PTAS for Weighted Set Cover on Unit Squares

Thomas Erlebach[1] and Erik Jan van Leeuwen[2]

[1] Department of Computer Science, University of Leicester, University Road, Leicester, LE1 7RH, UK. t.erlebach@mcs.le.ac.uk

[2] Department of Informatics, University of Bergen, P.O. Box 7803, N-5020 Bergen, Norway. E.J.van.Leeuwen@ii.uib.no

Abstract. We study the planar version of Minimum-Weight Set Cover, where one has to cover a given set of points with a minimum-weight subset of a given set of planar objects. For the unit-weight case, one PTAS (on disks) is known. For arbitrary weights however, the problem appears much harder, and in particular no PTASs are known. We present the first PTAS for Weighted Geometric Set Cover on planar objects, namely on axis-parallel unit squares. By extending the algorithm, we also obtain a PTAS for Minimum-Weight Dominating Set on intersection graphs of unit squares and Geometric Budgeted Maximum Coverage on unit squares. The running time of the developed algorithms is optimal under the exponential time hypothesis. We also show inapproximability results for Geometric Set Cover on various object shapes that are more general than unit squares.

1 Introduction

One of the most fundamental and best-known optimization problems is Minimum Set Cover. Given a universe $\mathbb{U}$, a set of elements $\mathcal{P} \subseteq \mathbb{U}$, and a set $\mathcal{S}$ of subsets of $\mathbb{U}$, one should find a minimum set $S \subseteq \mathcal{S}$ such that each element of $\mathcal{P}$ is contained in (covered by) a set in S. If $\mathbb{U} = \mathbb{R}^d$ for some $d > 0$, we talk about *Geometric Set Cover*. In particular, we are interested in the case where $d = 2$ and the sets in $\mathcal{S}$ are induced by simple geometric shapes, such as disks or squares. Geometric Set Cover can be better approximated than general Minimum Set Cover [2, 5, 23, 29], but for many object shapes the approximability has not been settled yet, particularly in the weighted case. In this paper, we consider the approximability of Geometric Set Cover and several of its variants, with emphasis on weighted cases.

Motivation Minimum Set Cover is known to be approximable within $1 + \ln |\mathcal{P}|$, even in the weighted case [20, 25, 4]. This algorithm is also optimal. That is, Minimum Set Cover has no polynomial-time algorithm attaining an approximation ratio of $(1 - \epsilon) \ln |\mathcal{P}|$ for any $\epsilon > 0$, unless NP $\subset$ DTIME($n^{\mathrm{O}(\log \log n)}$) [11]. Because of its applicability in the design of (wireless) networks, Geometric Set Cover has recently received a lot of attention. Geometric Set Cover is NP-hard on unit squares and on unit disks [12, 21], even if the point set $\mathcal{P}$ corresponds to the centers of the squares or disks. This has led to researchers studying approximation algorithms for variations of the problem.

The biggest focus of approximation algorithms for Geometric Set Cover has been on (unit) disks. Several constant-factor approximation algorithms were proposed on unweighted unit disks [2, 6, 30, 3]. This line of research recently culminated in the discovery of a PTAS for Geometric Set Cover on general disks [29] using a transformation into the geometric version of Minimum Hitting Set on three-dimensional half-spaces.

The above algorithms are not known to be applicable to the weighted case and only recently have algorithms approximating the problem started to appear. After a few iterations [1, 19, 7], a $(4 + \epsilon)$-approximation algorithm on unit disks, independently proposed by Zou et al. [34] and Erlebach and Mihalák [9], currently is the best known result. Varadarajan [32] gives a $2^{O(\log^* n)}$-approximation on general disks. On unit squares, a 2-approximation algorithm exists [28]. It seems however that past approaches are insufficient to reach a PTAS, except when the disk centers have a constant minimum distance from each other [13, 24].

We also consider the geometric version of the Budgeted Maximum Coverage problem. Here each element u of $\mathcal{P}$ has a profit $p(u)$, each set $\mathcal{S}_i$ of $\mathcal{S}$ a cost $c(\mathcal{S}_i)$, and we aim to maximize the total profit of the points covered by some $S \subseteq \mathcal{S}$, while the total cost of S is no more than a given budget B. Budgeted Maximum Coverage has a $1-\frac{1}{e}$-approximation algorithm in both the unit cost [33, 18, 16] and the general case [22]. Khuller, Moss, and Naor [22] proved that no polynomial-time algorithm can obtain an approximation ratio better than $(1-\frac{1}{e})$, unless NP $\subset$ DTIME($n^{\mathrm{O}(\log \log n)}$). As far as we know, Geometric Budgeted Maximum Coverage has not been studied yet. The problem can be shown to be NP-hard on unit squares by reduction from Geometric Set Cover.

Observe that Geometric Set Cover differs significantly from the Geometric Covering problem, where the position of the objects may be chosen freely. This problem has a well-known PTAS both on unit disks [17] and on unit squares [14].

Our Results In this paper, we present a PTAS for Geometric Set Cover on any set of axis-parallel unit squares. Using a novel dynamic programming idea, refining the classic sweep-line technique, we are able to solve this problem optimally in $n^{O(k)}$ time when the given sets of points lie within a horizontal strip of height k. Combining this with the well-known shifting technique then yields the PTAS. We also observe that it follows from Marx [26, 27] that the scheme has essentially optimal running time (up to constants), unless the exponential time hypothesis is false.

The presented scheme extends to the weighted case of Geometric Set Cover and Minimum Dominating Set on intersection graphs of unit squares and in fact to the more general Geometric Budgeted Maximum Coverage problem. We note that the optimality result for our PTAS continues to hold.

Beside these positive algorithmic results, we also give several negative results. In particular, we show that Geometric Set Cover is APX-hard on arbitrary four-sided convex polygons. We also obtain APX-hardness results on axis-parallel rectangles and ellipses. Finally, we show that on convex polygons, translated copies of a single polygon, rotated copies of a single convex polygon, and α-fat objects, Geometric Set Cover is as hard as Minimum Set Cover.

2 A PTAS on Unit Squares

We consider Geometric Set Cover on unit squares and show that it has a PTAS by applying the shifting technique. So let $\mathcal{P}$ be a set of points and $\mathcal{S}$ a set of axis-aligned unit squares. For sake of notation, when referring to the (x, y)-coordinates of a square, we mean the coordinates of the bottom left corner of that square. For a square s, the x-coordinate of s is denoted by $x(s)$, while the y-coordinate is denoted by $y(s)$. We can assume that no horizontal (vertical) boundary of a square is on the same line as the horizontal (vertical) boundary of another square, that no point lies on the boundary of a square, and that none of the square or point coordinates are integers.

Consider the horizontal lines $y = h$ ($h \in \mathbb{Z}$). They partition the plane into horizontal slabs of height 1. Any point is contained in a slab and every square intersects precisely one line. Let $k \geq 1$ be an integer determined later. Using the shifting technique, it suffices to prove that we can optimally solve Geometric Set Cover on unit squares if we restrict to k consecutive slabs and the $k + 1$ lines defining them.

Theorem 1. *For any instance of Geometric Set Cover on a set of unit squares $\mathcal{S}$ where all points of $\mathcal{P}$ are inside $k \geq 1$ consecutive height 1 horizontal slabs, one can find an optimal solution in* $\mathrm{O}((3|\mathcal{S}|)^{4k+4}\,|\mathcal{P}|)$ *time.*

The idea of the proof of this theorem will be to apply a sweep-line algorithm. To this end, consider the subset of squares of an optimum solution intersecting a horizontal line $y = h$ for some $h \in \mathbb{Z}$. Any such square must appear on the lower or upper envelope of this subset, or all points it covers would be covered by other squares. Following this observation, for each position of the sweep-line and for each of the $k+1$ integer horizontal lines, we should consider at most two squares intersecting the sweep-line: one that will appear on the upper envelope and one that will appear on the lower envelope of the final solution.

However, a square might appear on the lower envelope for some position of the sweep-line and on the upper envelope for a later position. This makes it difficult to avoid counting certain squares twice. To circumvent this, we split the sweep-line into k parts, one part per slab. We move these parts at different speeds, but always in such a way that if a square appears both on the lower and the upper envelope, then the split sweep-line is positioned such that it intersects the square both at the point where the square appears on the lower and on the upper envelope. We formalize this intuition below. We remark that the basic idea of having a split sweep-line was also used by Erlebach and Mihalák [9], but the details of how the split sweep-line is then handled by a dynamic programming approach are very different in our case.

Just as in any sweep-line algorithm, we maintain a data structure (the *front*) containing the squares that are 'active' at a given position of the sweep-lines and allow only a limited number of operations on it.

Let $\mathcal{S}^l$ and $\mathcal{S}^r$ be two dummy sets of $k+1$ squares each, such that the squares in $\mathcal{S}^l$ ($\mathcal{S}^r$) are to the left (right) of all squares in $\mathcal{S}$ and each integer horizontal line intersects precisely one square of $\mathcal{S}^l$ and one square of $\mathcal{S}^r$. Let $\overline{\mathcal{S}} = \mathcal{S} \cup \mathcal{S}^l \cup \mathcal{S}^r$. Given some set $S \subseteq \overline{\mathcal{S}}$, let S_i denote the set of squares in S intersecting line i. Let $R_i \subseteq S_i$ be the set containing precisely the rightmost square of S_i (denote it by s_i) and those squares s that overlap part of the left boundary of s_i and whose right boundary is not fully covered by squares of S_i.

We now define a front. For a better understanding of the definition, imagine that the squares are being inserted in order of increasing x-coordinate and that we want to keep track of the upper and lower envelope of each line i.

Definition 1. *Let S be the union of $\mathcal{S}^l$ and some subset of $\overline{\mathcal{S}}$. Then a* front $F = \{u_1, \ldots, u_{k+1}, l_1, \ldots, l_{k+1}, b_1, \ldots, b_{k+1}, x_1, \ldots, x_k\}$ *for S has the following:*

- *$u_i, l_i \in R_i$ with $u_i = s_i$ or $l_i = s_i$, $y(s) \leq y(u_i)$ for any $s \in S_i$ to the right of u_i (i.e. with $x(s) > x(u_i)$) and $y(s) \geq y(l_i)$ for any $s \in S_i$ to the right of l_i (i.e. with $x(s) > x(l_i)$),*
- *b_i equals the lowest square of S_i to the right of l_i if $x(u_i) > x(l_i)$, the highest square of S_i to the right of u_i if $x(l_i) > x(u_i)$, and s_i if $x(u_i) = x(l_i)$,*
- *x_i equals the larger x-coordinate from which l_{i+1} appears on the lower envelope of S_{i+1} and from which u_i appears on the upper envelope of S_i.*

Fronts are the representative of the current state of the sweep-line algorithm. The squares u_i and l_i track the current square on respectively the upper and the lower envelope of line i. The value of x_i is the x-coordinate of the part of the sweep-line between lines i and $i+1$. The square b_i is used in checking if a certain square may be inserted into the front or not. An example is depicted in Figure 1.

We make two observations about fronts. Firstly, $y(u_i) \geq y(l_i)$ and as $u_i, l_i \in R_i$, $|x(u_i) - x(l_i)| < 1$ for any $i = 1, \ldots, k+1$. Secondly, $y(l_i) \leq y(b_i) \leq y(u_i)$.

For a given front, we distinguish four types of insertions that are possible: an upper-insertion for squares that will appear only on the upper envelope for some line, a lower-insertion for squares appearing only on the lower envelope, and a middle-insertion and a skip-insertion for squares appearing on both envelopes. We define these four insertions, describe when they may be applied, and prove that any geometric set cover can be obtained using these insertions.

From now on, S will denote the union of $\mathcal{S}^l$ and some subset of $\overline{\mathcal{S}}$.

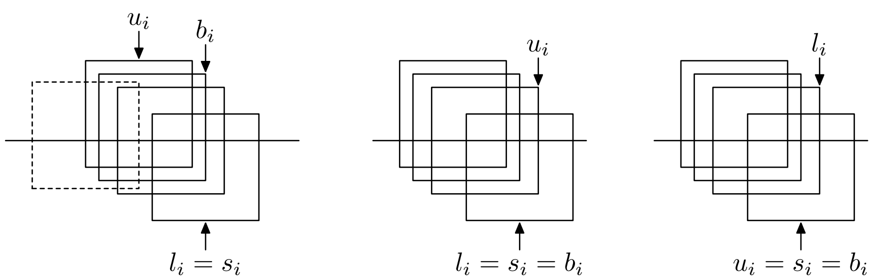

Fig. 1. The left figure shows a set S_i. The solid squares are in R_i, the dashed square is not. By Definition 1, the labeling of the left figure is correct. The middle figure shows the same set R_i, with a different and still correct labeling. The labeling in the right figure is incorrect.

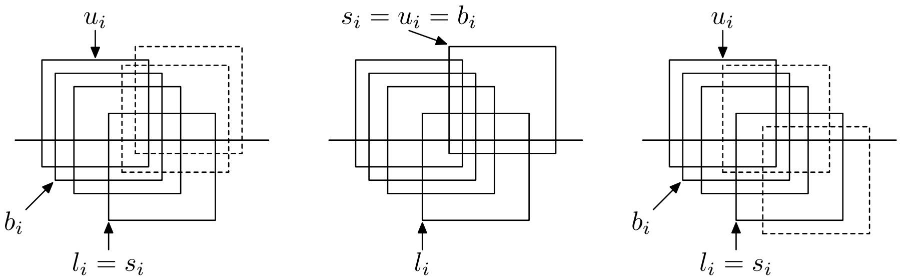

Fig. 2. The left figure shows two (dashed) squares that are upper-insertable. The middle figure shows the resulting front after upper-inserting the rightmost square. The right figure shows two (dashed) non-upper-insertable squares.

Definition 2. *Let F be a front for some S and let $s \notin S$ be a square intersecting line $i \in \{1, \ldots, k\}$. We say that s is* upper-insertable *into F if all of the following hold: 1) $y(s) > y(l_i)$ and if $x(l_i) > x(u_i)$, then $y(s) > y(b_i)$, 2) $x(s) \in (x(l_i), x(l_i) + 1]$ and $x(s) \in (x(u_i), x(u_i) + 1]$, 3) $x'_i > x_i$, 4) any point of $\mathcal{P}$ in $[x_i, x'_i] \times [i, i + 1]$ is covered by u_i or l_{i+1}, where x'_i is the x-coordinate from which s is on the upper envelope of $(S \cup \{s\})_i$.*

Condition 1 ensures that s lies above l_i and all squares between u_i and l_i (represented by b_i), Condition 2 ensures that s appears on the upper envelope of $(S \cup \{s\})_i$, Condition 3 ensures that this appearance happens after u_i appears on the upper envelope, and Condition 4 ensures that we cover all points between two consecutive sweep-line positions. An example of upper-insertable squares and squares that are not upper-insertable is given in Figure 2.

Lemma 1. *Let F be a front for some S and let $s \notin S$ be a square intersecting line $i \in \{1, \ldots, k\}$ that is upper-insertable into F. Then between the appearance of u_i and the appearance of s on the upper envelope of $(S \cup \{s\})_i$ no other squares appear on the upper envelope of $(S \cup \{s\})_i$.*

Proof. If $u_i = s_i$, this follows from $x(s) > x(u_i) = x(s_i)$ and $x'_i > x_i$. So assume that $u_i \neq s_i$. Then $l_i = s_i$ and $x(l_i) > x(u_i)$. Recall the definition of a front and observe that b_i is the highest square of S_i to the right of u_i. As $x(l_i) - x(u_i) < 1$ and $y(b_i) < y(u_i)$, it suffices for s to lie above b_i (i.e. $y(s) > y(b_i)$) and for s to cover the x-range $[x(u_i) + 1, x(l_i) + 1]$ (i.e. $x(l_i) < x(s) < x(u_i) + 1$). This holds from the definition of upper-insertable. $\square$

Lemma 2. *Let F be a front for some S and let $s \notin S$ be a square intersecting line $i \in \{1, \dots, k\}$ that is upper-insertable into F. Then $S \cup \{s\}$ has a front F' equal to F, except u_i is replaced by s, x_i is set to x'_i, where x'_i is equal to $x(s)$ if $y(s) > y(u_i)$ and to $x(u_i) + 1$ otherwise, and if $x(u_i) \leq x(l_i)$ or $y(s) \leq y(b_i)$, b_i is set to s.*

Proof. Since $x(s) > \max\{x(u_i), x(l_i)\} = x(s_i)$ by Condition 2 of Definition 2, we can replace u_i by s. Note that l_i can remain the same by Condition 1 and 2. By Lemma 1, x'_i is indeed the x-coordinate from which s appears on the upper envelope of $(S \cup \{s\})_i$. From Condition 3, x_i should be set to x'_i. If $x(u_i) \leq x(l_i)$, then as $x(s) > x(l_i)$, b_i should be set to s. If $x(u_i) > x(l_i)$, then b_i must be changed if s lies below b_i, i.e. if $y(s) \leq y(b_i)$. Then F' is indeed a front for $S \cup \{s\}$. ☐

Constructing the front F' from F as prescribed in the lemma statement is called the *upper-insertion* of s into F.

We can define the notions of lower-/middle-/skip-insertable and lower-/middle-/skip-insertions in similar ways. The definitions of upper- and lower-insertable/-insertion are similar, except that we check if the square we want to insert will appear on the lower envelope directly after l_i appears on the lower envelope. The definition of middle-insertable/-insertion combines the definitions of upper- and lower-insertable/-insertion. Skip-insertions are used when the square we want to insert does not intersect u_i or l_i, i.e. when $x(s) > 1 + \max\{x(u_i), x(l_i)\}$. Full definitions are in [31].

In general, we call an upper-/lower-/middle-/skip-insertion an *insertion* and we say s is *insertable* if it is upper-/lower-/middle-/skip-insertable. A *valid insertion* is the upper- (respectively lower-/middle-/skip-) insertion of a square that is upper- (respectively lower-/middle-/skip-) insertable.

Denote by F^l and F^r the fronts for $\mathcal{S}^l$ and $\overline{\mathcal{S}}$ respectively.

Lemma 3. *Assume $\mathcal{P} = \emptyset$. Let S be some set such that $S = \mathcal{S}^l \cup S_i \cup \mathcal{S}^r$ for some $i \in \{1, \dots, k+1\}$ and any square in S_i appears on the lower or the upper envelope of S_i. Then there is a sequence of $|S_i| + k - 1$ valid insertions starting from F^l, leading to fronts $F^l = F_0, F_1, \dots, F_{|S_i|+k-1} = F^r$ such that for any square $s \in S_i$, there is a front F_j containing s.*

Proof (Sketch). We assume that if $i = 1$, then no squares of S_i appear only on the lower envelope of S_i. Similarly, if $i = k+1$, assume that no squares of S_i appear only on the upper envelope of S_i. Order the squares in $S_i \backslash \mathcal{S}^l$ by increasing x-coordinate, i.e. $s_1, \dots, s_{|S_i|-1}$. Note that the squares appearing on the upper envelope form an increasing subsequence of S_i.

Similarly, the squares appearing on the lower envelope form an increasing subsequence. We claim that one can obtain the requested sequence of valid insertions by inserting s_j into F_{j-1} for all $j = 1, \ldots, |S_i| - 1$ as follows. If s_j appears only on the upper envelope of S_i, then s_j is upper-insertable and will be upper-inserted. If s_j appears only on the lower envelope of S_i, then s_j is lower-insertable and will be lower-inserted. If s_j appears on the upper and lower envelope of S_i and a square of S_i covers part of its left boundary, then s_j is middle-insertable and will be middle-inserted. If s_j appears on the upper and lower envelope of S_i and no square of S_i covers part of its left boundary, then s_j is skip-insertable and will be skip-inserted. Now apply induction on the number of inserted squares. □

Lemma 4. *Assume $\mathcal{P} = \emptyset$. Let S be some subset of $\overline{\mathcal{S}}$ containing $\mathcal{S}^l \cup \mathcal{S}^r$, such that for the set S_i of squares in S intersecting line i for $i \in \{1, \ldots, k+1\}$, any square in S_i appears on the upper or lower envelope of S_i. Then there is a sequence of $|S| - k - 1$ valid insertions starting from $F_0 = F^l$, leading to $F_1, \ldots, F_{|S|-k-1} = F^r$ such that for any square $s \in S$, there is a front F_j containing s.*

Proof (Sketch). By the previous lemma, we can insert the squares intersecting each horizontal line in order of increasing x-coordinate. However, we should interleave the sequences of the different lines. For any $i = 1, \ldots, k$, consider the squares appearing on the upper envelope of S_i and the lower envelope of S_{i+1}. Order these squares according to the x-coordinate from which they appear on the upper envelope of S_i or on the lower envelope of S_{i+1} respectively. Combining these two orders, we can extend this to an order by which to insert the squares of S. We can then prove that the j-th square s_j according to this order is insertable into F_{j-1} and that after inserting s_j, all squares $s_{j'}$ with $j' > j$ are still insertable. □

The next lemmas follow from the coverage constraints on valid insertions.

Lemma 5. *Let S be any smallest subset of $\overline{\mathcal{S}}$ containing $\mathcal{S}^l \cup \mathcal{S}^r$ and covering all points in $\mathcal{P}$. Then there is a sequence of $|S| - k - 1$ valid insertions starting from F^l, leading to $F_1, \ldots, F_{|S|-k-1} = F^r$ such that for any square $s \in S$, there is a front F_j containing s.*

Lemma 6. *Let $l \geq 0$. Then any sequence of $l + k + 1$ valid insertions starting from F^l and resulting in F^r corresponds to a set $S \subseteq \mathcal{S}$ of cardinality l covering all points in $\mathcal{P}$.*

Proof (of Theorem 1). Construct a directed graph G with $V(G)$ equal to the set of all fronts and a directed edge from front F to front F' if F' can be obtained from F by a single valid insertion. From the definition of a front, $|V(G)| = \mathrm{O}(|\overline{\mathcal{S}}|^{4k+3})$. As each front allows for at most $4|\overline{\mathcal{S}}|$ valid insertions, $|E(G)| = \mathrm{O}(|\overline{\mathcal{S}}|^{4k+4})$. Because the validity of an insertion can be checked in $\mathrm{O}(|\mathcal{P}|)$ time, G can be constructed in $\mathrm{O}(|\overline{\mathcal{S}}|^{4k+4}|\mathcal{P}|)$ time.

From Lemma 5 and 6, a shortest path in G from F^l to F^r corresponds to a minimum subset of $\mathcal{S}$ covering all points in $\mathcal{P}$. Then a shortest path can be found in $\mathrm{O}(|E(G)|) = \mathrm{O}(|\overline{\mathcal{S}}|^{4k+4})$ time. Observe that $|\overline{\mathcal{S}}| = |\mathcal{S}| + |\mathcal{S}^l| + |\mathcal{S}^r| \leq 3|\mathcal{S}|$, because if no square intersects a certain line, we may ignore it. The running time of the algorithm is $\mathrm{O}((3|\mathcal{S}|)^{4k+4}|\mathcal{P}|)$. $\square$

Using Theorem 1 with the shifting technique, we get a PTAS for Geometric Set Cover on unit squares. The proof of this theorem can be found in [31].

Theorem 2. *There is a PTAS for Geometric Set Cover on unit squares.*

3 Geometric Budgeted Maximum Coverage

The above PTAS easily extends to the weighted case of Geometric Set Cover, by weighting the graph constructed in the proof of Theorem 1. We can however extend to the more general budgeted case as well.

Let $\mathcal{S}$ be a set of unit squares, $\mathcal{P}$ a set of points, c a cost function over $\mathcal{S}$, p a nonnegative profit function over $\mathcal{P}$, and B a budget. Let $p_{\max}$ denote the maximum profit of any single point. We define the function $cov(s)$ as the set of points in $\mathcal{P}$ covered by a square $s \in \mathcal{S}$. This notation extends to $cov(S)$ for a set $S \subseteq \mathcal{S}$. Abusing notation, we will use $p(S)$ to denote $p(cov(S))$.

Theorem 3. *In Geometric Budgeted Maximum Coverage on a set of unit squares $\mathcal{S}$ where all points are inside $k-1$ consecutive height 1 horizontal slabs and all profits are positive integers, one can find a cheapest set of profit at least r (if one exists) for all $0 \leq r \leq |\mathcal{P}| \cdot p_{\max}$ in time $\mathrm{O}((3|\mathcal{S}|)^{4k}\,(|\mathcal{P}| \cdot p_{\max}))$.*

Proof. We modify the algorithm described above. Assume the cost of squares in $\mathcal{S}^l \cup \mathcal{S}^r$ to be zero. Remove the coverage constraints from the four definitions of insertable. Then, as in the proof of Theorem 1, we construct a directed graph G with $V(G)$ equal to the set of all fronts and an edge from F to F' if F' can be obtained from F by a single valid insertion.

We assign two weights, a *cost* and a *profit*, to each edge of this graph G. For any edge in $E(G)$ from some front F to a front F' that represents the insertion of a square s, the cost of the edge is the cost $c(s)$ of s and the profit of the edge is the total profit of the points covered by the insertion of s. For example, for an upper-insertion of a square s intersecting line i, the profit of the edge is the total profit of the points covered by u_i or l_{i+1} in $[x_i, x'_i] \times [i, i+1]$.

Now the sum of the profits of the edges on a F^l–F^r path equals the profit of the solution corresponding to this path. Moreover, the sum of the costs of the edges of the path equals the cost of that solution. Hence we aim to find for any $0 \leq r \leq |\mathcal{P}| \cdot p_{\max}$ a lightest path (with respect to edge costs) of total edge profit at least r. A straightforward dynamic programming algorithm for this problem takes $\mathrm{O}(|E(G)| \cdot |\mathcal{P}| \cdot p_{\max}) = \mathrm{O}((3|\mathcal{S}|)^{4k}(|\mathcal{P}| \cdot p_{\max}))$ time. □

We now apply the shifting technique and scaling to obtain a PTAS. First assume integer profits. For each integer $0 \leq a \leq k-1$, let N_a denote the set of points between lines $y = bk + a$ and $y = bk + a + 1$ for any $b \in \mathbb{Z}$. Moreover, for any $b \in \mathbb{Z}$, let $\mathcal{P}_a^b$ be the set of points between lines $y = bk + a + 1$ and $y = (b+1)k + a$.

For any $0 \leq r \leq |\mathcal{P}| \cdot p_{\max}$, let $C_a^b(r)$ denote the set returned by the algorithm of Theorem 3, applied on $\mathcal{S}$ and $\mathcal{P}_a^b$, attaining profit at least r. We assume that $c(C_a^b(r)) = \infty$ if profit at least r cannot be attained.

Let the nonempty sets $\mathcal{P}_a^b$ be numbered $\mathcal{P}_a^0, \ldots, \mathcal{P}_a^{l_a}$ in an arbitrary way, and let $C_a^0, \ldots, C_a^{l_a}$ be the corresponding solutions. Define

$$\begin{aligned} \mathsf{s}_a(0, r) &= c(C_a^0(r)) \\ \mathsf{s}_a(b, r) &= \min_{0 \leq r' \leq r} \{c(C_a^b(r')) + s_a(b-1, r-r')\} \end{aligned}$$

for $1 \leq b \leq l_a$ and $0 \leq r \leq |\mathcal{P}| \cdot p_{\max}$. Observe that computing s_a can be done in $\mathrm{O}(|\mathcal{P}| \cdot (|\mathcal{P}| \cdot p_{\max})^2)$ time.

Let C_a denote a set attaining $\max_{0 \leq r \leq |\mathcal{P}| \cdot p_{\max}} \{r \mid \mathsf{s}_a(l_a, r) \leq B\}$ and let $C_{\max}$ denote a most profitable such set. By definition, $c(C_{\max}) \leq B$.

Lemma 7. *$p(C_{\max}) \geq (1 - 1/k) \cdot p(OPT)$, where OPT is an optimal solution.*

Proof. Let $\mathcal{S}_a^b$ denote the set of squares in $\mathcal{S}$ covering at least one point in $\mathcal{P}_a^b$. Then it can be easily seen that $c\left(C_a^b\left(p\left(cov\left(OPT \cap \mathcal{S}_a^b\right) \cap \mathcal{P}_a^b\right)\right)\right) \leq c(OPT \cap \mathcal{S}_a^b)$ for any $0 \leq a \leq k-1$ and $0 \leq b \leq l_a$. Because for fixed a the sets $\mathcal{S}_a^b$ are pairwise disjoint, $\sum_{b=0}^{l_a} c(OPT \cap \mathcal{S}_a^b) \leq B$. Then it follows from the definition of s and by induction that $p(C_a) \geq \sum_{b=0}^{l_a} p(cov(OPT \cap \mathcal{S}_a^b) \cap \mathcal{P}_a^b)$. Since we can show that $\sum_{b=0}^{l_a} p(cov(OPT \cap \mathcal{S}_a^b) \cap \mathcal{P}_a^b) = p(OPT) - p(cov(OPT) \cap N_a)$ and any point is in N_a for precisely one value of a, $k \cdot p(C_{\max}) \geq \sum_{a=0}^{k-1} p(C_a) \geq (k-1) \cdot p(OPT)$. It follows immediately that $p(C_{\max}) \geq (1 - 1/k) \cdot p(OPT)$. $\square$

Using scaling for noninteger profits, we obtain a PTAS. Details can be found in [31].

Theorem 4. *There is a PTAS for Geometric Budgeted Maximum Coverage on unit squares.*

4 Optimality and Relation to Domination

Geometric Set Cover and the geometric version of Minimum Dominating Set are closely related. An instance of Minimum Dominating Set on an intersection graph of unit squares can be easily transformed into an instance of Geometric Set Cover on unit squares [28]. Then the following is immediate from Theorem 2 and the remarks at the beginning of Section 3.

Theorem 5. *There is a PTAS for Minimum-Weight Dominating Set on intersection graphs of unit squares.*

Theorem 5 is the first PTAS for Minimum-Weight Dominating Set on intersection graphs of two-dimensional objects. Another consequence of the above reduction from Minimum Dominating Set on unit square graphs to Geometric Set Cover on unit squares is the following. Recall that the exponential time hypothesis (ETH) states that n-variable 3SAT cannot be decided in $2^{o(n)}$ time.

Theorem 6. *If there exist constants $\delta \geq 1$, $0 < \beta < 1$ such that Geometric Set Cover or Geometric Budgeted Maximum Coverage on n unit squares has a PTAS with running time $2^{O(1/\epsilon)^{\delta}} n^{O(1/\epsilon)^{1-\beta}}$, then ETH is false.*

This holds, as Marx [27] showed that Minimum Dominating Set on intersection graphs of unit squares cannot have such a PTAS. Similarly, one can show from Marx [26] that Geometric Set Cover and Geometric Budgeted Maximum Coverage on unit squares have no EPTAS.

Theorem 7. *Geometric Set Cover and Geometric Budgeted Maximum Coverage on unit squares cannot have an EPTAS, unless FPT=$W[1]$.*

This is an indication that one cannot hope to improve the running time of the algorithms of Theorems 2 and 4.

5 Hardness of Approximation

Not many explicit inapproximability results for Geometric Set Cover problems can be found in the literature. Our approximation scheme settles the approximability of Geometric Set Cover on unit squares. In this section, we adapt known results for related problems to give several hardness results for more general shapes. A convex subset s of $\mathbb{R}^2$ is *α-fat* for some $\alpha \geq 1$ if the ratio between the radii of the smallest disk enclosing s and the largest disk inscribed in s is at most α [8].

Theorem 8. *Geometric Set Cover is not approximable within $(1-\epsilon)\ln n$ for any $\epsilon > 0$, unless NP $\subset$ DTIME($n^{O(\log\log n)}$), on convex polygons, translated copies of a single polygon, rotated copies of a single convex polygon, and α-fat objects for any $\alpha > 1$, where n is the number of points,*

Theorem 9. *Geometric Set Cover is APX-hard on convex polygons with $r \geq 4$ corners, α-fat objects of constant description complexity for any $\alpha > 1$, axis-parallel rectangles, and ellipses.*

Theorem 8 and Theorem 9 can be proved using constructions where points are arranged on a line or circle and the objects can cover arbitrary subsets (of bounded size, in case of Theorem 9) of these points. On axis-parallel rectangles and ellipses we need a more elaborate construction. The ideas are similar to ones used for the geometric version of Minimum Dominating Set [10]. See [31] for details. We remark that Har-Peled [15] recently showed that Geometric Set Cover is even APX-hard on fat convex polygons with $r \geq 3$ corners.

Using ideas from Khuller et al. [22], one can obtain the following.

Corollary 1. *Geometric Budgeted Maximum Coverage is not approximable with ratio better than $(1-1/e)$, unless NP $\subset$ DTIME($n^{O(\log\log n)}$), on convex polygons, translated copies of a single polygon, rotated copies of a single convex polygon, and α-fat objects for any $\alpha > 1$.*

6 Conclusions

We have given the first PTAS for Weighted Geometric Set Cover, in the case of axis-parallel unit squares. The scheme extends to Geometric Budgeted Maximum Coverage. Moreover, we presented evidence that one cannot hope to improve on the running time of these algorithms. This settles the approximability of Geometric Set Cover on unit squares.

Many problems surrounding Weighted Geometric Set Cover remain open however. In particular, the question of a PTAS on (unit) disks or arbitrary squares is very interesting. The techniques in this paper seem insufficient to deal with these problems and probably completely different insight is required. In general, it is an interesting question for which objects (Weighted) Geometric Set Cover can still be approximated well. The hardness results of this paper however set clear limits to its approximability.

References

1. Ambühl, C., Erlebach, T., Mihalák, M., Nunkesser, M., "Constant-Factor Approximation for Minimum-Weight (Connected) Dominating Sets in Unit Disk Graphs" in Díaz, J., Jansen, K., Rolim, J.D.P., Zwick, U. (eds.) *APPROX-RANDOM '06*, LNCS 4110, Springer-Verlag, Berlin, 2006, pp. 3–14.
2. Brönnimann, H., Goodrich, M.T., "Almost Optimal Set Covers in Finite VC-Dimension", *Discrete Comput. Geom.* 14:1 (1995), pp. 463–479.
3. Carmi, P., Katz, M.J., Lev-Tov, N., "Covering Points by Unit Disks of Fixed Location" in Tokuyama, T. (ed.) *Proc. ISAAC '07*, LNCS 4835, Springer-Verlag, Berlin, 2007, pp. 644–655.
4. Chvátal, V., "A Greedy Heuristic for the Set-Covering Problem", *Math. Oper. Research* 4:3 (1979), pp. 233–235.
5. Clarkson, K.L., Varadarajan, K.R., "Improved Approximation Algorithms for Geometric Set Cover", *Discrete Comput. Geom.* 37:1 (2007), pp. 43–58.
6. Călinescu, G., Mandoiu, I.I., Wan, P.J., Zelikovsky, A., "Selecting forwarding neighbors in wireless ad hoc networks" in *Proc. DIAL-M '01*, ACM, 2001, pp. 34–43.
7. Dai, D., Yu, C., "A $5+\epsilon$-approximation algorithm for minimum weighted dominating set in unit disk graph", *Theoret. Comput. Sci.* 410:8–10 (2009), pp. 756–765.
8. Efrat, A., Sharir, M., "The Complexity of the Union of Fat Objects in the Plane", *Discrete Comput. Geom.* 23:2 (2000), pp. 171–189.
9. Erlebach, T., Mihalák, M., "A (4+epsilon)-Approximation for the Minimum-Weight Dominating Set Problem in Unit Disk Graphs", *Proc. WAOA 2009*, LNCS 5893, Springer-Verlag, Berlin, pp. 135–146.
10. Erlebach, T., van Leeuwen, E.J., "Domination in Geometric Intersection Graphs" in Laber, E.S., Bornstein, C.F., Nogueira, L.T., Faria, L. (eds.) *Proc. LATIN 2008*, LNCS 4957, Springer-Verlag, Berlin, 2008, pp. 747–758.
11. Feige, U., "A Threshold of $\ln n$ for Approximating Set Cover", *J. ACM* 45:4 (1998), pp. 634–652.
12. Fowler, R.J., Paterson, M.S., Tanimoto, S.L., "Optimal Packing and Covering in the Plane are NP-Complete", *Inform. Process. Lett.* 12:3 (1981), pp. 133–137.
13. Glaßer, C., Reitwießner, C., Schmitz, H., "Multiobjective Disk Cover Admits a PTAS" in Hong, S.H., Nagamochi, H., Fukunaga, T. (eds.) *Proc. ISAAC '08*, LNCS 5369, Springer-Verlag, Berlin, 2008, pp. 40–51.
14. Gonzalez, T.F., "Covering a Set of Points in Multidimensional Space", *Inform. Process. Lett.* 40:4 (1991), pp. 181–188.
15. Har-Peled, S., "Being Fat and Friendly is Not Enough", arXiv:0908.2369v1 [cs.CG], 2009.

16. Hochbaum, D.S., "Approximating Covering and Packing Problems: Set Cover, Vertex Cover, Independent Set, and Related Problems" in Hochbaum, D.S. (ed.) *Approximation Algorithms for NP-hard Problems*, PWS Publishing Company, Boston, 1997, pp. 46–93.
17. Hochbaum, D.S., Maass, W., "Approximation Schemes for Covering and Packing Problems in Image Processing and VLSI", *J. ACM* 32:1 (1985), pp. 130–136.
18. Hochbaum, D.S., Pathria, A., *Analysis of the greedy approach in covering problems*, unpublished manuscript, 1994.
19. Huang, Y., Gao, X., Zhang, Z., Wu, W., "A Better Constant-Factor Approximation for Weighted Dominating Set in Unit Disk Graph", *J. Combin. Optim.* 18:2 (2009), pp. 179–194.
20. Johnson, D.S., "Approximation algorithms for combinatorial problems", *J. Comput. System Sci.* 9:3 (1974), pp. 256–278.
21. Johnson, D.S., "The NP-Completeness Column: An Ongoing Guide", *J. Algorithms* 3:2 (1982), pp. 182–195.
22. Khuller, S., Moss, A., Naor, J.S., "The Budgeted Maximum Coverage Problem", *Inform. Process. Lett.* 70:1 (1999), pp. 39–45.
23. Laue, S., "Geometric Set Cover and Hitting Sets for Polytopes in $\mathbb{R}^3$" in Albers, S., Weil, P. (eds.) *Proc. STACS 2008*, LIPIcs 1, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2008, pp. 479–490.
24. Liao, C., Hu, S., "Polynomial time approximation schemes for minimum disk cover problems", *J. Global Optim.* (2009), doi: 10.1007/s10878-009-9216-y.
25. Lovász, L., "On the ratio of optimal integral and fractional covers", *Discrete Math.* 13:4 (1975), pp. 383–390.
26. Marx, D., "Parameterized Complexity of Independence and Domination Problems on Geometric Graphs" in Bodlaender, H.L., Langston, M.A. (eds.) *Proc. IWPEC '06*, LNCS 4169, Springer-Verlag, Berlin, 2006, pp. 154–165.
27. Marx, D., "On the Optimality of Planar and Geometric Approximation Schemes" in *Proc. FOCS '07*, IEEE, 2007, pp. 338–348.
28. Mihalák, M., *Optimization Problems in Communication Networks*, PhD thesis, Univ. Leicester, 2006.
29. Mustafa, N., Ray, S., "PTAS for Geometric Hitting Set Problems via Local Search" in *Proc. SoCG '09*, ACM, New York, 2009, pp. 17–22.
30. Narayanappa, S., Vojtěchovský, P., "An Improved Approximation Factor For The Unit Disk Covering Problem" in *Proc. CCCG '06*, 2006.
31. van Leeuwen, E.J., *Optimization and Approximation on Systems of Geometric Objects*, PhD thesis, University of Amsterdam, 2009.
32. Varadarajan, K., "Weighted Geometric Set Cover via Quasi-Uniform Sampling" in *Proc. STOC 2010*, ACM, New York, 2010, pp. 641–648.
33. Vohra, R., Hall, N.G., "A Probabilistic Analysis of the Maximal Covering Location Problem", *Discrete Appl. Math.* 43:2 (1993), pp. 175–183.
34. Zou, F., Wang, Y., Xu, X.-H., Li, X., Du, H., Wan, P., Wu, W., "New approximations for minimum-weighted dominating sets and minimum-weighted connected dominating sets on unit disk graphs", *Theoret. Comput. Sci.*, to appear.