\mathcal{P}, 1)$. A Monte-Carlo $(1 - \varepsilon)$ -approximation algorithm for $\text{MaxCov}_D(\mathcal{P}, 1)$ was shown in [1], where \mathcal{P} is an unweighted point set. Aronov and Harpeled [3] showed that for unweighted point sets an $O(n\varepsilon^{-2} \log n)$ time Monte-Carlo $(1 - \varepsilon)$ -approximation algorithm exists, and also provided some results for other shapes. Berg et al. [13] provided an $O(n \log n + n\varepsilon^{-3})$ time deterministic $(1 - \varepsilon)$ -approximation algorithm.

For $m > 1$, $\text{MaxCov}_D(\mathcal{P}, m)$ has only a few results. For $m = 2$, Cabello et al. [8] gave an exact algorithm for this problem when the two disks are disjoint in $O(n^{8/3} \log^2 n)$ time. Berg et al. [13] gave $(1 - \varepsilon)$ -approximation algorithms that runs in $O(n \log n + n\varepsilon^{-4m+4} \log^{2m-1}(1/\varepsilon))$ time for $m > 3$ and in $O(n \log n + n\varepsilon^{-6m+6} \log(1/\varepsilon))$ time for $m = 2, 3$.

³ In fact, their dynamic programming runs in time at least $\Omega(m^2)$. Since they focused on constant m , this term is negligible in their running time. But if $m > \sqrt{n}$, the term can not be ignored and may become the dominating term.

The dual of the maximum coverage problem is the classical set cover problem. The geometric set cover problem has enjoyed extensive study in the past two decades. The literature is too vast to list exhaustively here. See e.g., [7, 11, 17, 25, 29, 9] and the references therein.

2 Preliminaries

We first define some notations and mention some results that are needed in our algorithm. Denote by $G_\delta(a, b)$ the square grid with mesh size δ such that the vertical and horizontal lines are defined as follows

$$G_\delta(a, b) = \{(x, y) \in \mathbb{R}^2 \mid y = b + k \cdot \delta, k \in \mathbb{Z}\} \cup \{(x, y) \in \mathbb{R}^2 \mid x = a + k \cdot \delta, k \in \mathbb{Z}\}$$

Given $G_\delta(a, b)$ and a point $p = (x, y)$, we call the integer pair $(\lfloor x/\delta \rfloor, \lfloor y/\delta \rfloor)$ the *index* of p (the index of the cell in which p lies in).

Perfect Hashing: M. Dietzfelbinger et.al [15] shows that if each basic algebraic operation (including $\{+, -, \times, \div, \log_2, \exp_2\}$) can be done in constant time, we can get a perfect hash family so that each insertion and membership query takes $O(1)$ expected time. In particular, using this hashing scheme, we can hash the indices of all points, so that we can obtain the list of all non-empty cells in $O(n)$ expected time. Moreover, for any non-empty cell, we can retrieve all points lies in it in time linear in the number of such points.

Linear Time Weighted Median and Selection: It is well known that finding the weighted median for an array of numbers can be done in deterministic worst-case linear time. The setting is as follows: Given n distinct elements x_1, x_2, \dots, x_n with positive weights w_1, w_2, \dots, w_n . Let $w = \sum_{i=1}^n w_i$. The *weighted median* is the element x_k satisfying $\sum_{x_i < x_k} w_i < w/2$ and $\sum_{x_i > x_k} w_i \leq w/2$. Finding the k th smallest elements for any array can also be done in deterministic worst-case linear time. See e.g. [12].

An Exact Algorithm for $\text{MaxCov}_R(\mathcal{P}, 1)$: Nandy and Bhattacharya [26] gives an $O(n \log n)$ exact algorithm for the $\text{MaxCov}_R(\mathcal{P}, 1)$ problem.

3 A Linear Time Algorithm for $\text{MaxCov}_R(\mathcal{P}, 1)$

Notations: Without loss of generality, we can assume that $a = b = 1$, i.e., all the rectangles are 1×1 squares, (by properly scaling the input). We also assume that all points are in general positions. In particular, all coordinates of all points are distinct. For a unit square r , we use $w(r)$ to denote the sum of the weights of the points covered by r . We say a unit square r is located at (x, y) if the top-left corner of r is (x, y) .

Now we present our approximation algorithm for $\text{MaxCov}_R(\mathcal{P}, 1)$.

3.1 Grid Shifting

Recall the definition of a grid $G_\delta(a, b)$ (in Section 2). Consider the following four grids: $G_2(0, 0), G_2(0, 1), G_2(1, 0), G_2(1, 1)$ with $\delta = 2$. We can easily see that for

any unit square r , there exists one of the above grids that does not intersect r (i.e., r is inside some cell of the grid). This is also the case for the optimal solution.

Now, we describe the overall framework, which is similar to that in [28]. Our algorithm differs in several details. $\text{MAXCOVCELL}(c)$ is a subroutine that takes a 2×2 cell c as input and returns a unit square r that is a $(1-\varepsilon)$ -approximate solution if the problem is restricted to cell c . We present the details of MAXCOVCELL in the next subsection.

Algorithm 1 $\text{MaxCov}_R(\mathcal{P}, 1)$

```

 $w_{\max} \leftarrow 0$ 
for each  $G \in \{G_2(0,0), G_2(0,1), G_2(1,0), G_2(1,1)\}$  do
    Use perfect hashing to find all the non-empty cells of  $G$ .
    for each non-empty cell  $c$  of  $G$  do
         $r \leftarrow \text{MAXCOVCELL}(c)$ .
        If  $w(r) > w_{\max}$ , then  $w_{\max} \leftarrow w(r)$  and  $r_{\max} \leftarrow r$ .
    end for;
end for;
return  $r_{\max}$ ;
    
```

As we argued above, there exists a grid G such that the optimal solution is inside some cell $c^* \in G$. Therefore, $\text{MAXCOVCELL}(c^*)$ should return a $(1-\varepsilon)$ -approximation for the original problem $\text{MaxCov}_R(\mathcal{P}, 1)$.

3.2 MaxCovCell

In this section, we present the details of the subroutine MAXCOVCELL . Now we are dealing with the problem restricted to a single 2×2 cell c . Denote the number of point in c by n_c , and the sum of the weights of points in c by W_c . We distinguish two cases, depending on whether n_c is larger or smaller than $(\frac{1}{\varepsilon})^2$. If $n_c < (\frac{1}{\varepsilon})^2$, we simply apply the $O(n \log n)$ time exact algorithm. [26]

The other case requires more work. In this case, we further partition cell c into many smaller cells. First, we need the following simple lemma.

Lemma 1. *Given n points in \mathbb{R}^2 with positive weights w_1, w_2, \dots, w_n , $\sum_{i=1}^n w_i = w$. Assume that x_1, x_2, \dots, x_n are their distinct x -coordinates. We are also given a value w_d such that $\max(w_1, w_2, \dots, w_n) \leq w_d \leq w$, Then, we can find at most $2w/w_d$ vertical lines such that the sum of the weights of points strictly between (we do not count the points on these lines) any two adjacent lines is at most w_d in time $O(n \log(w/w_d))$.*

Proof. See Algorithm 2. In this algorithm, we apply the weighted median algorithm recursively. Initially we have a global variable $\mathcal{L} = \emptyset$, which upon termination is the set of x -coordinates of the selected vertical lines. Each time we find

Algorithm 2 PARTITION($\{x_1, x_2, \dots, x_n\}$)

Find the weighted median x_k (w.r.t. w -weight);
 $\mathcal{L} = \mathcal{L} \cup \{x_k\}$;
Generate $S = \{x_i \mid w_i < x_k\}$, $L = \{x_i \mid w_i > x_k\}$;
If the sum of the weights of the points in S is larger than w_d , run PARTITION(S);
If the sum of the weights of the points in L is larger than w_d , run PARTITION(L);

the weighted median x_k and separate the point with the vertical line $x = x_k$, which we add into \mathcal{L} . The sum of the weights of points in either side is at most half of the sum of the weights of all the points. Hence, the depth of the recursion is at most $\lceil \log(w/w_d) \rceil$. Thus, the size of \mathcal{L} is at most $2^{\lceil \log(w/w_d) \rceil} \leq 2w/w_d$, and the running time is $O(n \log(w/w_d))$. \square

Now, we describe how to partition cell c into smaller cells. First we partition c with some vertical lines. Let \mathcal{L}_v to denote a set of vertical lines. Initially, $\mathcal{L} = \emptyset$. Let $w_d = \frac{\varepsilon \cdot W_c}{16}$. We find all the points whose weights are at least w_d . For each such point, we add to \mathcal{L}_v the vertical line that passes through the point. Then, we apply Algorithm 2 to all the points with weights less than w_d . Next, we add a set \mathcal{L}_h of horizontal lines in exactly the same way.

Lemma 2. *The sum of the weights of points strictly between any two adjacent lines in \mathcal{L}_v is at most $w_d = \frac{\varepsilon \cdot W_c}{16}$. The number of vertical lines in \mathcal{L}_v is at most $\frac{32}{\varepsilon}$. Both statements hold for \mathcal{L}_h as well.*

Proof. The first statement is straightforward from the description of the algorithm. We only need to prove the upper bound of the number of the vertical lines. Assume the sum of the weights of those points considered in the first (resp. second) step is W_1 (resp. W_2), $W_1 + W_2 = W_c$. The number of vertical lines in \mathcal{L}_v is at most

$$W_1 / \left(\frac{\varepsilon \cdot W_c}{16} \right) + 2W_2 / \left(\frac{\varepsilon \cdot W_c}{16} \right) \leq \frac{32}{\varepsilon}.$$

The first term is due to the fact that the weight of each point we found in the first step has weight at least $\frac{\varepsilon \cdot W_c}{16}$, and the second term directly follows from Lemma 1. \square

We add both vertical boundaries of cell c into \mathcal{L}_v and both horizontal boundaries of cell c into \mathcal{L}_h . Now $\mathcal{L} = \mathcal{L}_v \cup \mathcal{L}_h$ forms a grid of size at most $(\frac{32}{\varepsilon} + 2) \times (\frac{32}{\varepsilon} + 2)$. Assume $\mathcal{L} = \{(x, y) \in \mathbb{R}^2 \mid y = y_j, j \in \{1, \dots, u\}\} \cup \{(x, y) \in \mathbb{R}^2 \mid x = x_i, i \in \{1, \dots, v\}\}$, with both $\{y_i\}$ and $\{x_i\}$ are sorted. \mathcal{L} partitions c into *small cells*. The final step of our algorithm is simply enumerating all the unit squares located at $(x_i, y_j), i \in \{1, \dots, u\}, j \in \{1, \dots, v\}$, and return the one with the maximum coverage. However, computing the coverage exactly for all these unit squares is expensive. Instead, we only calculate the weight of these unit square approximately as follows. For each unit square r , we only count the

weight of points that are in some small cell fully covered by r . Now, we show this can be done in $O\left(n_c \log\left(\frac{1}{\varepsilon}\right) + \left(\frac{1}{\varepsilon}\right)^2\right)$ time.

After sorting $\{y_i\}$ and $\{x_i\}$, we can use binary search to identify which small cell each point lies in. So we can calculate the sum of the weights of points at the interior, edges or corners of all small cells in $O(n_c \log\left(\frac{1}{\varepsilon}\right))$ times. Searching the unit square with the maximum (approximate) coverage can be done with a standard incremental algorithm. For completeness, we provide the details as follows. In fact, the problem can be reduced to the following problem: We have a $k \times k$ grid, each cell (i, j) has a given weight $w(i, j)$. We are also given two non-decreasing integer function ⁴ $X : [k] \rightarrow [k], Y : [k] \rightarrow [k]$ such that $X(i) \geq i, Y(i) \geq i$ for all $i \in [k]$. The goal is to find i^*, j^* such that

$$I(i, j) = \sum_{i \leq i' \leq X(i)} \sum_{j \leq j' \leq Y(j)} w_{i', j'}$$

is maximized. First, we define an auxiliary function $H(i, j) = \sum_{j \leq j' \leq Y(j)} w_{i, j'}$. This can be computed separately for each column. We maintain two pointers p and q , such that $q = Y(p)$ holds through the process. Initially, $p = 1$ and $q = Y(1)$. In each iteration, we increment p by 1 and q by $Y(p+1) - Y(p)$. Each $w(i, j)$ will be encountered at most twice. So we can calculate all $H(i, j)$ values in $O\left(\frac{1}{\varepsilon}\right)^2$ time. Then in the same way, We can use $H(i, j)$ values to calculate all $I(i, j)$ values in exactly the same way. This takes $O\left(\frac{1}{\varepsilon}\right)^2$ time as well. We return the maximum $I(i, j)$ value as the result for $\text{MAXCOVCELL}(c)$.

Putting everything together, we conclude that if $n_c \geq \left(\frac{1}{\varepsilon}\right)^2$, the running time of $\text{MAXCOVCELL}(c)$ is $O\left(n_c \log\left(\frac{1}{\varepsilon}\right) + \left(\frac{1}{\varepsilon}\right)^2\right)$. We can conclude the main result of this section with the following theorem. Due to space constraints, the proof is deferred to the appendix.

Theorem 1. *Algorithm 1 returns a $(1-\varepsilon)$ -approximate answer for $\text{MaxCov}_R(\mathcal{P}, 1)$ in $O(n \log\left(\frac{1}{\varepsilon}\right))$ time.*

4 Linear Time Algorithms for $\text{MaxCov}_R(\mathcal{P}, m)$

4.1 Grid Shifting

For general m , we need the shifting technique [21]. Consider grids with a different side length: $G_{6/\varepsilon}(a, b)$. We shift the grid to $\frac{6}{\varepsilon}$ different positions: $(0, 0), (1, 1), \dots, \left(\frac{6}{\varepsilon} - 1, \frac{6}{\varepsilon} - 1\right)$. (For simplicity, we assume that $\frac{1}{\varepsilon}$ is an integer and no point in \mathcal{P} has an integer coordinate, so points in \mathcal{P} will never lie on the grid line. Let

$$\mathbb{G} = \left\{G_{6/\varepsilon}(0, 0), \dots, G_{6/\varepsilon}\left(\frac{6}{\varepsilon} - 1, \frac{6}{\varepsilon} - 1\right)\right\}.$$

The following lemma is quite standard. We provide a proof in Appendix B for completeness.

⁴ $[k] = \{1, 2, \dots, k\}$

Lemma 3. *There exist $G^* \in \mathbb{G}$ and a $(1 - \frac{2\varepsilon}{3})$ -approximate solution R such that none of the unit squares in R intersects G^* .*

We present a subroutine which can approximately solve the problem for a grid, and then apply it to each grid in \mathbb{G} . We return the maximum obtained solution as our final output.

4.2 Dynamic Programming

Now consider a fixed grid $G \in \mathbb{G}$. Let c_1, \dots, c_t be the cells of grid G and Opt be the optimal solution that does not intersect G . Obviously, $(\frac{6}{\varepsilon})^2$ unit squares are enough to cover an entire $\frac{6}{\varepsilon} \times \frac{6}{\varepsilon}$ cell. Thus the maximum number of unit squares we need to place in one single cell is $m_c = \min\{m, (\frac{6}{\varepsilon})^2\}$.

Let $\text{Opt}(c_i, k)$ be the maximum weight we can cover with k unit squares in cell c_i . For each nonempty cell c_i and for each $k \in [m_c]$, we find a $(1 - \frac{\varepsilon}{3})$ -approximation $F(c_i, k)$ to $\text{Opt}(c_i, k)$. We will show how to achieve this later. Now assume that we can do it.

Let $\text{Opt}_F(m)$ be the optimal solution we can get from the values $F(c_i, k)$. More precisely,

$$\text{Opt}_F(m) = \max_{k_1, \dots, k_t \in [m_c]} \left\{ \sum_{i=1}^t F(c_i, k_i) \mid \sum_{i=1}^t k_i = m \right\} \quad (1)$$

We can see that $\text{Opt}_F(m)$ must be a $(1 - \frac{\varepsilon}{3})$ -approximation to Opt . We can easily use dynamic programming to calculate the exact value of $\text{Opt}_F(m)$. Denote by $A(i, k)$ the maximum weight we can cover with k unit squares in cells c_1, c_2, \dots, c_i . We have the following DP recursion:

$$A(i, k) = \begin{cases} \max_{j=0}^{\min(k, m_c)} \{A(i-1, k-j) + F(c_i, j)\} & \text{if } i > 1 \\ F(c_1, k) & \text{if } i = 1 \end{cases}$$

The running time of the above simple dynamic programming is $O(m^2 \cdot m_c)$. One may notice that each step of the DP is computing a $(+, \max)$ convolution. However, existing algorithms (see e.g., [6, 30]) only run slightly better than quadratic time. So the improvement would be quite marginal. But in the next section, we show that if we would like to settle for an approximation to $\text{Opt}_F(m)$, the running time can be dramatically improved to linear.

4.3 A Greedy Algorithm

We first apply our $\text{MaxCov}_R(\mathcal{P}, 1)$ algorithm in Section 3 to each cell c_i , to compute a $(1 - \frac{\varepsilon^2}{9})$ -approximation of $\text{Opt}(c_i, 1)$. Let $f(c_i, 1)$ be the return values.⁵ This takes $O(n \log \frac{1}{\varepsilon})$ time. Then, we use the selection algorithm to find out the m cells with the largest $f(c_i, 1)$ values. Assume that those cells are $c_1, \dots, c_m, c_{m+1}, \dots, c_t$, sorted from largest to smallest by $f(c_i, 1)$.

⁵ Both $f(c_i, 1)$ and $F(c_i, 1)$ are approximations of $\text{Opt}(c_i, 1)$, with slightly different approximation ratios.




Fig. 1. $F(c_i, k)$ (left) and $\widehat{F}(c_i, k)$ (right)

Lemma 4. Let Opt' be the maximum weight we can cover using m unit squares in c_1, \dots, c_m . Then $\text{Opt}' \geq (1 - \frac{\epsilon^2}{9})\text{Opt}$

Proof. Let k be the number of unit squares in Opt that are chosen from c_{m+1}, \dots, c_t . This means there must be at least k cells in $\{c_1, \dots, c_m\}$ such that Opt does not place any unit square. Therefore we can always move all k unit squares placed in c_{m+1}, \dots, c_t to these empty cells such that each empty cell contains only one unit square. Denote the weight of this modified solution by A . Obviously, $\text{Opt}' \geq A$. For any i, j such that $1 \leq i \leq m < j \leq t$, we have $\text{Opt}(c_i, 1) \geq f(c_i, 1) \geq f(c_j, 1) \geq (1 - \frac{\epsilon^2}{9})\text{Opt}(c_j, 1)$. Combining with a simple observation that $\text{Opt}(c_i, k) \leq k\text{Opt}(c_i, 1)$, we can see that $A \geq (1 - \frac{\epsilon^2}{9})\text{Opt}$. Therefore, $\text{Opt}' \geq (1 - \frac{\epsilon^2}{9})\text{Opt}$. \square

Hence, from now on, we only need to consider the first m cells $\{c_1, \dots, c_m\}$. We distinguish two cases. If $m \leq 324(\frac{1}{\epsilon})^4$, we just apply the dynamic program to c_1, \dots, c_m . The running time of the above dynamic programming is $O((\frac{1}{\epsilon})^{O(1)})$.

If $m > 324(\frac{1}{\epsilon})^4$, we can use a greedy algorithm to find an answer of weight at least $(1 - \frac{\epsilon^2}{9})\text{Opt}_F(m)$.

Let $\mathbf{b} = (\frac{6}{\epsilon})^2$. For each cell c_i , we find the upper convex hull of 2D points $\{(0, F(c_i, 0)), (1, F(c_i, 1)), \dots, (\mathbf{b}, F(c_i, \mathbf{b}))\}$. See Figure 1. Suppose the convex hull points are $\{(t_{i,0}, F(c_i, t_{i,0})), (t_{i,1}, F(c_i, t_{i,1})), \dots, (t_{i,s_i}, F(c_i, t_{i,s_i}))\}$, where $t_{i,0} = 0, t_{i,s_i} = \mathbf{b}$. For each cell, since the above points are already sorted from left to right, we can compute the convex hull in $O(\mathbf{b})$ time by Graham's scan[20]. Therefore, computing the convex hulls for all these cells takes $O(m\mathbf{b})$ time.

For each cell c_i , we maintain a value p_i representing that we are going to place t_{i,p_i} squares in cell c_i . Initially for all $i \in [m]$, $p_i = 0$. In each stage, we find the cell c_i such that current slope (the slope of the next convex hull edge)

$$\frac{F(c_i, t_{i,p_i+1}) - F(c_i, t_{i,p_i})}{t_{i,p_i+1} - t_{i,p_i}}$$

is maximized. Then we add 1 to p_i , or equivalently we assign $t_{i,p_i+1} - t_{i,p_i}$ more squares into cell c_i . We repeat this step until we have already placed at least $m - \mathbf{b}$ squares. We can always achieve this since we can place at most \mathbf{b} squares in one single cell in each iteration. Let m' the number of squares we have placed

($m = \mathbf{b} \leq m' \leq m$). For the remaining $m - m'$ squares, we allocate them arbitrarily. We denote the algorithm by GREEDY and let the value obtained be $\text{Greedy}(m')$. Having the convex hulls, the running time of the greedy algorithm is $O(m)$.

Now we analyze the performance of the greedy algorithm.

Lemma 5. *The above greedy algorithm computes an $(1 - \varepsilon^2/9)$ -approximation to $\text{Opt}_{\mathbf{F}}(m)$.*

Proof. Define an auxiliary function $\widehat{\mathbf{F}}(\mathbf{c}_i, k)$ as follows: If $k = t_{i,j}$ for some j , $\widehat{\mathbf{F}}(\mathbf{c}_i, k) = F(\mathbf{c}_i, k)$. Otherwise, suppose $t_{i,j} < k < t_{i,j+1}$, then

$$\widehat{\mathbf{F}}(\mathbf{c}_i, k) = F(\mathbf{c}_i, t_{i,j}) + \frac{F(\mathbf{c}_i, t_{i,j+1}) - F(\mathbf{c}_i, t_{i,j})}{t_{i,j+1} - t_{i,j}} \times (k - t_{i,j}).$$

Intuitively speaking, $\widehat{\mathbf{F}}(\mathbf{c}_i, k)$ (See Figure 1) is the function defined by the upper convex hull at integer points.⁶ Thus, for all $i \in [m]$, $\widehat{\mathbf{F}}(\mathbf{c}_i, k)$ is a concave function. Obviously, $\widehat{\mathbf{F}}(\mathbf{c}_i, k) \geq F(\mathbf{c}_i, k)$ for all $i \in [m]$ and all $k \in [\mathbf{b}]$.

Let $\text{Opt}_{\widehat{\mathbf{F}}}(i)$ be the optimal solution we can get from the values $\widehat{\mathbf{F}}(\mathbf{c}_i, k)$ by placing i squares. By the convexity of $\widehat{\mathbf{F}}(\mathbf{c}_i, k)$, the following greedy algorithm is optimal: as long as we still have budget, we assign 1 more square to the cell which provides the largest increment of the objective value. In fact, this greedy algorithm runs in almost the same way as GREEDY. The only difference is that GREEDY only picks an entire edge of the convex hull, while the greedy algorithm here may stop in the middle of an edge (only happen for the last edge). Since the marginal increment never increases, we can see that $\text{Opt}_{\widehat{\mathbf{F}}}(i)$ is concave.

By the way of choosing cells in our greedy algorithm, we make the following simple but important observation:

$$\text{Greedy}(m') = \text{Opt}_{\widehat{\mathbf{F}}}(m') = \text{Opt}_{\mathbf{F}}(m').$$

So, our greedy algorithm is in fact optimal for m' . Combining with $m - m' \leq \mathbf{b}$ and the concavity of $\text{Opt}_{\widehat{\mathbf{F}}}$, we can see that

$$\text{Opt}_{\widehat{\mathbf{F}}}(m') \geq \frac{m - \mathbf{b}}{m} \text{Opt}_{\widehat{\mathbf{F}}}(m) \geq \left(1 - \frac{\varepsilon^2}{9}\right) \text{Opt}_{\mathbf{F}}(m).$$

The last inequality holds because $\text{Opt}_{\widehat{\mathbf{F}}}(i) \geq \text{Opt}_{\mathbf{F}}(i)$ for any i . □

4.4 Computing $\mathbf{F}(\mathbf{c}, k)$

Now we show the subroutine MAXCOVCELLM for computing $\mathbf{F}(\mathbf{c}, k)$. We use a similar partition algorithm as Section 3.2. The only difference is that this time we need to partition the cell finer so that the maximum possible weight of points

⁶ At first sight, it may appear that $\mathbf{F}(\mathbf{c}_i, k)$ should be a concave function. However, this is not true. See Figure 4 in Appendix B for an example.

between any two adjacent parallel partition lines is $(\frac{\varepsilon^3 W_\varepsilon}{864})$. After partitioning the cell, we enumerate all the possible ways of placing k unit squares at the grid point. Similarly, for each unit square r , we only count the weight of points that are in some cell fully covered by r . The algorithm is summarized in Appendix B.

We can adapt the algorithm in [13] to enumerate these possible choices in $O((\frac{1}{\varepsilon})^\Delta)$ time where $\Delta = O(\min(\sqrt{m}, \frac{1}{\varepsilon}))$. We briefly sketch the algorithm in Appendix C. Now we prove the correctness of this algorithm.

Lemma 6. *MAXCOVCELLM returns a $(1 - \frac{\varepsilon}{3})$ approximate answer for $\text{Opt}(c_i, k)$.*

Proof. We can use $(\frac{6}{\varepsilon})^2$ unit squares to cover the entire cell, so $\text{Opt}(c_i, k) \geq \frac{k\varepsilon^2 W_\varepsilon}{72}$. By the same argument as in Theorem 1, the difference between $\text{Opt}(c_i, k)$ and the answer from Algorithm 3 are at most $4k$ times the maximum possible weight of points between two adjacent parallel partition lines. Therefore, the algorithm returns a $(1 - \frac{\varepsilon}{3})$ -approximate answer of $\text{Opt}(c_i, k)$. \square

Theorem 2. *Algorithm 3 returns a $(1 - \varepsilon)$ -approximate answer for $\text{MaxCov}_R(\mathcal{P}, m)$ in time*

$$O\left(\frac{n}{\varepsilon} \log \frac{1}{\varepsilon} + m \left(\frac{1}{\varepsilon}\right)^\Delta\right),$$

where $\Delta = O(\min(\sqrt{m}, \frac{1}{\varepsilon}))$.

The proof is similar to Theorem 1, which is given in Appendix B.

References

1. P. K. Agarwal, T. Hagerup, R. Ray, M. Sharir, M. Smid, and E. Welzl. Translating a planar object to maximize point containment. In *ESA*, pages 42–53. Springer, 2002.
2. P. K. Agarwal and C. M. Procopiuc. Exact and approximation algorithms for clustering. *Algorithmica*, 33(2):201–226, 2002.
3. B. Aronov and S. Har-Peled. On approximating the depth and related problems. *SICOMP*, 38(3):899–921, 2008.
4. G. Barequet, M. Dickerson, and P. Pau. Translating a convex polygon to contain a maximum number of points. *Computational Geometry*, 8(4):167–179, 1997.
5. M. Ben-Or. Lower bounds for algebraic computation trees. In *STOC*, pages 80–86. ACM, 1983.
6. D. Bremner, T. M. Chan, E. D. Demaine, J. Erickson, F. Hurtado, J. Iacono, S. Langerman, and P. Taslakian. Necklaces, convolutions, and $x+y$. In *ESA*, pages 160–171. Springer, 2006.
7. H. Brönnimann and M. Goodrich. Almost optimal set covers in finite vc-dimension. *DCG*, 14(1):463–479, 1995.
8. S. Cabello, J. M. Díaz-Báñez, C. Seara, J. A. Sellares, J. Urrutia, and I. Ventura. Covering point sets with two disjoint disks or squares. *Computational Geometry*, 40(3):195–206, 2008.
9. T. M. Chan, E. Grant, J. Könemann, and M. Sharpe. Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling. In *SODA, SODA '12*, pages 1576–1585. SIAM, 2012.

10. B. M. Chazelle and D.-T. Lee. On a circle placement problem. *Computing*, 36(1-2):1–16, 1986.
11. K. L. Clarkson and K. Varadarajan. Improved approximation algorithms for geometric set cover. *DCG*, 37(1):43–58, 2007.
12. T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. The MIT Press, 3rd edition, 2009.
13. M. de Berg, S. Cabello, and S. Har-Peled. Covering many or few points with unit disks. *Theory of Computing Systems*, 45(3):446–469, 2009.
14. M. Dickerson and D. Scharstein. Optimal placement of convex polygons to maximize point containment. *Computational Geometry*, 11(1):1–16, 1998.
15. M. Dietzfelbinger, T. Hagerup, J. Katajainen, and M. Penttonen. A reliable randomized algorithm for the closest-pair problem. *Journal of Algorithms*, 25(1):19–51, 1997.
16. Z. Drezner. Note on a modified one-center model. *Management Science*, 27(7):848–851, 1981.
17. G. Even, D. Rawitz, and S. M. Shahar. Hitting sets when the vc-dimension is small. *IPL*, 95(2):358–362, 2005.
18. U. Feige. A threshold of $\ln n$ for approximating set cover. *JACM*, 45(4):634–652, 1998.
19. E. Fogel, D. Halperin, L. Kettner, M. Teillaud, R. Wein, and N. Wolpert. Arrangements. In J.-D. Boissonat and M. Teillaud, editors, *Effective Computational Geometry for Curves and Surfaces*, Mathematics and Visualization, chapter 1, pages 1–66. Springer, 2007.
20. R. L. Graham. An efficient algorithm for determining the convex hull of a finite planar set. *Information processing letters*, 1(4):132–133, 1972.
21. D. S. Hochbaum and W. Maass. Approximation schemes for covering and packing problems in image processing and vlsi. *JACM*, 32(1):130–136, 1985.
22. D. S. Hochbaum and A. Pathria. Analysis of the greedy approach in problems of maximum k-coverage. *Naval Research Logistics*, 45(6):615–627, 1998.
23. H. Imai and T. Asano. Finding the connected components and a maximum clique of an intersection graph of rectangles in the plane. *Journal of algorithms*, 4(4):310–323, 1983.
24. N. Megiddo and K. J. Supowit. On the complexity of some common geometric location problems. *SICOMP*, 13(1):182–196, 1984.
25. N. H. Mustafa and S. Ray. Ptas for geometric hitting set problems via local search. In *SCG*, pages 17–22. ACM, 2009.
26. S. C. Nandy and B. B. Bhattacharya. A unified algorithm for finding maximum and minimum object enclosing rectangles and cuboids. *Computers & Mathematics with Applications*, 29(8):45–61, 1995.
27. G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions. *Mathematical Programming*, 14(1):265–294, 1978.
28. Y. Tao, X. Hu, D.-W. Choi, and C.-W. Chung. Approximate maxrs in spatial databases. *Proceedings of the VLDB Endowment*, 6(13):1546–1557, 2013.
29. K. Varadarajan. Weighted geometric set cover via quasi-uniform sampling. In *STOC*, pages 641–648. ACM, 2010.
30. R. Williams. Faster all-pairs shortest paths via circuit complexity. In *STOC*, pages 664–673. ACM, 2014.




Fig. 2. Proof of Theorem 1. Difference between the optimal solution (the top-left one) and our solution (the lower-right one)

A Missing Proofs from Section 3

Theorem 1 *Algorithm 1 returns a $(1-\varepsilon)$ -approximate answer for $\text{MaxCov}_R(\mathcal{P}, 1)$ in $O(n \log(\frac{1}{\varepsilon}))$ time.*

Proof. We only need to prove that $\text{MAXCOVCELL}(c)$ returns a $(1-\varepsilon)$ -approximation for cell c . The case $n_c < (\frac{1}{\varepsilon})^2$ is trivial since we apply the exact algorithm. So we only need to prove the case of $n_c \geq (\frac{1}{\varepsilon})^2$.

Suppose the optimal unit square is r . Denote by Opt the weight of the optimal solution. The size of c is 2×2 , so we can use 4 unit squares to cover the entire cell. Therefore, $\text{Opt} \geq (\frac{W_c}{4})$. Suppose r is located at a point p , which is in the strict interior of a small cell B separated by \mathcal{L} .⁷ Suppose the index of B is (i, j) . We compare the weight of r with $I(i, j)$ (which is the approximate weight of the unit square located at the top-left corner of B). See Figure 2. By the rule of our partition, the weight difference is at most 4 times the maximum possible weight of points between two adjacent lines in \mathcal{L} . So $I(i, j) \geq \text{Opt} - 4 \cdot \frac{\varepsilon \cdot W_c}{16} \geq (1-\varepsilon)\text{Opt}$. This proves the approximation guarantee of the algorithm.

Now, we analyze the running time. The running time consists of two parts: cells with number of points more than $(\frac{1}{\varepsilon})^2$ and cells with number of points less than $(\frac{1}{\varepsilon})^2$. Let $n_1 \geq n_2 \geq \dots \geq n_j \geq (\frac{1}{\varepsilon})^2 > n_{j+1} \geq n_{j+2}, \dots, \geq n_{j+k}$ be the

⁷ If p lies on the boundary of B , the same argument still works.

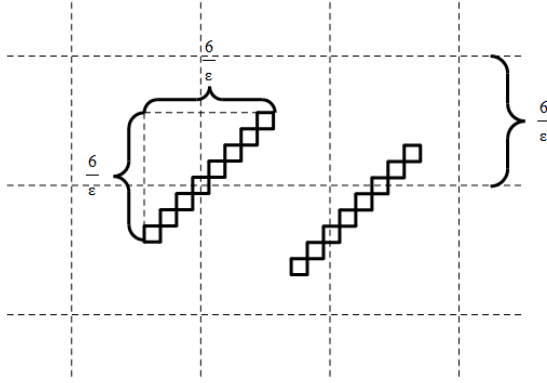


Fig. 3. Proof of Lemma 3: the shifting technique.

sorted sequence of the number of points in all cells. Then, we have that

$$\begin{aligned}
 \text{Running time} &\leq \sum_{i=1}^j O\left(n_i \log\left(\frac{1}{\varepsilon}\right) + \left(\frac{1}{\varepsilon}\right)^2\right) + \sum_{i=1}^k O(n_{i+j} \log(n_{i+j})) \\
 &= O\left(\log\left(\frac{1}{\varepsilon}\right) \sum_{i=1}^j n_i + j \left(\frac{1}{\varepsilon}\right)^2 + \sum_{i=1}^k n_{i+j} \log(n_{i+j})\right) \\
 &\leq O\left(\log\left(\frac{1}{\varepsilon}\right) \sum_{i=1}^j (n_i) + n + \sum_{i=1}^k n_{i+j} \log\left(\frac{1}{\varepsilon}\right)\right) \\
 &= O\left(\log\left(\frac{1}{\varepsilon}\right) \sum_{i=1}^{j+k} (n_i) + n\right) = O\left(n \log\left(\frac{1}{\varepsilon}\right)\right).
 \end{aligned}$$

This completes the proof. \square

B Missing Proofs from Section 4

Lemma 3 *There exist $G^* \in \mathbb{G}$ and a $(1 - \frac{2\varepsilon}{3})$ approximate solution R such that none of the unit squares in R intersects G^* .*

Proof. For any point p , we can always use four unit squares to cover the 2×2 square centered at p . Therefore, there exists an optimal solution OPT such that each covered point is cover by at most 4 unit squares in OPT . For each grid $G_{\frac{\varepsilon}{2}}(i, i) \in \mathbb{G}$, we build a modified answer R_i from OPT in the following way. For each square r that intersects with $G_{\frac{\varepsilon}{2}}(i, i)$, there are two different situations. If r only intersects with one vertical line or one horizontal line. We move the square to one side of the line with bigger weight. In this case we will lose at most half of the weight of r . Notice that this kind of squares can only intersect with two grids in \mathbb{G} . Similarly, If r intersects with one vertical line and one




Fig. 4. $F(c_i, k)$ may not be concave: $F(c_i, 1) = 3$, $F(c_i, 2) = 4$, $F(c_i, 3) = 6$.

horizontal line at the same time, we move it to one of the four quadrants derived by these two lines. In this case we will lose at most $3/4$ of the weight of r . This kind of squares can only intersect with one grid in \mathbb{G} . (see Figure 3) Now we calculate the sum of the weights we lose from $R_0, R_2, \dots, R_{\frac{6}{\varepsilon}-1}$, which is at most $\max\{1/2 \times 2, 3/4 \times 1\} = 1$ times the sum of weights of squares in OPT . By the definition of OPT , it is at most $4w(OPT)$. So the sum of the weights of $R_0, R_2, \dots, R_{\frac{6}{\varepsilon}-1}$ is at least $(\frac{6}{\varepsilon} - 4)w(OPT)$. Therefore there exists some i such that R_i (which does not intersect $G_{\frac{6}{\varepsilon}}(i, i)$) is a $(1 - \frac{2\varepsilon}{3})$ approximate answer. \square

Theorem 2 *Algorithm 3 returns a $(1 - \varepsilon)$ -approximate answer for $\text{MaxCov}_R(\mathcal{P}, m)$ in time*

$$O\left(\frac{n}{\varepsilon} \log \frac{1}{\varepsilon} + m \left(\frac{1}{\varepsilon}\right)^\Delta\right),$$

where $\Delta = O(\min(\sqrt{m}, \frac{1}{\varepsilon}))$.

Proof. We know that $\text{Opt}_F(m)$ of m selected cells is a $(1 - \frac{\varepsilon}{3})$ -approximation to Opt' , so the greedy algorithm returns a $(1 - \frac{\varepsilon}{3})(1 - \frac{\varepsilon^2}{9})$ -approximation to Opt' . Combining with Lemma 3 and Lemma 4, Algorithm 3 returns a $(1 - \frac{2\varepsilon}{3})(1 - \frac{\varepsilon}{3})(1 - \frac{\varepsilon^2}{9})(1 - \frac{\varepsilon^2}{9})$ approximate solution of the original problem. Since $(1 - \frac{2\varepsilon}{3})(1 - \frac{\varepsilon}{3})(1 - \frac{\varepsilon^2}{9})(1 - \frac{\varepsilon^2}{9}) > (1 - \varepsilon)$, Algorithm 3 does return a $(1 - \varepsilon)$ -approximate solution.

The algorithm is summarized in Algorithm 3. We only need to calculate the overall running time. Solving the values $F(c_i, 1)$ and finding out the top m results require $O(n \log \frac{1}{\varepsilon})$ time. We compute the values $F(c_i, k)$ of m cells. For each cell c_i , we partition it only once and calculate $F(c_i, 1), \dots, F(c_i, \Delta)$ using the same partition. Computing the values $F(c_i, k)$ of all m cells requires $O(n \log(\frac{1}{\varepsilon}) + m(\frac{1}{\varepsilon})^\Delta)$ time. We do the same for $\frac{1}{\varepsilon}$ different grids. Therefore, the over all running time is as we state in the theorem. \square

C Enumeration In MaxCovCellM

We can adapt the algorithm in [13] to enumerate these possible ways of placing k unit squares at the grid point in $O((\frac{1}{\varepsilon})^\Delta)$ time where $\Delta = O(\min(\sqrt{m}, \frac{1}{\varepsilon}))$. We

Algorithm 3 MaxCov_R(\mathcal{P}, m)

```

 $w_{\max} \leftarrow 0$ 
for each  $G \in \{G_{\frac{6}{\varepsilon}}(0, 0), \dots, G_{\frac{6}{\varepsilon}}(\frac{6}{\varepsilon} - 1, \frac{6}{\varepsilon} - 1)\}$  do
  Use perfect hashing to find all the non-empty cells of  $G$ .
  for each non-empty cell  $\mathbf{c}$  of  $G$  do
     $r_{\mathbf{c}} \leftarrow$  Algorithm 1 for  $\mathbf{c}$  with approximate ratio  $(1 - \frac{\varepsilon^2}{9})$ 
  end for;
  Find the  $m$  cells with the largest  $r_{\mathbf{c}}$ . Suppose they are  $\mathbf{c}_1, \dots, \mathbf{c}_m$ .
  for  $i \leftarrow 1$  to  $m$  do
    for  $k \leftarrow 1$  to  $b$  do  $F(\mathbf{c}_i, k) \leftarrow$  MAXCOVCELLM( $\mathbf{c}, k$ )
    end for
  end for;
  if  $m \leq 324(\frac{1}{\varepsilon})^4$ , then  $r \leftarrow$  DP( $\{F(\mathbf{c}_i, k)\}$ ) else  $r \leftarrow$  GREEDY( $\{F(\mathbf{c}_i, k)\}$ )
  if  $w(r) > w_{\max}$ , then  $w_{\max} \leftarrow w(r)$  and  $r_{\max} \leftarrow r$ 
end for;
return  $r_{\max}$ ;

```

briefly sketch the algorithm. We denote the optimal solution as $\text{Opt}_{\mathbf{c}}$. From [2] we know that for any optimal solution, there exists a line of integer grid (either horizontal or vertical) that intersects with $O(\sqrt{m})$ squares in $\text{Opt}_{\mathbf{c}}$, denoted as the *parting line*. So we can use dynamic programming. At each stage, we enumerate the parting line, and the $O(\sqrt{m})$ squares intersecting the parting line. We also enumerate the number of squares in each side of the parting line in the optimal solution. The total number of choices is $O((\frac{1}{\varepsilon})^{\Delta})$. Then, we can solve recursively for each side. In the recursion, we should consider that a subproblem which is composed of a smaller rectangle, and an enumeration of $O(\sqrt{m})$ squares of the optimal solution intersecting the boundary of the rectangle and at most m squares fully contained in the rectangle. Overall, the dynamic programming can be carried out in $O((\frac{1}{\varepsilon})^{\Delta})$ time.

D Extension to Other Shapes

Our algorithm can be easily extended to other shapes, such as unit disks. For ease of description, we only consider unit disks, i.e., the MaxCov_D(\mathcal{P}, m) problem. The algorithm framework is almost the same as before, except several implementation details required changing from rectangles to disks. The major difference is the way for building an ε -approximation in each cell (the partition scheme in Section 4.4 works only for rectangles).

Now, we sketch our algorithm, focusing only on the differences. Again, we first adopt the shifting technique. For unit disk, we consider grids with a different side length: $G_{56/\varepsilon}(a, b)$. Again for simplicity, we assume that $\frac{1}{\varepsilon}$ is an integer and no point in \mathcal{P} has an integer coordinate. We shift the grid to $\frac{28}{\varepsilon}$ different positions: $(0, 0), (2, 2), \dots, (\frac{56}{\varepsilon} - 2, \frac{56}{\varepsilon} - 2)$. Let $\mathbb{G} = \{G_{56/\varepsilon}(0, 0), \dots, G_{56/\varepsilon}(56/\varepsilon - 2, 56/\varepsilon - 2)\}$. We can prove the following lemma, which is similar to Lemma 3. The only difference in the proof is that for unit disks, there exists an optimal answer Opt

such that each covered point is covered by at most 7 disks in Opt instead of 4 for unit squares.

Lemma 7. *There exist $G^* \in \mathbb{G}$ and a $(1 - \frac{\varepsilon}{2})$ -approximate answer R such that all the disks in R do not intersect any line in G^* .*

Now consider a fixed grid $G \in \mathbb{G}$. For each cell c , we compute a $(1 - \frac{\varepsilon^2}{8})$ -approximation to $\text{Opt}(c, 1)$, in $O(n\varepsilon^{-6})$ time, by applying the following result in [13].

Lemma 8. [13] *Give a set \mathcal{P} of n weighted points and $\varepsilon(0 < \varepsilon < 1)$, we can find a $(1 - \varepsilon)$ -approximation for $\text{MaxCov}_D(\mathcal{P}, 1)$ in $O(n\varepsilon^{-3})$ time.*

Then, we use the selection algorithm to find out the m cells with the largest returned values. The rest of our algorithm will only consider those m cells. For each such cell c , we want to compute $(1 - \frac{\varepsilon}{2})$ -approximations $F(c, k)$ of $\text{Opt}(c, k)$. For this purpose, we can use the following lemma from [13].

Lemma 9. [13] *Give a set \mathcal{P} of n weighted points, an integer $m \geq 1$ and $r \geq 1$. Let U be the collection of sets that are the union of m unit disks. In range space $(\mathcal{P}, \{\mathcal{P} \cap s \mid s \in U\})$, Let \mathcal{A} be a $(1/2r)$ -approximation for the range space, where $r = w(\mathcal{P})/(\varepsilon \text{MaxCov}_D(\mathcal{P}, m))$. If $\text{Opt}_{\mathcal{A}}$ is a optimal solution for $\text{MaxCov}_D(\mathcal{A}, m)$, then $\text{Opt}_{\mathcal{A}}$ is a $(1 - \varepsilon)$ -approximation for $\text{MaxCov}_D(\mathcal{P}, m)$.*

By the above lemma, we only need to build a $(1/4r)$ -approximation \mathcal{A} for the cell c , where $r = W_c/(\varepsilon \text{Opt}(c, k))$ and apply an exact algorithm to \mathcal{A} . Since we can use $O(1/\varepsilon^2)$ unit disks to cover the whole cell, we can see that $r < O(1/k\varepsilon^3)$. We can build the $1/4r$ -approximation \mathcal{A} of size $O(\frac{k^2}{\varepsilon^6} \log \frac{k}{\varepsilon})$ in $O(n_c \cdot (\frac{k}{\varepsilon})^{O(1)})$ time (see e.g., [19]). Using the exact algorithm in [13] for $\text{MaxCov}_D(\mathcal{P}, m)$, we can compute the exact solution for \mathcal{A} in $O((\frac{k}{\varepsilon})^{O(\sqrt{k})})$ time.

Suppose we have computed all $F(c, k)$ values for the first m cells. If $m \leq O(1/\varepsilon^4)$, we apply the dynamic programming, as in Section 4.2. Otherwise, we apply the algorithm GREEDY in Section 4.3. The correctness proof is exactly the same, except that the approximation ratio in Lemma 5 changes slightly to $(1 - \frac{\varepsilon^2}{8})$.

Now we can summarize the overall running time. Computing the value $F(c, 1)$ for each cell requires $O(n\varepsilon^{-6})$ time. We build ε -approximations in m different cells and k is at most $\min\{O(1/\varepsilon^2), m\}$. Therefore the overall running time is $O(n(\frac{1}{\varepsilon})^{O(1)} + m(\frac{1}{\varepsilon})^\Delta)$, where $\Delta = O(\min(\sqrt{m}, \frac{1}{\varepsilon}))$.