

Faster Approximation Algorithms for Geometric Set Cover

Timothy M. Chan 

Department of Computer Science, University of Illinois at Urbana-Champaign, IL, USA
tmc@illinois.edu

Qizheng He

Department of Computer Science, University of Illinois at Urbana-Champaign, IL, USA
qizheng6@illinois.edu

Abstract

We improve the running times of $O(1)$ -approximation algorithms for the set cover problem in geometric settings, specifically, covering points by disks in the plane, or covering points by halfspaces in three dimensions. In the unweighted case, Agarwal and Pan [SoCG 2014] gave a randomized $O(n \log^4 n)$ -time, $O(1)$ -approximation algorithm, by using variants of the multiplicative weight update (MWU) method combined with geometric data structures. We simplify the data structure requirement in one of their methods and obtain a *deterministic* $O(n \log^3 n \log \log n)$ -time algorithm. With further new ideas, we obtain a still faster randomized $O(n \log n (\log \log n)^{O(1)})$ -time algorithm.

For the weighted problem, we also give a randomized $O(n \log^4 n \log \log n)$ -time, $O(1)$ -approximation algorithm, by simple modifications to the MWU method and the quasi-uniform sampling technique.

2012 ACM Subject Classification Theory of computation \rightarrow Computational geometry

Keywords and phrases Set cover, approximation algorithms, multiply weight update method, random sampling, shallow cuttings

Digital Object Identifier 10.4230/LIPIcs.SoCG.2020.27

Related Version A full version of this paper is available at <http://arxiv.org/abs/2003.13420>.

Funding *Timothy M. Chan*: Supported in part by NSF Grant CCF-1814026.

1 Introduction

Unweighted geometric set cover. In this paper we study one of the most fundamental classes of geometric optimization problems: *geometric set cover*. Given a set X of $O(n)$ points and a set S of $O(n)$ geometric objects, find the smallest subset of objects from S to cover all points in X . In the dual set system, the problem corresponds to *geometric hitting set* (finding the smallest number of points from X that hit all objects in S).

This class of problems has been *extensively* investigated in the computational geometry literature. Since they are NP-hard in most scenarios, attention is turned towards approximation algorithms. Different types of objects give rise to different results. Typically, approximation algorithms fall into the following categories:

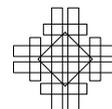
1. Simple heuristics, e.g., greedy algorithms.
2. Approaches based on solving the linear programming (LP) relaxation (i.e., fractional set cover) and rounding the LP solution.
3. Polynomial-time approximation schemes (PTASs), e.g., via local search, shifted grids/quadtrees (sometimes with dynamic programming), or separator-based divide-and-conquer.



© Timothy M. Chan and Qizheng He;
licensed under Creative Commons License CC-BY
36th International Symposium on Computational Geometry (SoCG 2020).
Editors: Sergio Cabello and Danny Z. Chen; Article No. 27; pp. 27:1–27:14
Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



Generally, greedy algorithms achieve only logarithmic approximation factors (there are some easy cases where they give $O(1)$ approximation factors, e.g., hitting set for fat objects such as disks/balls in the “continuous” setting with $X = \mathbb{R}^d$ [20]). The LP-based approaches give better approximation factors in many cases, e.g., $O(1)$ approximation for set cover and hitting set for disks in 2D and halfspaces in 3D, set cover for objects in 2D with linear “union complexity”, and hitting set for pseudodisks in 2D [7, 19, 5, 32, 29]. Subsequently, local-search PTASs have been found by Mustafa and Ray [27] in some cases, including set cover and hitting set for disks in 2D and halfspaces in 3D (earlier, PTASs were known for hitting set only in the continuous setting for unit disks/balls [21], and for arbitrary disks/balls and fat objects [11]).

Historically, the focus has been on obtaining good approximation factors. Here, we are interested in obtaining approximation algorithms with good – ideally, near linear – running time. Concerning the efficiency of known approximation algorithms:

1. Certain simple heuristics can lead to fast $O(1)$ -approximation algorithms in some easy cases (e.g., continuous hitting set for unit disks or unit balls by using grids), but generally, even simple greedy algorithms may be difficult to implement in near linear time (as they may require nontrivial dynamic geometric data structures).
2. LP-based approaches initially may not seem to be the most efficient, because of the need to solve an LP. However, a general-purpose LP solver can be avoided. The set-cover LP can alternatively be solved (approximately) by the *multiplicative weight update* (MWU) method. In the computational geometry literature, the technique has been called *iterative reweighting*, and its use in geometric set cover was explored by Brönnimann and Goodrich [7] (one specific application appeared in an earlier work by Clarkson [18]), although the technique was known even earlier outside of geometry. On the other hand, the LP-rounding part corresponds to the well-known geometric problem of constructing ε -nets, for which efficient algorithms are known [15, 25].
3. PTAS approaches generally have large polynomial running time, even when specialized to specific approximation factors. For example, see [8] for efforts in improving the degree of the polynomial.

In this paper we design faster approximation algorithms for geometric set cover via the LP/MWU-based approaches. There has been a series of work on speeding up MWU methods for covering or packing LPs (e.g., see [17, 22, 33]). In geometric settings, we would like more efficient algorithms (as generating the entire LP explicitly would already require quadratic time), by somehow exploiting geometric data structures. The main previous work was by Agarwal and Pan [4] from SoCG 2014, who showed how to compute an $O(1)$ -approximation for set cover for 2D disks or 3D halfspaces, in $O(n \log^4 n)$ randomized time.

Agarwal and Pan actually proposed two MWU-based algorithms: The first is a simple variant of the standard MWU algorithm of Brönnimann and Goodrich, which proceeds in logarithmically many rounds. The second views the problem as a 2-player zero-sum game, works quite differently (with weight updates to both points and objects), and uses randomization; the analysis is more complicated. Because the first algorithm requires stronger data structures – notably, for approximate weighted range counting with dynamic changes to the weights – Agarwal and Pan chose to implement their second algorithm instead, to get their $O(n \log^4 n)$ result for 3D halfspaces.

New results. In this paper we give:

- a *deterministic* near-linear $O(1)$ -approximation algorithm for set cover for 3D halfspaces. Its running time is $O(n \log^3 n \log \log n)$, which besides eliminating randomization is also a little faster than Agarwal and Pan’s;

- a still faster randomized near-linear $O(1)$ -approximation algorithm for set cover for 3D halfspaces. Its running time is $O(n \log n \log^{O(1)} \log n)$, which is essentially optimal¹ ignoring minor $\log \log n$ factors.

Although generally shaving logarithmic factors may not be the most important endeavor, the problem is fundamental enough that we feel it worthwhile to find the most efficient algorithm possible.

Our approach interestingly is to go back to Agarwal and Pan's first MWU algorithm. We show that with one simple modification, the data structure requirement can actually be relaxed: namely, for the approximate counting structure, there is no need for weights, and the only update operation is insertion. By standard techniques, insertion-only data structures reduce to static data structures. This simple idea immediately yields our deterministic result. (Before, Bus et al. [9] also aimed to find variants of Agarwal and Pan's first algorithm with simpler data structures, but they did not achieve improved theoretical time bounds.) Our best randomized result requires a more sophisticated combination of several additional ideas. In particular, we incorporate random sampling in the MWU algorithm, and extensively use *shallow cuttings*, in both primal and dual space.

We have stated our results for set cover for 3D halfspaces. This case is arguably the most central. It is equivalent to hitting set for 3D halfspaces, by duality, and also includes set cover and hitting set for 2D disks as special cases, by the standard lifting transformation. The case of 3D dominance ranges is another special case, by a known transformation [14, 28] (although for the dominance case, word-RAM techniques can speed up the algorithms further). The ideas here are likely useful also in the other cases considered in Agarwal and Pan's paper (e.g., hitting set for rectangles, set cover for fat triangles, etc.), but in the interest of keeping the paper focused, we will not discuss these implications.

Weighted geometric set cover. Finally, we consider the weighted version of set cover: assuming that each object is given a weight, we now want a subset of the objects of R with the minimum total weight that covers all points in X . The weighted problem has also received considerable attention: Varadarajan [31] and Chan et al. [13] used the LP-based approach to obtain $O(1)$ -approximation algorithms for weighted set cover for 3D halfspaces (or for objects in 2D with linear union complexity); the difficult part is in constructing ε -nets with small weights, which they solved by the *quasi-random sampling* technique. Later, Mustafa, Raman, and Ray [26] discovered a quasi-PTAS for 3D halfspaces by using geometric separators; the running time is very high ($n^{\log^{O(1)} n}$).

Very recently, Chekuri, Har-Peled, and Quanrud [16] described new randomized MWU methods which can efficiently solve the LP corresponding to various generalizations of geometric set cover, by using appropriate geometric data structures. In particular, for weighted set cover for 3D halfspaces, they obtained a randomized $O(n \log^{O(1)} n)$ -time algorithm to solve the LP but with an unspecified number of logarithmic factors. They did not address the LP-rounding part, i.e., construction of an ε -net of small weight – a direct implementation of the quasi-uniform sampling technique would not lead to a near-linear time bound.

We observe that a simple direct modification of the standard MWU algorithm of Brönnimann and Goodrich, or Agarwal and Pan's first algorithm, can also solve the LP for weighted geometric set cover, with arguably simpler data structures than Chekuri et al.'s. Secondly,

¹ Just deciding whether a solution exists requires $\Omega(n \log n)$ time in the algebraic decision-tree model, even for 1D intervals.

we observe that an ε -net of small weight can be constructed in near-linear time, by using quasi-uniform sampling more carefully. This leads to a randomized $O(n \log^4 n \log \log n)$ -time, $O(1)$ -approximation algorithm for weighted set cover for 3D halfspaces (and thus for 2D disks).

2 Preliminaries

Let X be a set of points and S be a set of objects. For a point p , its *depth* in S refers to the number of objects in S containing p . A point p is said to be ε -*light* in S , if it has depth $\leq \varepsilon|S|$ in S ; otherwise it is ε -*heavy*. A subset of objects $T \subseteq S$ is an ε -*net* of S if T covers all points that are ε -heavy in S .

It is known that there exists an ε -net with size $O(\frac{1}{\varepsilon})$ for any set of halfspaces in 3D or disks in 2D [24] (or more generally for objects in the plane with linear union complexity [19]).

2.1 The Basic MWU Algorithm

We first review the standard multiplicative weight² update (MWU) algorithm for geometric set cover, as described by Brönnimann and Goodrich [7] (which generalizes an earlier algorithm by Clarkson [18], and is also well known outside of computational geometry).

Let X be the set of input points and S be the set of input objects, with $n = |X| + |S|$. Let OPT denote the size of the minimum set cover. We assume that a value $t = \Theta(\text{OPT})$ is known; this assumption will be removed later by a binary search for t . In the following pseudocode, we work with a multiset \hat{S} ; in measuring size or counting depth, we include multiplicities (e.g., $|\hat{S}|$ is the sum of the multiplicities of all its elements).

-
- 1: Guess a value $t \in [\text{OPT}, 2 \text{OPT}]$ and set $\varepsilon = \frac{1}{2t}$.
 - 2: Define a multiset \hat{S} where each object i in S initially has multiplicity $m_i = 1$.
 - 3: **while** we can find a point $p \in X$ which is ε -light in \hat{S} **do**
 - 4: **for** each object i containing p **do** \triangleright call lines 4–5 a *multiplicity-doubling step*
 - 5: Double its multiplicity m_i .
 - 6: Return an ε -net of the multiset \hat{S} .
-

Since at the end all points in X are ε -heavy in \hat{S} , the returned subset is a valid set cover of X . For halfspaces in 3D or disks in 2D, its size is $O(\frac{1}{\varepsilon}) = O(t) = O(\text{OPT})$.

A standard analysis shows that the algorithm always terminates after $O(t \log \frac{n}{t})$ multiplicity-doubling steps. We include a quick proof: Each multiplicity-doubling step increases $|\hat{S}|$ by a factor of at most $1 + \varepsilon$, due to the ε -lightness of p . Thus, after z doubling steps, $|\hat{S}| \leq n(1 + \varepsilon)^z \leq ne^{\varepsilon z} = ne^{z/(2t)}$. On the other hand, consider a set cover T^* of size t . In each multiplicity-doubling step, at least one of the objects in T^* has its multiplicity doubled. So, after z multiplicity-doubling steps, the total multiplicity in T^* is at least $t2^{z/t}$. We conclude that $t2^{z/t} \leq |\hat{S}| \leq ne^{z/(2t)}$, implying that $z = O(t \log \frac{n}{t})$.

² In our algorithm description, we prefer to use the term “multiplicity” instead of “weight”, to avoid confusion with the weighted set cover problem later.

2.2 Agarwal and Pan's (First) MWU Algorithm

Next, we review Agarwal and Pan's first variant of the MWU algorithm [4]. One issue in implementing the original algorithm lies in the test in line 3: searching for one light point by scanning all points in X from scratch every time seems inefficient. In Agarwal and Pan's refined approach, we proceed in a small number of rounds, where in each round, we examine the points in X in a fixed order and test for lightness in that order.

```

1: Guess a value  $t \in [\text{OPT}, 2\text{OPT}]$  and set  $\varepsilon = \frac{1}{2t}$ .
2: Define a multiset  $\hat{S}$  where each object  $i$  in  $S$  initially has multiplicity  $m_i = 1$ .
3: loop ▷ call this the start of a new round
4:   for each point  $p \in X$  in any fixed order do
5:     while  $p$  is  $\varepsilon$ -light in  $\hat{S}$  do
6:       for each object  $i$  containing  $p$  do ▷ call lines 6–7 a multiplicity-doubling step
7:         Double its multiplicity  $m_i$ .
8:       if the number of multiplicity-doubling steps in this round exceeds  $t$  then
9:         Go to line 3 and start a new round.
10:  Terminate and return an  $\frac{\varepsilon}{2}$ -net of the multiset  $\hat{S}$ .

```

To justify correctness, observe that since each round performs at most t multiplicity-doubling steps, $|\hat{S}|$ increases by a factor of at most $(1 + \varepsilon)^t \leq e^{\varepsilon t} \leq e^{1/2} < 2$. Thus, a point p that is checked to be ε -heavy in \hat{S} at any moment during the round will remain $\frac{\varepsilon}{2}$ -heavy in \hat{S} at the end of the round.

Since all but the last round performs t multiplicity-doubling steps and we have already shown that the total number of such steps is $O(t \log \frac{n}{t})$, the number of rounds is $O(\log \frac{n}{t})$.

3 “New” MWU Algorithm

Agarwal and Pan's algorithm still requires an efficient data structure to test whether a given point is light, and the data structure needs to support dynamic changes to the multiplicities. We propose a new variant that requires simpler data structures.

Our new algorithm is almost identical to Agarwal and Pan's, but with just one very simple change! Namely, after line 3, at the beginning of each round, we add the following line, to readjust all multiplicities:

```

3.5: for each object  $i$ , reset its multiplicity  $m_i \leftarrow \lceil m_i \frac{10n}{|\hat{S}|} \rceil$ .

```

To analyze the new algorithm, consider modifying the multiplicity m_i instead to $\lceil m_i \frac{10n}{|\hat{S}|} \rceil \cdot \frac{|\hat{S}|}{10n}$. The algorithm behaves identically (since the multiplicities are identical except for a common rescaling factor), but is more convenient to analyze. In this version, multiplicities are nondecreasing over time (though they may be non-integers). After the modified line 3.5, the new $|\hat{S}|$ is at most $\sum_i \left(m_i \frac{10n}{|\hat{S}|} + 1 \right) \cdot \frac{|\hat{S}|}{10n} \leq 11n \cdot \frac{|\hat{S}|}{10n} = 1.1|\hat{S}|$. If the algorithm makes z multiplicity-doubling steps, then it performs line 3.5 at most z/t times and we now have $|\hat{S}| \leq n(1 + \varepsilon)^z \cdot 1.1^{z/t} \leq ne^{z/(2t)} \cdot 1.1^{z/t}$. This is still sufficient to imply that $z = O(t \log \frac{n}{t})$, and so the number of rounds remains $O(\log \frac{n}{t})$.

Now, let's go back to line 3.5 as written. The advantage of this multiplicity readjustment step is that it decreases $|\hat{S}|$ to $\sum_i \left(m_i \frac{10n}{|\hat{S}|} + 1\right) = O(n)$. At the end of the round, $|\hat{S}|$ increases by a factor of at most $(1 + \varepsilon)^t < 2$ and so remains $O(n)$. Thus, in line 7, instead of doubling the multiplicity of an object, we can just repeatedly increment the multiplicity (i.e., insert one copy of an object) to reach the desired value. The total number of increments per round is $O(n)$.

Note that in testing for ε -lightness in line 5, a constant-factor approximation of the depth is sufficient, with appropriate adjustments of constants in the algorithm. Also, although the algorithm as described may test the same point p for lightness several times in a round, this can be easily avoided: we just keep track of the increase D in the depth of the current point p ; the new depth of p can be 2-approximated by the maximum of the old depth and D .

To summarize, an efficient implementation of each round of the new algorithm requires solving the following geometric data structure problems (REPORT for line 6, and APPROX-COUNT-DECISION for line 5):

Problem Report: Design a data structure to store a static set S of size $O(n)$ so that given a query point $p \in X$, we can report all objects in S containing the query point p . Here, the output size of a query is guaranteed to be at most $O(k)$ where $k := \frac{n}{t}$ (since ε -lightness of p implies that its depth is at most $\varepsilon|\hat{S}| = \Theta(\frac{n}{t})$ even including multiplicities).

Problem Approx-Count-Decision: Design a data structure to store a multiset \hat{S} of size $O(n)$ so that given a query point $p \in X$, we can either declare that the number of objects in \hat{S} containing p is less than a fixed threshold value k , or that the number is more than $\frac{k}{c}$, for some constant $c > 1$. Here, the threshold again is $k := \frac{n}{t}$ (since $\varepsilon|\hat{S}| = \Theta(\frac{n}{t})$). The data structure should support the following type of updates: insert one copy of an object to \hat{S} . (Deletions are not required.) Each point in X is queried once.

To bound the cost of the algorithm:

- Let T_{report} denote the total time for $O(t)$ queries in Problem REPORT.
- Let T_{count} denote the total time for $O(n)$ queries and $O(n)$ insertions in Problem APPROX-COUNT-DECISION. (Note that the initialization of \hat{S} at the beginning of the round can be done by $O(n)$ insertions.)
- Let T_{net} denote the time for computing an ε -net of size $O(\frac{1}{\varepsilon})$ for a given multiset \hat{S} of size $O(n)$.

The total running time over all $O(\log \frac{n}{t})$ rounds is

$$O((T_{\text{report}} + T_{\text{count}}) \log \frac{n}{t} + T_{\text{net}}). \quad (1)$$

4 Implementations

In this section, we describe specific implementations of our MWU algorithm when the objects are halfspaces in 3D (which include disks in 2D as a special case by the standard lifting transformation). We first consider deterministic algorithms.

4.1 Deterministic Version

Shallow cuttings. We begin by reviewing an important tool that we will use several times later. For a set of n planes in \mathbb{R}^3 , a k -shallow ε -cutting is a collection of interior-disjoint polyhedral cells, such that each cell intersects at most εn planes, and the union of the cells cover all points of level at most k (the *level* of a point refers to the number of planes below it). The list of all planes intersecting a cell Δ is called the *conflict list* of Δ . Matoušek [23]

proved the existence of a k -shallow $(\frac{ck}{n})$ -cutting with $O(\frac{n}{k})$ cells for any constant c . Chan and Tsakalidis [15] gave an $O(n \log \frac{n}{k})$ -time deterministic algorithm to construct such a cutting, along with all its conflict lists (an earlier randomized algorithm was given by Ramos [30]). If c is sufficiently large, the cells may be made “downward”, i.e., they all contain $(0, 0, -\infty)$.

Constructing ε -nets. The best known deterministic algorithm for constructing ε -nets for 3D halfspaces is by Chan and Tsakalidis [15] and runs in $T_{\text{net}} = O(n \log \frac{1}{\varepsilon}) = O(n \log n)$ time.

The result follows directly from their shallow cutting algorithm (using a simple argument of Matoušek [23]): Without loss of generality, assume that all halfspaces are upper halfspaces, so depth corresponds to level with respect to the bounding planes (we can compute a net for lower halfspaces separately and take the union, with readjustment of ε by a factor of 2). We construct an (εn) -shallow $\frac{\varepsilon}{2}$ -cutting with $O(\frac{1}{\varepsilon})$ cells, and for each cell, add a plane completely below the cell (if it exists) to the net. To see correctness, for a point p with level εn , consider the cell Δ containing p ; at least $\varepsilon n - \frac{\varepsilon n}{2} > 0$ planes are completely below Δ , and so the net contains at least one plane below p .

Solving Problem Report. This problem corresponds to 3D *halfspace range reporting* in dual space, and by known data structures [10, 2, 15], the total time to answer $O(t)$ queries is $T_{\text{report}} = O(t \cdot (\log n + k)) = O(t \log n + n)$, assuming an initial preprocessing of $O(n \log n)$ time (which is done only once).

This result also follows directly from shallow cuttings (since space is not our concern, the solution is much simplified): Without loss of generality, assume that all halfspaces are upper halfspaces. We construct a k -shallow $O(\frac{k}{n})$ -cutting with $O(\frac{n}{k})$ downward cells. Given a query point $p \in X$, we find the cell containing p , which can be done in $O(\log n)$ time by planar point location; we then do a linear search over its conflict list, which has size $O(k)$.

Note that the point location operations can be actually be done during preprocessing in $O(n \log n)$ time since X is known in advance. This lowers the time bound for $O(t)$ queries to $T_{\text{report}} = O(tk) = O(n)$.

Solving Problem Approx-Count-Decision. This problem corresponds to the decision version of 3D *halfspace approximate range counting* in dual space, and several deterministic and randomized data structures have already been given in the static case [1, 3], achieving $O(\log n)$ query time and $O(n \log n)$ preprocessing time.

This result also follows directly from shallow cuttings: Without loss of generality, assume that all halfspaces are upper halfspaces. We construct a $\frac{n}{b^i}$ -shallow $O(\frac{1}{b^i})$ -cutting with $O(b^i)$ downward cells for every $i = 1, \dots, \log_b n$ for some constant b . Chan and Tsakalidis’s algorithm can actually construct all $O(\log n)$ such cuttings in $O(n \log n)$ total time. With these cuttings, we can compute an $O(1)$ -approximation to the depth/level of a query point p by simply finding the largest i such that p is contained in a cell of the $\frac{n}{b^i}$ -shallow cutting (the level of p would then be $O(\frac{n}{b^i})$ and at least $\frac{n}{b^{i+1}}$). In Chan and Tsakalidis’s construction, each cell in one cutting intersects $O(1)$ cells in the next cutting, and so we can locate the cells containing p in $O(1)$ time per i , for a total of $O(\log n)$ time.

To solve Problem APPROX-COUNT-DECISION, we still need to support insertion. Although the approximate decision problem is not decomposable, the above solution solves the approximate counting problem, which is decomposable, so we can apply the standard *logarithmic method* [6] to transform the static data structure into a semi-dynamic, insertion-only data structure. The transformation causes a logarithmic factor increase, yielding in our case $O(\log^2 n)$ query time and $O(\log^2 n)$ insertion time. Thus, the total time for $O(n)$ queries and insertions is $T_{\text{count}} = O(n \log^2 n)$.

Conclusion. By (1), the complete algorithm has running time $O((T_{\text{report}} + T_{\text{count}}) \log \frac{n}{t} + T_{\text{net}}) = O((n + n \log^2 n) \log \frac{n}{t} + n \log n) = O(n \log^3 n)$.

One final issue remains: we have assumed that a value $t \in [\text{OPT}, 2\text{OPT}]$ is given. In general, either the algorithm produces a solution of size $O(t)$, or (if it fails to complete within $O(\log \frac{n}{t})$ rounds) the algorithm may conclude that $\text{OPT} > t$. We can thus find an $O(1)$ -approximation to OPT by a binary search over t among the $O(\log n)$ possible powers of 2, with $O(\log \log n)$ calls to the algorithm. The final time bound is $O(n \log^3 n \log \log n)$.

► **Theorem 1.** *Given $O(n)$ points and $O(n)$ halfspaces in \mathbb{R}^3 , we can find a subset of halfspaces covering all points, of size within $O(1)$ factor of the minimum, in deterministic $O(n \log^3 n \log \log n)$ time.*

4.2 Randomized Version 1

We now describe a better solution to Problem APPROX-COUNT-DECISION, by using randomization and the fact that all query points (namely, X) are given in advance.

Reducing the number of insertions in Problem Approx-Count-Decision. In solving Problem APPROX-COUNT-DECISION, one simple way to speed up insertions is to work with a random sample R of \hat{S} . When we insert an object to \hat{S} , we independently decide to insert it to the sample R with probability $\rho := \frac{c_0 \log n}{k}$, or ignore it with probability $1 - \rho$, for a sufficiently large constant c_0 . (Different copies of an object are treated as different objects here.) It suffices to solve the problem for the sample R with the new threshold around ρk .

To justify correctness, consider a fixed query point $p \in X$. Let x_1, x_2, \dots be the sequence of objects in \hat{S} that contain p , in the order in which they are inserted (extend the sequence arbitrarily to make its length greater than k). Let $y_i = 1$ if object x_i is chosen to be in the sample R , or 0 otherwise. Note that the x_i 's may not be independent (since the object we insert could depend on random choices made before); however, the y_i 's are independent. By the Chernoff bound, $\sum_{i=1}^{k/c} y_i \leq \frac{(1+\delta)\rho k}{c}$ and $\sum_{i=1}^k y_i \geq (1-\delta)\rho k$ with probability $1 - e^{-\Omega(\rho k)} = 1 - n^{-\Omega(c_0)}$ for any fixed constant $\delta > 0$. Thus, with high probability, at any time, if the number of objects in R containing p is more than $\frac{(1+\delta)\rho k}{c}$, then the number of objects in \hat{S} containing p is more than $\frac{k}{c}$; if the former number is less than $(1-\delta)\rho k$, then the later number is less than k . Since there are $O(n)$ possible query points, all queries are correct with high probability.

By this strategy, the number of insertions is reduced to $O(\rho n) = O(\frac{n}{k} \log n) = O(t \log n)$ with high probability.

Preprocessing step. Next we use a known preprocessing step to ensure that each object contains at most $\frac{n}{t}$ points, in the case of 3D halfspaces. This subproblem was addressed in Agarwal and Pan's paper [4] (where it was called "(P5)" – curiously, they used it to implement their second algorithm but not their first MWU-based algorithm.) We state a better running time:

► **Lemma 2.** *In $O(n \log t)$ time, we can find a subset $T_0 \subseteq S$ of $O(t)$ halfspaces, such that after removing all points in X covered by T_0 , each halfspace of S contains at most $\frac{n}{t}$ points.*

Proof. We may assume that all halfspaces are upper halfspaces. We work in dual space, where S is now a set of points and X is a set of planes. The goal is to find a subset $T_0 \subseteq S$ of $O(t)$ points such that after removing all planes of X that are below some points of T_0 , each point of S has depth/level at most $\frac{n}{t}$.

We proceed in rounds. Let b be a constant. In the i -th round, assume that all points of S have level $\leq \frac{n}{b^i}$. Compute a $\frac{n}{b^i}$ -shallow $\frac{1}{b^{i+1}}$ -cutting with $O(b^i)$ cells. In each cell, add an arbitrary point of S (if exists) to the set T_0 . In total $O(b^i)$ points are added. Remove all planes that are below these added points from X .

Consider a point p of S . Let Δ be the cell containing p , and let q be the point in Δ that was added to T_0 . Any plane that is below p but not removed (and thus above q) must intersect Δ , so there can be at most $\frac{n}{b^{i+1}}$ such planes. Thus, after the round, the level of p is at most $\frac{n}{b^{i+1}}$. We terminate when b^i reaches t . The total size of T_0 is $O(\sum_{i=1}^{\log_b t} b^i) = O(t)$.

Naively computing each shallow cutting from scratch by Chan and Tsakalidis's algorithm would require $O(n \log n \cdot \log n) = O(n \log^2 n)$ total time. But Chan and Tsakalidis's approach can compute multiple shallow cuttings more quickly: given a $\frac{n}{b^i}$ -shallow cutting along with its conflict lists, we can compute the next $\frac{n}{b^{i+1}}$ -shallow cutting along with its conflict lists in $O(n + b^i \log b^i)$ time. However, in our application, before computing the next cutting, we also remove some of the input planes. Fortunately, this type of scenario has been examined in a recent paper by Chan [12], who shows that the approach still works, provided that the next cutting is relaxed to cover only points covered by the previous cutting (see Lemma 8 in his paper); this is sufficient in our application. In our application, we also need to locate the cell containing each point of S . This can still be done in $O(n)$ time given the locations in the previous cutting. Thus, the total time is $O(\sum_{i=1}^{\log_b t} (n + b^i \log b^i)) = O(n \log t)$. ◀

At the end, we add T_0 back to the solution, which still has $O(t)$ total size.

Solving Problem Approx-Count-Decision. We now propose a very simple approach to solve Problem APPROX-COUNT-DECISION: just explicitly maintain the depth of all points in X . Each query then trivially takes $O(1)$ time. When inserting an object, we find all points contained in the object and increment their depths.

Due to the above preprocessing step, the number of points contained in the object is $O(\frac{n}{t})$. For the case of 3D halfspaces, we can find these points by halfspace range reporting; as explained before for Problem REPORT, this can be done in $O(\frac{n}{t})$ time by using shallow cuttings, after an initial preprocessing in $O(n \log n)$ time. Thus, each insertion takes $O(\frac{n}{t})$ time. Since the number of insertions has been reduced to $O(t \log n)$ by sampling, the total time for Problem APPROX-COUNT-DECISION is $T_{\text{count}} = O((t \log n) \cdot \frac{n}{t}) = O(n \log n)$.

Conclusion. By (1), the complete randomized algorithm has running time $O((T_{\text{report}} + T_{\text{count}}) \log \frac{n}{t} + T_{\text{net}}) = O((n + n \log n) \log \frac{n}{t} + n \log n) = O(n \log n \log \frac{n}{t}) = O(n \log^2 n)$ (even including the $O(n \log n)$ -time preprocessing step). Including the binary search for t , the time bound is $O(n \log^2 n \log \log n)$.

4.3 Randomized Version 2

Finally, we combine the ideas from both the deterministic and randomized implementations, to get our fastest randomized algorithm for 3D halfspaces.

Solving Problem Approx-Count-Decision. We may assume that all halfspaces are upper halfspaces. We work in dual space, where \hat{S} is now a multiset of points and X is a set of planes. In a query, we want to approximately count the number of points in \hat{S} that are above a query plane in X . By the sampling reduction from Section 4.2, we may assume that the number of insertions to \hat{S} is $O(t \log n)$. By the preprocessing step from Section 4.2, we may assume that all points in \hat{S} have level at most $\frac{n}{t}$.

27:10 Faster Approximation Algorithms for Geometric Set Cover

Compute a $\frac{n}{t}$ -shallow $O(\frac{1}{t})$ -cutting with $O(t)$ downward cells, along with its conflict lists. For each point $p \in R$, locate the cell containing p . All this can be done during a (one-time) preprocessing in $O(n \log n)$ time.

For each cell Δ , we maintain $\hat{S} \cap \Delta$ in a semi-dynamic data structure for 3D approximate halfspace range counting. As described in Section 4.1, we get $O(\log^2 n_\Delta)$ query and insertion time, where $n_\Delta = |\hat{S} \cap \Delta|$.

In an insertion of a point p to \hat{S} , we look up the cell Δ containing p and insert the point to the approximate counting structure in Δ .

In a query for a plane $h \in X$, we look up the cells Δ whose conflict lists contain h , answer approximate counting queries in these cells, and sum the answers.

We bound the total time for all insertions and queries. For each cell Δ , the number of insertions in its approximate counting structure is n_Δ and the number of queries is $O(\frac{n}{t})$ (since each plane $h \in X$ is queried once). The total time is

$$O\left(\sum_{\Delta} \left(\frac{n}{t} + n_\Delta\right) \log^2 n_\Delta\right).$$

Since there are $O(t)$ terms and $\sum_{\Delta} n_\Delta = O(t \log n)$, we have $n_\Delta = O(\log n)$ “on average”; applying Jensen’s inequality to the first term, we can bound the sum by $O(n \log^2 \log n + t \log^3 n)$. Thus, $T_{\text{count}} = O(n \log^2 \log n + t \log^3 n)$.

Conclusion. By (1), the complete randomized algorithm has running time $O((T_{\text{report}} + T_{\text{count}}) \log \frac{n}{t} + T_{\text{net}}) = O((n + n \log^2 \log n + t \log^3 n) \log \frac{n}{t} + n \log n) = O(n \log n \log^2 \log n + t \log^4 n)$. If $t \leq n / \log^3 n$, the first term dominates. On the other hand, if $t > n / \log^3 n$, our earlier randomized algorithm has running time $O(n \log n \log \frac{n}{t}) = O(n \log n \log \log n)$. In any case, the time bound is at most $O(n \log n \log^2 \log n)$. Including the binary search for t , the time bound is $O(n \log n \log^3 \log n)$.

► **Theorem 3.** *Given $O(n)$ points and $O(n)$ halfspaces in \mathbb{R}^3 , we can find a subset of halfspaces covering all points, of size within $O(1)$ factor of the minimum, in $O(n \log n \log^3 \log n)$ time by a randomized Monte-Carlo algorithm with error probability $O(n^{-c_0})$ for any constant c_0 .*

Remark. The number of the $\log \log n$ factors is improvable with still more effort, but we feel it is of minor significance.

5 Weighted Set Cover

In this final section, we consider the weighted set cover problem. We define ε -lightness and ε -nets as before, ignoring the weights. It is known that there exists an ε -net of S with total weight $O(\frac{1}{\varepsilon} \cdot \frac{w(S)}{|S|})$, for any set of 3D halfspaces or 2D disks (or objects in 2D with linear union complexity) [13]. Here, the weight $w(S)$ of a set S refers to the sum of the weights of the objects in S .

5.1 MWU Algorithm in the Weighted Case

Let X be the set of input points and S be the set of weighted input objects, where object i has weight w_i , with $n = |X| + |S|$. Let OPT be the weight of the minimum-weight set cover. We assume that a value $t \in [\text{OPT}, 2 \text{OPT}]$ is given; this assumption can be removed by a binary search for t .

We may delete objects with weights $> t$. We may automatically include all objects with weights $< \frac{1}{n}t$ in the solution, and delete them and all points covered by them, since the total weight of the solution increases by only $O(n \cdot \frac{1}{n}t) = O(t)$. Thus, all remaining objects have weights in $[\frac{1}{n}t, t]$. By rescaling, we may now assume that all objects have weights in $[1, n]$ and that $t = \Theta(n)$.

In the following, for a multiset \hat{S} where object i has multiplicity m_i , the weight of the multiset is defined as $w(\hat{S}) = \sum_i m_i w_i$.

We describe a simple variant of the basic MWU algorithm to solve the weighted set cover problem. (A more general, randomized MWU algorithm for geometric set cover was given recently by Chekuri, Har-Peled, and Quanrud [16], but our algorithm is simpler to describe and analyze.) The key innovation is to replace doubling with multiplication by a factor $1 + \frac{1}{w_i}$, where w_i is the weight of the concerned object i . (Note that multiplicities may now be non-integers.)

-
- 1: Guess a value $t \in [\text{OPT}, 2\text{OPT}]$.
 - 2: Define a multiset \hat{S} where each object i in S initially has multiplicity $m_i = 1$.
 - 3: **repeat**
 - 4: Find a point p which is ε -light in \hat{S} with $\varepsilon = \frac{1}{2t} \cdot \frac{w(\hat{S})}{|\hat{S}|}$.
 - 5: **for** each object i containing p **do** ▷ call lines 5–6 a “multiplicity-increasing step”
 - 6: Multiply its multiplicity m_i by $1 + \frac{1}{w_i}$.
 - 7: **until** all points are ε -heavy in \hat{S} .
 - 8: Return an ε -net of the multiset \hat{S} .
-

Since at the end all points are ε -heavy in \hat{S} , the returned subset is a valid set cover of X . For halfspaces in 3D or disks in 2D, its weight is $O(\frac{1}{\varepsilon} \cdot \frac{w(\hat{S})}{|\hat{S}|}) = O(\text{OPT})$.

We now prove that the algorithm terminates in $O(t \log n) = O(n \log n)$ multiplicity-increasing steps.

In each multiplicity-increasing step, $w(\hat{S})$ increases by

$$\sum_{\text{object } i \text{ containing } p} m_i \cdot \frac{1}{w_i} \cdot w_i = \sum_{\text{object } i \text{ containing } p} m_i \leq \frac{w(\hat{S})}{2t},$$

i.e., $w(\hat{S})$ increases by a factor of at most $1 + \frac{1}{2t}$. Initially, $w(\hat{S}) \leq n^2$. Thus, after z multiplicity-increasing steps, $w(\hat{S}) \leq n^2(1 + \frac{1}{2t})^z \leq n^2 e^{z/(2t)}$.

On the other hand, consider the optimal set cover T^* . Suppose that object i has its multiplicity increased z_i times. In each multiplicity-increasing step, at least one object in T^* has its multiplicity increased. So, after z multiplicity-increasing steps, $\sum_{i \in T^*} z_i \geq z$ and $\sum_{i \in T^*} w_i \leq t$. In particular, $z_i/w_i \geq z/t$ for some $i \in T^*$. Therefore, $w(\hat{S}) \geq (1 + \frac{1}{w_i})^{z_i} w_i \geq (1 + \frac{1}{w_i})^{z_i} \geq 2^{z_i/w_i} \geq 2^{z/t}$ (since $w_i \geq 1$). We conclude that $2^{z/t} \leq w(\hat{S}) \leq n^2 e^{z/(2t)}$, implying that $z = O(t \log n)$.

Similar to Agarwal and Pan’s first MWU algorithm, we can also divide the multiplicity-increasing steps into rounds, with each round performing up to t multiplicity-increasing steps. Within each round, the total weight $w(\hat{S})$ increases by at most $(1 + \frac{1}{2t})^t = O(1)$. Also if $|\hat{S}|$ increases by a constant factor, we immediately start a new round: because $|\hat{S}| \leq w(\hat{S})$ and $w(\hat{S})$ may be doubled at most $O(\log n)$ times, this case can happen at most $O(\log n)$ times. This ensures that if a point is checked to be ε -heavy at any moment during a round, it will remain $\Omega(\varepsilon)$ -heavy at the end of the round. There are only $O(\log n)$ rounds.

Additional ideas are needed to speed up implementation (in particular, our modified MWU algorithm with multiplicity-readjustment steps does not work as well now). First, we work with an approximation \tilde{m}_i to the multiplicity m_i of each object i . By rounding, we may assume all weights w_i are powers of 2. In the original algorithm, $m_i = (1 + \frac{1}{w_i})^{z_i}$, where z_i is the number of points $p \in Z$ that are contained in object i , and Z be the multiset consisting of all points p that have undergone multiplicity-increasing steps so far. Note that since the total multiplicity is $n^{O(1)}$, we have $z_i = O(w_i \log n)$. Let $Y^{(w_i)}$ be a random sample of Z where each point $p \in Z$ is included independently with probability $\frac{\log^2 n}{w_i}$ (if $w_i = O(\log^2 n)$, we can just set $Y^{(w_i)} = Z$). Let y_i be the number of points $p \in Y^{(w_i)}$ that are contained in object i . By the Chernoff bound, since $\frac{\log^2 n}{w_i} z_i = O(\log^3 n)$, we have $|y_i - \frac{\log^2 n}{w_i} z_i| \leq O(\log^2 n)$ with high probability. By letting $\tilde{m}_i = (1 + \frac{1}{w_i})^{y_i w_i / \log^2 n}$, it follows that \tilde{m}_i and m_i are within a factor of $O(1)$ of each other, with high probability, at all times, for all i . Thus, our earlier analysis still holds when working with \tilde{m}_i instead of m_i . Since $z_i = O(w_i \log n)$, we have $y_i = O(\log^3 n)$ with high probability. So, the total number of increments to all y_i and updates to all \tilde{m}_i is $O(n \log^3 n)$. In lines 5–6, we flip a biased coin to decide whether p should be placed in the sample $Y^{(2^j)}$ (with probability $\frac{\log^2 n}{2^j}$) for each j , and if so, we use halfspace range reporting in the dual to find all objects i of weight 2^j containing p , and increment y_i and update \tilde{m}_i . Over all $O(n \log n)$ executions of lines 5–6 and all $O(\log n)$ indices j , the cost of these halfspace range reporting queries is $O(n \log n \cdot \log n \cdot \log n)$ plus the output size. As the total output size for the queries is $O(n \log^3 n)$, the total cost is $O(n \log^3 n)$.

We also need to redesign a data structure for lightness testing subject to multiplicity updates: For each j , we maintain a subset $S^{(j)}$ containing all objects i with multiplicity at least 2^j , in a data structure to support approximate depth (without multiplicity). The depth of a point p in \hat{S} can be $O(1)$ -approximated by $\sum_j 2^j \cdot (\text{depth of } p \text{ in } S^{(j)})$. Each subset $S^{(j)}$ undergoes insertion only, and the logarithmic method can be applied to each $S^{(j)}$. Since $|\hat{S}| \leq w(\hat{S}) \leq n^{O(1)}$, there are $O(\log n)$ values of j . This slows down lightness testing by a logarithmic factor, and so in the case of 3D halfspaces, the overall time bound is $O(n \log^4 n \log \log n)$, excluding the ε -net construction time.

We can efficiently construct an ε -net of the desired weight for 3D halfspaces in $O(n \log n)$ randomized time, by using the quasi-random sampling technique of Varadarajan [31] and Chan et al. [13] in a more careful way. Due to lack of space, we defer the description to the full paper. We conclude:

► **Theorem 4.** *Given $O(n)$ points and $O(n)$ weighted halfspaces in \mathbb{R}^3 , we can find a subset of halfspaces covering all points, of total weight within $O(1)$ factor of the minimum, in $O(n \log^4 n \log \log n)$ expected time by a randomized Las Vegas algorithm.*

Remark. A remaining open problem is to find efficient deterministic algorithms for the weighted problem. Chan et al. [13] noted that the quasi-uniform sampling technique can be derandomized via the method of conditional probabilities, but the running time is high.

References

- 1 Peyman Afshani and Timothy M. Chan. On approximate range counting and depth. *Discrete & Computational Geometry*, 42(1):3–21, 2009. doi:10.1007/s00454-009-9177-z.
- 2 Peyman Afshani and Timothy M. Chan. Optimal halfspace range reporting in three dimensions. In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 180–186, 2009.

- 3 Peyman Afshani, Chris Hamilton, and Norbert Zeh. A general approach for cache-oblivious range reporting and approximate range counting. *Computational Geometry*, 43(8):700–712, 2010.
- 4 Pankaj K. Agarwal and Jiangwei Pan. Near-linear algorithms for geometric hitting sets and set covers. In *Proceedings of the 30th Symposium on Computational Geometry (SoCG)*, page 271, 2014.
- 5 Boris Aronov, Esther Ezra, and Micha Sharir. Small-size ε -nets for axis-parallel rectangles and boxes. *SIAM Journal on Computing*, 39(7):3248–3282, 2010. doi:10.1137/090762968.
- 6 Jon Louis Bentley and James B. Saxe. Decomposable searching problems I. static-to-dynamic transformation. *Journal of Algorithms*, 1(4):301–358, 1980.
- 7 Hervé Brönnimann and Michael T. Goodrich. Almost optimal set covers in finite VC-dimension. *Discrete & Computational Geometry*, 14(4):463–479, 1995.
- 8 Norbert Bus, Shashwat Garg, Nabil H. Mustafa, and Saurabh Ray. Limits of local search: Quality and efficiency. *Discrete & Computational Geometry*, 57(3):607–624, 2017. doi:10.1007/s00454-016-9819-x.
- 9 Norbert Bus, Nabil H. Mustafa, and Saurabh Ray. Practical and efficient algorithms for the geometric hitting set problem. *Discrete Applied Mathematics*, 240:25–32, 2018.
- 10 Timothy M. Chan. Random sampling, halfspace range reporting, and construction of ($\leq k$)-levels in three dimensions. *SIAM Journal on Computing*, 30(2):561–575, 2000. doi:10.1137/S0097539798349188.
- 11 Timothy M. Chan. Polynomial-time approximation schemes for packing and piercing fat objects. *Journal of Algorithms*, 46(2):178–189, 2003. doi:10.1016/S0196-6774(02)00294-8.
- 12 Timothy M. Chan. Dynamic geometric data structures via shallow cuttings. In *Proceedings of the 35th Symposium on Computational Geometry (SoCG)*, pages 24:1–24:13, 2019.
- 13 Timothy M. Chan, Elyot Grant, Jochen Könemann, and Malcolm Sharpe. Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1576–1585, 2012.
- 14 Timothy M. Chan, Kasper Green Larsen, and Mihai Pătraşcu. Orthogonal range searching on the RAM, revisited. In *Proceedings of the 27th Symposium on Computational Geometry (SoCG)*, pages 1–10, 2011.
- 15 Timothy M. Chan and Konstantinos Tsakalidis. Optimal deterministic algorithms for 2-d and 3-d shallow cuttings. *Discrete & Computational Geometry*, 56(4):866–881, 2016.
- 16 Chandra Chekuri, Sarel Har-Peled, and Kent Quanrud. Fast LP solving and approximation algorithms for geometric packing and covering problems. In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1019–1038, 2020.
- 17 Chandra Chekuri and Kent Quanrud. Randomized MWU for positive LPs. In *Proceedings of the 29th annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 358–377, 2018.
- 18 Kenneth L. Clarkson. Algorithms for polytope covering and approximation. In *Workshop on Algorithms and Data Structures*, pages 246–252, 1993.
- 19 Kenneth L. Clarkson and Kasturi Varadarajan. Improved approximation algorithms for geometric set cover. *Discrete & Computational Geometry*, 37(1):43–58, 2007.
- 20 Alon Efrat, Matthew J. Katz, Frank Nielsen, and Micha Sharir. Dynamic data structures for fat objects and their applications. *Computational Geometry*, 15(4):215–227, 2000. doi:10.1016/S0925-7721(99)00059-0.
- 21 Dorit S. Hochbaum and Wolfgang Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *Journal of the ACM (JACM)*, 32(1):130–136, 1985. doi:10.1145/2455.214106.
- 22 Christos Koufogiannakis and Neal E. Young. A nearly linear-time PTAS for explicit fractional packing and covering linear programs. *Algorithmica*, 70(4):648–674, 2014. doi:10.1007/s00453-013-9771-6.

- 23 Jiří Matoušek. Reporting points in halfspaces. *Computational Geometry*, 2(3):169–186, 1992.
- 24 Jiří Matoušek, Raimund Seidel, and Emo Welzl. How to net a lot with little: Small ϵ -nets for disks and halfspaces. In *Proceedings of the 6th Symposium on Computational Geometry (SoCG)*, pages 16–22, 1990.
- 25 Nabil H. Mustafa. Computing optimal epsilon-nets is as easy as finding an unhit set. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 87:1–87:12, 2019. doi:10.4230/LIPIcs.ICALP.2019.87.
- 26 Nabil H. Mustafa, Rajiv Raman, and Saurabh Ray. Quasi-polynomial time approximation scheme for weighted geometric set cover on pseudodisks and halfspaces. *SIAM J. Comput.*, 44(6):1650–1669, 2015. doi:10.1137/14099317X.
- 27 Nabil H. Mustafa and Saurabh Ray. PTAS for geometric hitting set problems via local search. In *Proceedings of the 25th Symposium on Computational Geometry (SoCG)*, pages 17–22, 2009.
- 28 János Pach and Gábor Tardos. Tight lower bounds for the size of epsilon-nets. In *Proceedings of the 27th Symposium on Computational Geometry (SoCG)*, pages 458–463, 2011. doi:10.1145/1998196.1998271.
- 29 Evangelia Pyrga and Saurabh Ray. New existence proofs for ϵ -nets. In *Proceedings of the 24th Symposium on Computational Geometry (SoCG)*, pages 199–207, 2008. doi:10.1145/1377676.1377708.
- 30 Edgar A. Ramos. On range reporting, ray shooting and k -level construction. In *Proceedings of the 15th Symposium on Computational Geometry (SoCG)*, pages 390–399, 1999.
- 31 Kasturi Varadarajan. Weighted geometric set cover via quasi-uniform sampling. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC)*, pages 641–648, 2010.
- 32 Kasturi R. Varadarajan. Epsilon nets and union complexity. In *Proceedings of the 25th Symposium on Computational Geometry (SoCG)*, pages 11–16, 2009. doi:10.1145/1542362.1542366.
- 33 Neal E. Young. Sequential and parallel algorithms for mixed packing and covering. In *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 538–546, 2001. doi:10.1109/SFCS.2001.959930.