

# More Dynamic Data Structures for Geometric Set Cover with Sublinear Update Time

Timothy M. Chan ✉ 

Department of Computer Science, University of Illinois at Urbana-Champaign, IL, USA

Qizheng He ✉

Department of Computer Science, University of Illinois at Urbana-Champaign, IL, USA

---

## Abstract

We study geometric set cover problems in dynamic settings, allowing insertions and deletions of points and objects. We present the first dynamic data structure that can maintain an  $O(1)$ -approximation in sublinear update time for set cover for axis-aligned squares in 2D. More precisely, we obtain randomized update time  $O(n^{2/3+\delta})$  for an arbitrarily small constant  $\delta > 0$ . Previously, a dynamic geometric set cover data structure with sublinear update time was known only for unit squares by Agarwal, Chang, Suri, Xiao, and Xue [SoCG 2020]. If only an approximate size of the solution is needed, then we can also obtain sublinear amortized update time for disks in 2D and halfspaces in 3D. As a byproduct, our techniques for dynamic set cover also yield an optimal randomized  $O(n \log n)$ -time algorithm for static set cover for 2D disks and 3D halfspaces, improving our earlier  $O(n \log n (\log \log n)^{O(1)})$  result [SoCG 2020].

**2012 ACM Subject Classification** Theory of computation → Computational geometry

**Keywords and phrases** Geometric set cover, approximation algorithms, dynamic data structures, sublinear algorithms, random sampling

**Digital Object Identifier** 10.4230/LIPIcs.SoCG.2021.25

**Related Version** *Full Version*: <https://arxiv.org/abs/2103.07857>

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

## 1 Introduction

Approximation algorithms for NP-hard problems and dynamic data structures are two of the major themes studied by the algorithms community. Recently, problems at the intersection of these two threads have gained much attention, and researchers in computational geometry have also started to systematically explore such problems. For example, at SoCG last year, two papers appeared, one on dynamic geometric set cover by Agarwal et al. [2], and another on dynamic geometric independent set by Henzinger, Neumann, and Wiese [19]. In this paper, we continue the study by Agarwal et al. [2] and investigate dynamic data structures for approximating the minimum set cover in natural geometric instances.

**Static geometric set cover.** In the static (unweighted) *geometric set cover* problem, we are given a set  $X$  of  $O(n)$  points and a set  $S$  of  $O(n)$  geometric objects, and we want to find the smallest subset of objects in  $S$  that covers all points of  $X$ . The problem is fundamental and has many applications. Let  $\text{OPT}$  denote the value (i.e., cardinality) of the optimal solution. As the problem is NP-hard for many classes of geometric objects, we are interested in efficient approximation algorithms.

Many classes of objects, such as squares and disks in 2D, objects with linear union complexity, and halfspaces in 3D, admit polynomial-time  $O(1)$ -approximation algorithms (i.e., computing a solution of size  $O(1) \cdot \text{OPT}$ ), by using  $\varepsilon$ -nets and LP rounding or the



© Timothy M. Chan and Qizheng He;

licensed under Creative Commons License CC-BY 4.0


37th International Symposium on Computational Geometry (SoCG 2021).

Editors: Kevin Buchin and Éric Colin de Verdière; Article No. 25; pp. 25:1–25:14

Leibniz International Proceedings in Informatics



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



multiplicative weight update (MWU) method [6, 14, 15]. Some classes of objects, such as disks in 2D and halfspaces in 3D, even have PTASs [26] or quasi-PTASs [25], though with large polynomial running time.

Agarwal and Pan [4] worked towards finding *faster* approximation algorithms that run in near linear time. They gave randomized  $O(n \log^4 n)$ -time algorithms (based on MWU) for computing  $O(1)$ -approximations, for example, for disks in 2D and halfspaces in 3D. At last year's SoCG [11], we described further improvements to the running time for 2D disks and 3D halfspaces, including a deterministic  $O(n \log^3 n \log \log n)$ -time algorithm and a randomized  $O(n \log n (\log \log n)^{O(1)})$ -time algorithm.

**Dynamic geometric set cover.** It is natural to consider the *dynamic* setting of the geometric set cover problem. Here, we want to support insertions and deletions of points in  $X$  as well as insertions and deletions of objects in  $S$ , while maintaining an approximate solution. Note that the solution may have linear size in the worst case. In the simplest version of the problem, we may just want to output the value of the solution. More strongly, we may want some representation of the solution itself, so that afterwards, the objects in the solution can be reported in constant time per element when needed.

Agarwal et al. [2] gave a number of results on dynamic geometric set cover. They showed that for intervals in 1D, a  $(1 + \varepsilon)$ -approximation can be maintained in  $O((1/\varepsilon)n^\delta)$  time per insertions and deletions of points and intervals. (Throughout the paper,  $\delta > 0$  denotes an arbitrarily small constant.) In 2D, they had only one main result: a fully dynamic  $O(1)$ -approximation algorithm for unit axis-aligned squares, with  $O(n^{1/2+\delta})$  update time.

**New results.** We present several new results on dynamic geometric set cover. The first is a fully dynamic, randomized,  $O(1)$ -approximation algorithm for the more general case of arbitrary axis-aligned squares in 2D, with  $O(n^{2/3+\delta})$  amortized update time. Though our time bound is a little worse than Agarwal et al.'s for the unit square case, the arbitrary square case is more challenging. (The unit square case reduces to case of dominance ranges, i.e., quadrants, via a standard grid approach; since the union of such ranges forms a “staircase” sequence of vertices, the problem is in some sense “1.5-dimensional”. In contrast, the arbitrary square problem requires truly “2-dimensional” ideas.)

We then consider the case of halfspaces in 3D. This case is fundamental, as set cover for 2D disks reduces to set cover for 3D halfspaces by the standard lifting transformation [17]. Also, by duality, hitting set for 3D halfspaces is equivalent to set cover for 3D halfspaces, and hitting set for 2D disks reduces to set cover for 3D halfspaces as well.

For 3D halfspaces, we obtain a fully dynamic, randomized algorithm with  $O(n^{12/13+\delta}) \leq O(n^{0.924})$  amortized update time. Our result here is slightly weaker: it only finds the value of an  $O(1)$ -approximate solution (which could be good enough in some applications). If a solution itself is required, we can still get sublinear update time as long as OPT is sublinear (below  $n^{1-\delta}$ ). This assumption seems reasonable, since sublinear reporting time is not possible otherwise. (However, we currently do not know how to obtain a stronger time bound of the form  $O(n^\alpha + \text{OPT})$  with  $\alpha < 1$  to report a solution for 3D halfspaces.)

**Remarks.** Our results are randomized in the Monte Carlo sense: the computed solution may not always be correct, but it is correct with high probability (w.h.p.), i.e., probability at least  $1 - 1/n^c$  for an arbitrarily large constant  $c$ . The error probability bounds hold even when the user has knowledge of the random choices made by the algorithms. (We do not assume an “oblivious adversary”; in fact, the algorithms make a new set of random choices each time it computes a solution, and the data structures themselves are not randomized.)

The update times are better than stated in many cases, notably, when OPT is small or when OPT is large. More precise OPT-sensitive bounds are given in lemmas and theorems throughout the paper. (The  $O(n^{2/3+\delta})$  and  $O(n^{12/13+\delta})$  bounds are obtained by “balancing”.)

We have assumed that we are required to compute a solution after every update. In some (but not all) cases, the update cost is smaller than the cost of computing a solution (this may be useful if we are executing a batch of updates).

**Techniques.** Our algorithms are obtained by handling two cases differently: when OPT is small and when OPT is large. Intuitively, the small OPT case is easier since we are generating fewer objects, but the large OPT case also seems potentially easier since we can tolerate a larger additive error when targeting an  $O(1)$ -factor approximation – so, we are in a “win-win” situation. For our algorithms for 3D halfspaces, we even find it necessary to handle an intermediate case when OPT is medium (aiming for sublinear time for sublinear OPT).

Our algorithms for the small OPT case are based on the previous static MWU algorithms [6, 4, 11]. The adaptation of these static algorithms is not straightforward, and requires using various known techniques in new ways ( $\leq k$ -levels in arrangements, for our 2D square algorithm in Section 3.1, and “augmented” partition trees, for our 3D halfspace algorithm in Section 4.1). The medium case for 3D halfspaces (in Section 4.2) is technically even more challenging (where we use “shallow” partition trees and other ideas).

Our algorithm for squares in the large OPT case (in Section 3.2) is different (not based on MWU), and interestingly uses quadtrees in a non-obvious way. For 3D halfspaces, our algorithm in the large OPT case (in Section 4.3) can compute only the value of an approximate solution, and is based on random sampling. However, the obvious way to use a random sample (just solving the problem on a random subset of points and objects) does not work. We use sampling in a nontrivial way, combining geometric cuttings with planar graph separators.

In some of our algorithms, notably the small OPT algorithms for squares and 3D halfspaces (in Sections 3.1 and 4.1) and the large OPT algorithm for 3D halfspaces (in Section 4.3), the data structure part is “minimal”: we just assume that points and objects are stored separately in standard range searching data structures. We describe sublinear-time algorithms to compute a solution from scratch, using range searching as oracles. Dynamization becomes trivial, since range searching data structures typically are already known to support insertions and deletions. (It also potentially enables other operations like merging sets, or solving the set cover problem for range-restricted subsets of points and subsets of objects.)

The topic of *sublinear-time algorithms* has received considerable attention in the algorithms community, due to applications to big data (where we want to solve problems without examining the entire input). A similar model of sublinear-time algorithms where the input is augmented with range searching data structures was proposed by Czumaj et al. [16], who presented results on approximating the weight of the Euclidean minimum spanning tree in any constant dimension under this model.

**Application to static geometric set cover.** Although we did not intend to revisit the static problem, our techniques can lead a randomized  $O(1)$ -approximation algorithm for set cover for 3D halfspaces running in  $O(n \log n)$  time, which completely eliminates the extra  $\log \log n$  factors in our previous result from SoCG’20 [11] and is optimal (in comparison-based models)! This bonus result is interesting in its own right, but due to lack of space, we defer the description to the full paper.

## 2 Review of an MWU algorithm

We begin by briefly reviewing a known static approximation algorithm for geometric set cover, based on the *multiplicative weight updates (MWU)* method. Some of our dynamic algorithms will be built upon this algorithm.

Specifically, we consider the following randomized algorithm from our previous SoCG'20 paper [11], which is a variant of a standard algorithm by Brönnimann and Goodrich [6] or Clarkson [14] (see also Agarwal and Pan [4]). Below,  $c_0$  is a sufficiently large constant. The *depth* of a point  $p$  in a set  $S$  of objects is the number of objects of  $S$  containing  $p$ . A subset of objects  $T \subseteq S$  is called an  $\varepsilon$ -*net* of  $S$  if all points with depth  $\geq \varepsilon|S|$  in  $S$  are covered by  $T$ . The size  $|\hat{S}|$  of a multiset  $\hat{S}$  refers to the sum of the multiplicities of its elements  $\sum_{i \in S} m_i$ .

■ **Algorithm 1** MWU for set cover.

---

```

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: loop ▷ call this the start of a new round
4:   Fix  $\rho := \frac{c_0 t \log n}{|\hat{S}|}$  and take a random sample  $R$  of  $\hat{S}$  with sampling probability  $\rho$ .
5:   while there exists a point  $p \in X$  with depth in  $R$  at most  $\frac{c_0}{2} \log n$  do
6:     for each object  $i$  containing  $p$  do ▷ call lines 6–8 a multiplicity-doubling step
7:       Double its multiplicity  $m_i$ , i.e., insert  $m_i$  new copies of object  $i$  into  $\hat{S}$ .
8:       For each copy, independently decide to insert it into  $R$  with probability  $\rho$ .
9:     if the number of multiplicity-doubling steps in this round exceeds  $t$  then
10:      Go to line 3 and start a new round.
11:  Terminate and return a  $\frac{1}{8t}$ -net of  $R$ .
```

---

In the standard version of MWU, the “lightness” condition in line 5 was whether the depth of  $p$  in  $\hat{S}$  is at most  $\frac{|\hat{S}|}{2t}$ . The main difference in the above randomized version is that lightness is tested with respect to the sample  $R$ , which is computationally easier to work with –  $R$  has size  $O(t \log n)$  with high probability (w.h.p.). Justification of this randomized variant follows from a Chernoff bound, and was shown in [11, Section 4.2].

It is known that the algorithm terminates in  $O(\log \frac{n}{t})$  rounds and  $O(t \log \frac{n}{t})$  multiplicity-doubling steps [6, 4, 11]. Furthermore,  $|\hat{S}|$  increases by at most a factor of 2 in each round; in particular,  $|\hat{S}|$  is bounded by  $n^{O(1)}$  at the end of  $O(\log \frac{n}{t})$  rounds.

In the standard version of MWU, line 11 returns a  $\Theta(\frac{1}{t})$ -net of the multiset  $\hat{S}$ , but a  $\Theta(\frac{1}{t})$ -net of  $R$  works just as well: in the end, the depth in  $R$  of all points of  $X$  is at least  $\frac{c_0}{2} \log n > \frac{|R|}{8t}$  w.h.p., and so  $X$  will be covered by the net. For objects that are axis-aligned squares in 2D, disks in 2D, or halfspaces in 3D,  $\varepsilon$ -nets of size  $O(\frac{1}{\varepsilon})$  exist, and thus the above algorithm yields a set cover of size  $O(t)$ , i.e., a constant-factor approximation. By known algorithms (e.g., [12]), the net for  $R$  in line 11 can be constructed in  $\tilde{O}(|R|) = \tilde{O}(t)$  time.

Several modifications have been explored in previous work. For example, in the first algorithm of Agarwal and Pan [4], each round examines all points of  $X$  in a fixed order and test for lightness of the points one by one (based on the observation that a point found to have large depth will still have large depth by the end of the round). In our previous paper [11], we have also added a step at the beginning of each round, where the multiplicities are rescaled and rounded, so as to keep  $|\hat{S}|$  bounded by  $O(n)$ . These modifications led to a number of different static implementations running in  $\tilde{O}(n)$  time.

### 3 Axis-aligned squares

Our first result for dynamic set cover is for axis-aligned squares. Previously, a sublinear time algorithm for dynamic set cover was known only for unit squares. Our method will be divided into two cases: when OPT is small and when OPT is large. Let  $t$  be a guess on OPT. We will run our algorithm for each possible  $t = 2^i$  in parallel. (When the guess is wrong, our algorithm will be able to tell whether  $t$  is approximately smaller or larger than OPT w.h.p.) The update time will increase only by a factor of  $O(\log n)$ .

#### 3.1 Algorithm for small OPT

Our algorithm for the small OPT case will be based on the randomized MWU algorithm described in Section 2. The key is to realize that this algorithm can actually be implemented to run in *sublinear* time, assuming that the points and the objects have been preprocessed in standard range searching data structures. Since these structures are dynamizable, we can just re-run the MWU algorithm from scratch after every update.

##### 3.1.1 Data structures

Our data structures are simple. We store the point set  $X$  in the standard 2D *range tree* [17]. For each square  $s$  with center  $(x, y)$  and side length  $2z$ , map  $s$  to a point  $s^\uparrow = (x - z, x + z, y - z, y + z)$  in 4D. We also store the lifted point set  $S^\uparrow = \{s^\uparrow : s \in S\}$  in a 4D range tree. Range trees support insertions and deletions in  $X$  and  $S$  in polylogarithmic time.

##### 3.1.2 Computing a solution

We now show how to compute an  $O(1)$ -approximate solution in sublinear time when OPT is small by running Algorithm 1 using the above data structures. At first glance, linear time seems unavoidable: (i) the obvious way to find low-depth points in line 5 is to scan through all points of  $X$  in each round (as was done in previous algorithms [4, 11]), and (ii) explicitly maintaining the multiplicities of all points would also require linear time.

To overcome these obstacles, we observe that (i) we can use data structures to find the next low-depth point  $p$  without testing each point one by one (recall that there are only  $\tilde{O}(t)$  multiplicity-doubling steps), and (ii) multiplicities do not need to be maintained explicitly, so long as in line 8 we can generate a multiplicity-weighted random sample among the objects containing a given point  $p$  efficiently (recall that the sample  $R$  has only  $\tilde{O}(t)$  size). The subproblem in (ii) is a *weighted range sampling* problem.

**Finding a low-depth point.** Let  $b := \frac{c_0}{2} \log n$ . Each time we want to find a low-depth point in line 5, we compute (from scratch)  $\mathcal{L}_{\leq b}(R)$ , the  $(\leq b)$ -level of  $R$ , i.e., the collection of all cells in the arrangement of the squares of depth at most  $b$ . It is known [27] that  $\mathcal{L}_{\leq b}(R)$  has  $O(|R|b)$  cells and can be constructed in  $\tilde{O}(|R|b)$  time, which is  $\tilde{O}(t)$  (since  $|R| = \tilde{O}(t)$  and  $b = \tilde{O}(1)$ ). To find a point  $p$  of  $X$  that has depth in  $R$  at most  $b$ , we simply examine each cell of  $\mathcal{L}_{\leq b}(R)$ , and perform an orthogonal range query to test if the cell contains a point of  $X$ . All this takes  $\tilde{O}(t)$  time.

As there are  $\tilde{O}(t)$  multiplicity-doubling steps, the total cost is  $\tilde{O}(t^2)$ .

**Weighted range sampling.** For each square  $s$  with center  $(x, y)$  and side length  $2z$ , define its *dual* point  $s^*$  to be  $(x, y, z)$  in 3D. For each point  $p = (p_x, p_y)$ , define its *dual* region  $p^*$  to be  $\{(x, y, z) : z \geq \max\{|x - p_x|, |y - p_y|\}\}$  in 3D. Then a point  $p$  is in the square  $s$  iff the point  $s^*$  is in the region  $p^*$ .

Let  $Q$  be the set of all points  $p$  for which we have performed multiplicity-doubling steps thus far. Note that  $|Q| = \tilde{O}(t)$ . Let  $b' = c'_0 \log n$  for a sufficiently large constant  $c'_0$ . Each time we perform a multiplicity-doubling step, we compute (from scratch)  $\mathcal{L}_{\leq b'}(Q^*)$ , the  $(\leq b')$ -level of the dual regions  $Q^* = \{p^* : p \in Q\}$ . This structure corresponds to planar order- $(\leq b')$   $L_\infty$  Voronoi diagrams. By known results [27, 7, 21],  $\mathcal{L}_{\leq b'}(Q^*)$  has  $O(|Q|(b')^2)$  cells and can be constructed in  $\tilde{O}(|Q|(b')^2)$  time, which is  $\tilde{O}(t)$  (since  $|Q| = \tilde{O}(t)$  and  $b' = \tilde{O}(1)$ ). The multiplicity of a square  $s \in S$  is equal to  $2^{\text{depth of } s^* \text{ in } Q^*}$  (since each  $p^*$  in  $Q^*$  containing  $s^*$  doubles the multiplicity of  $s$ ). In particular, since multiplicities are bounded by  $n^{O(1)}$ , the depth (i.e., level) of  $s^*$  must be logarithmically bounded. So, each  $s^*$  is covered by  $\mathcal{L}_{\leq b'}(Q^*)$ , and the multiplicity of  $s$  is determined by which cell of  $\mathcal{L}_{\leq b'}(Q^*)$  the point  $s^*$  is in.

To generate a multiplicity-weighted sample of the squares containing  $p$  for line 8, after  $p$  has been inserted to  $Q$ , we examine all cells of  $\mathcal{L}_{\leq b'}(Q^*)$  contained in  $p^*$ . For each such cell  $\gamma$ , we identify the squares  $s \in S$  for which  $s^* \in \gamma$ ; this reduces to an orthogonal range query in  $S^\uparrow$ , and the answer can be expressed as a disjoint union of  $\tilde{O}(1)$  canonical subsets. Knowing the sizes and multiplicities of these canonical subsets for all such  $\tilde{O}(t)$  cells, we can then generate the weighted sample in time  $\tilde{O}(t)$  plus the size of the sample.

Hence, the total cost of all  $\tilde{O}(t)$  multiplicity-doubling steps is  $\tilde{O}(t^2)$ .

In addition, we need to generate a new weighted sample  $R$  (with a new sampling probability  $\rho$ ) in line 4 at the beginning of each round; this can be done similar to above, in  $\tilde{O}(t)$  time plus the size of the sample ( $\tilde{O}(t)$ ), for each of the  $O(\log n)$  rounds. As mentioned, the final net computation in line 11 takes  $\tilde{O}(t)$  time. We have thus obtained:

► **Lemma 1.** *There exists a data structure for the dynamic set cover problem for  $O(n)$  axis-aligned squares and  $O(n)$  points in 2D that supports insertions and deletions in  $\tilde{O}(1)$  time and can find an  $O(1)$ -approximate solution w.h.p. to the set cover problem in  $\tilde{O}(\text{OPT}^2)$  time.*

### 3.2 Algorithm for large OPT

To complement our solution for the small OPT case, we now show that the problem also gets easier when OPT is large, mainly because we can afford a large additive error. We describe a different, self-contained algorithm for this case (not based on modifying MWU), interestingly by using quadtrees in a novel way.


To allow for both multiplicative and additive error, we use the term  $(\alpha, \beta)$ -approximation to refer to a solution with cost at most  $\alpha \text{OPT} + \beta$ .

For simplicity, we assume that all coordinates are integers bounded by  $U = \text{poly}(n)$ . At the end, we will comment on how to remove this assumption.

In the standard *quadtree*, we start with a bounding square cell and recursively divide a square cell into four square subcells. We define the size of a cell  $\Gamma$  to be the number of vertices in  $\Gamma$  among the squares of  $S$ , plus the number of points of  $X$  in  $\Gamma$ . We stop subdividing when a leaf cell has size at most  $b$ , where  $b$  is a parameter to be set later. This yields a subdivision into  $O(\frac{n}{b})$  cells per level, and  $O(\frac{n}{b} \log U)$  cells in total. The quadtree decomposition can be easily made dynamic under insertions and deletions of points and squares.

For each leaf cell  $\Gamma$ , a square in  $S$  intersecting  $\Gamma$  is called *short* if at least one of its vertices is in the cell, and *long* otherwise, as shown in Fig. 1. Note that a long square can have at most one side crossing the cell, because the quadtree cell  $\Gamma$  is also a square. The union of the long squares within the cell is defined by at most 4 long squares – call these the *maximal* long squares. (If there is a square containing  $\Gamma$ , we can designate one such square as the maximal long square.) For each leaf cell  $\Gamma$ , it suffices to approximate the optimal set cover for the input points in  $X \cap \Gamma$  using only the short squares plus the at most 4 maximal long

squares in  $\Gamma$ . By charging each square in the optimal solution to the cells containing its 4 vertices, we see that the sum of the sizes of the optimal covers in the leaf cells is at most  $4 \text{OPT} + O(\frac{n}{b} \log U)$ , which is indeed an  $O(1)$ -approximation if we choose  $b \geq \frac{n \log n}{\text{OPT}}$ .



■ **Figure 1** Short square (left) and long square (right). The quadtree cell is shaded.

Note that the complement of the union of the at most 4 maximal long squares in a cell  $\Gamma$  is a rectangle  $r_\Gamma$ . We will store the short squares in the cell  $\Gamma$  in a data structure  $\mathcal{S}_\Gamma$  to answer the following type of query:

Given any query rectangle  $r$ , compute an  $O(1)$ -approximation to the optimal set cover for the points in  $X \cap r$  using only the short squares.

Assuming the availability of such data structures  $\mathcal{S}_\Gamma$ , we can solve the dynamic set cover problem as follows:

- An insertion/deletion of a square  $s$  in  $S$  requires updating 4 of these data structures  $\mathcal{S}_\Gamma$  for the leaf cells  $\Gamma$  containing the 4 vertices of  $s$ , and also updating the maximal long squares for all leaf cells in  $\tilde{O}(\frac{n}{b})$  time.
- An insertion/deletion of a point  $p$  in  $X$  requires updating one data structure  $\mathcal{S}_\Gamma$  for the leaf cell  $\Gamma$  containing  $p$ .
- Whenever we want to compute a set cover solution, we examine all  $\tilde{O}(\frac{n}{b})$  cells  $\Gamma$  and query the data structure  $\mathcal{S}_\Gamma$  for  $r_\Gamma$ , and return the union of the answers.

**First implementation of  $\mathcal{S}_\Gamma$ .** A simple way to implement the data structure  $\mathcal{S}_\Gamma$  is as follows: in an update, we just recompute an approximate solution from scratch for every possible query rectangle in  $\Gamma$ . Since there are only  $O(b^4)$  combinatorially different query rectangles, and static approximate set cover on  $O(b)$  squares and points takes  $\tilde{O}(b)$  time [4, 11], the update time is  $\tilde{O}(b^5)$ , and the query time is trivially  $O(1)$ .

As a result, the cost of insertion/deletion in the overall method is  $\tilde{O}(b^5 + \frac{n}{b})$ .

**Improved implementation for  $\mathcal{S}_\Gamma$ .** We further improve the update time for the data structure  $\mathcal{S}_\Gamma$ . Instead of recomputing solutions for all  $O(b^4)$  rectangles, the idea is to recompute solutions for a smaller number of “canonical rectangles”. More precisely, by using a 2D range tree [3, 17] for the  $O(b)$  points in  $X \cap \Gamma$ , with branching factor  $a$  in each dimension, we can form a set of canonical rectangles with total size  $O(a^{O(1)}b(\log_a b)^2)$ , such that every query rectangle can be decomposed into  $O((\log_a b)^2)$  canonical rectangles, ignoring portions that are empty of points. We set  $a := b^\delta$  for an arbitrarily small constant  $\delta > 0$ .

For each canonical rectangle  $r$  with size  $b_i$ , there are at most  $O(b_i)$  maximal long squares with respect to  $r$ : the union of the long squares that cut across  $r$  horizontally have at most two edges between any two consecutive points/vertices; and a similar statement holds for the long squares that cut across  $r$  vertically. These maximal long squares can be found in  $\tilde{O}(b_i)$  time by standard orthogonal range searching. We can thus approximate the optimal set cover for the points in the canonical rectangle  $r$ , using the  $O(b_i)$  short squares and maximal long squares with respect to  $r$ , in  $\tilde{O}(b_i)$  time by known static set cover algorithms [4, 11]. The total time over all canonical rectangles is  $\tilde{O}(a^{O(1)}b(\log_a b)^2) = \tilde{O}(b^{1+O(\delta)})$ .

Given a query rectangle, we can decompose it into  $O((\log_a b)^2) = O(1)$  canonical rectangles and return the union of the optimal solutions in the canonical rectangles, which is an  $O(1)$ -approximation (more precisely, an  $O(\delta^{-2})$ -approximation).

As a result, the cost of insertion/deletion in the overall method is  $\tilde{O}(b^{1+O(\delta)} + \frac{n}{b})$ .

**Removing the dependency on  $U$ .** When  $U$  may be large, we can reduce the tree depth from  $O(\log U)$  to  $O(\log n)$  by replacing the quadtree with the *BBD tree* of Arya et al. [5]. Each cell in the BBD tree is the set difference of two quadtree squares, one contained in the other. The size of each child cell is at most a fraction of the size of the parent cell. As before, we stop subdividing when a leaf cell has size at most  $b$ . Since any such leaf cell has size  $\Theta(b)$  now, the number of leaf cells is  $O(\frac{n}{b})$ . The BBD tree can be maintained dynamically in polylogarithmic time (for example, by periodically rebuilding when subtrees become unbalanced). Since a leaf cell  $\Gamma$  is the difference of two quadtree squares, it is not difficult to see that the number of maximal long squares in  $\Gamma$  remains  $O(1)$ . So, our previous analysis remains valid.

► **Lemma 2.** *Given a parameter  $b$ , there exists a data structure for the dynamic set cover problem for  $O(n)$  axis-aligned squares and  $O(n)$  points in 2D that maintains an  $(O(1), O(\frac{n}{b}))$ -approximate solution with  $\tilde{O}(b^{1+O(\delta)} + \frac{n}{b})$  insertion and deletion time.*

**Combining the algorithms.** When  $\text{OPT} \leq n^{1/3}$ , we use the algorithm for small OPT; the running time is  $\tilde{O}(\text{OPT}^2) \leq \tilde{O}(n^{2/3})$ . When  $\text{OPT} > n^{1/3}$ , we use the algorithm for large OPT with  $b = n^{2/3}$ , so that an  $(O(1), O(\frac{n}{b}))$ -approximation is indeed an  $O(1)$ -approximation; the running time is  $\tilde{O}(b^{1+O(\delta)} + \frac{n}{b}) = O(n^{2/3+O(\delta)})$ .

► **Theorem 3.** *There exists a data structure for the dynamic set cover problem for  $O(n)$  axis-aligned squares and  $O(n)$  points in 2D that maintains an  $O(1)$ -approximate solution w.h.p. with  $O(n^{2/3+\delta})$  insertion and deletion time for any constant  $\delta > 0$ .*

The case of fat rectangles can be reduced to squares, since such rectangles can be replaced by  $O(1)$  squares (increasing the approximation factor by only  $O(1)$ ). Our approach can be modified to work more generally for homothets of a fixed fat convex polygon with a constant number of vertices.

## 4 Halfspaces in 3D

In this section, we study dynamic geometric set cover for the more challenging case of 3D halfspaces. Using the standard lifting transformation [17], we can transform 2D disks to 3D upper halfspaces. For simplicity, we assume that all halfspaces are upper halfspaces; we discuss how to modify our algorithms when there are both upper and lower halfspaces in the full paper. Our method will be divided into three cases: small, medium, and large OPT.

### 4.1 Algorithm for small OPT

Similar to the small OPT algorithm for axis-aligned squares in Section 3.1, we describe a small OPT algorithm for halfspaces based on the randomized MWU algorithm in Section 2. Although our earlier approach using levels in arrangements could be generalized, we describe a better approach based on augmenting partition trees with counters.



### 4.1.1 Data structures

We store the 3D point set  $X$  in Matoušek's *partition tree* [22]: The tree has height  $O(\log n)$  and degree  $r$  for a sufficiently large constant  $r$ . Each node  $v$  stores a simplicial cell  $\Gamma_v$  and a "canonical subset"  $X_v \subset \Gamma_v$ , where  $X_v = X$  at the root  $v$ , and  $\Gamma_v$  is contained in  $\Gamma_{\text{parent}(v)}$ ,  $X_v$  is the disjoint union of  $X_{v'}$  over all children  $v'$  of  $v$ , and  $X_v$  has constant size at each leaf  $v$ . Furthermore, any halfspace crosses  $O(n^{2/3+\delta})$  cells of the tree for any arbitrarily small constant  $\delta > 0$  (depending on  $r$ ). Here a halfspace  $h$  *crosses* a cell  $\Gamma$  iff the boundary of  $h$  intersects  $\Gamma$ .

For each upper halfspace  $h$ , let  $h^*$  denote its dual point; for each point  $p$ , let  $p^*$  denote its dual upper halfspace. (Duality [17] is defined so that  $p$  is in  $h$  iff  $h^*$  is in  $p^*$ .)

We also store the 3D dual point set  $S^* = \{h^* : h \in S\}$  in Matoušek's partition tree. Each node  $v$  stores a cell  $\Gamma_v$  and a canonical subset  $S_v^* \subset \Gamma_v$  like above.

Matoušek's partition trees can be built in  $\tilde{O}(n)$  time and support insertions and deletions in  $X$  and  $S$  in polylogarithmic time. (In the static case, there are slightly improved partition trees reducing the  $n^\delta$  factor in the crossing number bound [22, 24, 10], but these will not be important to us.)

### 4.1.2 Computing a solution

We now show how to compute an  $O(1)$ -approximate solution in sublinear time when OPT is small by running Algorithm 1 using the above data structures. As in Section 3.1, the main subproblems are (i) finding a low-depth point with respect to  $R$ , and (ii) weighted range sampling, where the weights are the multiplicities (which are not explicitly stored).

**Finding a low-depth point.** We maintain two values at each node  $v$  of the partition tree of  $X$ :

- $c_v$  is the number of halfspaces of  $R$  containing  $\Gamma_v$  but not containing  $\Gamma_{\text{parent}(v)}$ .
- $d_v$  is the minimum depth among all points in  $X_v$  with respect to the halfspaces of  $R$  crossing  $\Gamma_v$ .

The overall minimum depth with respect to  $R$  is given by the value  $d_v$  at the root  $v$ . Whenever we insert a halfspace  $h$  to  $R$ , for each of the  $O(n^{2/3+\delta})$  cells  $\Gamma_v$  crossed by  $h$ , we update the counters  $c_{v'}$  for the children  $v'$  of the nodes  $v$ ; we also update the value  $d_v$  bottom-up according to the formula  $d_v = \min_{\text{child } v' \text{ of } v} (d_{v'} + c_{v'})$ . Thus, all values can be maintained in  $O(n^{2/3+\delta})$  time per insertion to  $R$ .

As there are  $\tilde{O}(t)$  insertions to  $R$ , the total cost is  $\tilde{O}(tn^{2/3+\delta})$ . This cost covers the resetting of counters at every round.

(We remark that the idea of augmenting nodes of partition trees with counters appeared before in at least one prior work on dynamic geometric data structures [8, Theorem 4.1].)

**Weighted range sampling.** Let  $Q$  be the set of all points  $p$  for which we have performed multiplicity-doubling steps thus far. Note that  $|Q| = \tilde{O}(t)$ . The multiplicity of a halfspace  $h \in S$  is  $2^{\text{depth of } h^* \text{ in } Q^*}$ . To implicitly represent the multiplicities and their sum, we maintain two values at each node  $v$  of the partition tree for  $S^*$ :

- $c_v$  is the number of dual halfspaces of  $Q^*$  containing  $\Gamma_v$  but not containing  $\Gamma_{\text{parent}(v)}$ .
- $m_v$  is the sum of  $2^{\text{depth of } h^*}$  among the halfspaces of  $Q^*$  crossing  $\Gamma_v$  over all  $h^* \in S_v^*$ .

Whenever we insert a point  $p$  to  $Q$ , for each of the  $O(n^{2/3+\delta})$  cells  $\Gamma_v$  crossed by the dual halfspace  $p^*$ , we update the counters  $c_{v'}$  for the children  $v'$  of the nodes  $v$ ; we also update the value  $m_v$  bottom-up according to the formula  $m_v = \sum_{\text{child } v' \text{ of } v} 2^{c_{v'}} m_{v'}$ . Thus, all values can be maintained in  $O(n^{2/3+\delta})$  time per insertion to  $Q$ .

To generate a weighted sample of the halfspaces of  $S$  containing  $p$  for line 8, we find all  $O(n^{2/3+\delta})$  cells  $\Gamma_v$  crossed by the dual halfspace  $p^*$ , and consider the canonical subsets  $S_v^*$  for the children  $v'$  of  $v$  with  $\Gamma_{v'}$  contained in  $p^*$ . We can then sample from these canonical subsets, weighted by  $m_{v'} 2^{\sum_u c_u}$ , where the sum is over all ancestors  $u$  of  $v'$ . All this takes time  $O(n^{2/3+\delta})$  plus the size of the sample.

Hence, the total cost of all  $\tilde{O}(t)$  multiplicity-doubling steps is  $\tilde{O}(tn^{2/3+\delta})$ .

► **Lemma 4.** *There exists a data structure for the dynamic set cover problem for  $O(n)$  upper halfspaces and  $O(n)$  points in 3D that supports insertions and deletions in  $\tilde{O}(1)$  time and can find an  $O(1)$ -approximate solution w.h.p. in  $\tilde{O}(\text{OPT} \cdot n^{2/3+\delta})$  time.*

## 4.2 Algorithm for medium OPT

The preceding algorithm works well only when OPT is smaller than about  $n^{1/3}$ . We show that a more involved algorithm, also based on MWU, can achieve sublinear time even when OPT approaches  $n^{1-\delta}$ . The basic approach is to use *shallow* versions of the partition trees [23]. Several technical new ideas are required, but due to space limitation, we defer the details to the full paper and just state the result:

► **Lemma 5.** *There exists a data structure for the dynamic set cover problem for  $O(n)$  upper halfspaces and  $O(n)$  points in 3D that maintains an  $O(1)$ -approximate solution w.h.p. with  $\tilde{O}(\frac{n^{7/8}}{\text{OPT}^{1/8}} + \frac{n}{\sqrt{\text{OPT}}} + \text{OPT}^{1/3} n^{2/3+O(\delta)})$  amortized insertion and deletion time.*

## 4.3 Algorithm for large OPT

Lastly, we give an algorithm for the large OPT case, which is very different from the algorithms in the previous subsections (and not based on modifying MWU). Here, we can only compute the size of the approximate set cover, not the cover itself. Like before, we will show that the problem gets easier for large OPT, because we can afford a large additive error. The idea is to decompose the problem into subproblems via geometric sampling and planar separators, and then approximate the sum of the subproblems' answers by sampling again.

### 4.3.1 Data structures

We just store the dual point set  $S^*$  in a known 3D halfspace range reporting structure. The data structure by Chan [9] supports queries in  $O((\log n + k) \log n)$  time for output size  $k$ , and insertions and deletions in  $S$  in polylogarithmic amortized time.

We store the  $xy$ -projection of the point set  $X$  in a known 2D triangle range searching structure [22] that supports queries in  $O(\frac{n^{1/2+\delta}}{z^{1/2}} + k)$  time for output size  $k$ , and insertions and deletions in  $X$  in  $\tilde{O}(z)$  time for a given trade-off parameter  $z \in [1, n]$ .

### 4.3.2 Approximating the optimal value

Let  $b$  and  $g$  be parameters to be set later. Take a random sample  $R$  of the halfspaces  $S$  with size  $\frac{n}{b}$ . Imagine that  $R$  is included in the solution. The remaining uncovered space is the complement of the union of  $R$ , which is a 3D convex polyhedron. There are  $O(|R|) = O(\frac{n}{b})$  cells in the vertical decomposition  $\text{VD}(R)$  of this polyhedron (formed by triangulating each face and drawing a vertical wall at each edge of the triangulation). Each cell is crossed by  $O(b \log n)$  halfspaces w.h.p., by well-known geometric sampling analysis [13]. The decomposition  $\text{VD}(R)$  can be constructed in  $\tilde{O}(\frac{n}{b})$  time.

Our key idea is to use planar graph separators to divide into smaller subproblems. The following is a multi-cluster version of the standard planar separator theorem [20] (sometimes known as “ $r$ -divisions” [18]):

► **Lemma 6** (Planar Separator Theorem, Multi-Cluster Version). *Given a planar graph  $G = (V, E)$  with  $n$  vertices, and a parameter  $g$ , we can partition  $V$  into  $\frac{n}{g}$  subsets  $V_1, \dots, V_{n/g}$  of size  $O(g)$  each, and an extra “boundary set”  $B$  of size  $O(\frac{n}{\sqrt{g}})$ , such that no two vertices from different subsets  $V_i$  and  $V_j$  are adjacent. The partition can be constructed in  $\tilde{O}(n)$  time.*

(We remark that the general idea of combining cuttings/geometric sampling with planar graph separators appeared in some geometric approximation algorithms before, e.g., [1].)

We apply Lemma 6 to the dual graph of  $\text{VD}(R)$  (which has size  $O(\frac{n}{b})$ ), yielding  $O((n/b)/g)$  “clusters” of  $O(g)$  cells each, and a set  $B$  of  $O((n/b)/\sqrt{g})$  “boundary cells”, in  $\tilde{O}(\frac{n}{b})$  time.

Let  $S_B$  be the subset of all halfspaces of  $S$  that cross boundary cells of  $B$ . Note that  $|S_B| = O((n/b)/\sqrt{g} \cdot b \log n) = \tilde{O}(\frac{n}{\sqrt{g}})$  w.h.p.

For each cluster  $\gamma$ , let  $X_\gamma$  denote the subset of all points of  $X$  whose  $xy$ -projections lie in the  $xy$ -projection of the cells of  $\gamma$ , and let  $S_\gamma$  denote the subset of all halfspaces of  $S$  that cross the cells of  $\gamma$ . Note that  $|S_\gamma| = O(g \cdot b \log n) = \tilde{O}(bg)$  w.h.p. Let  $\text{OPT}_\gamma$  denote the optimal value for the set cover problem for the halfspaces of  $S_\gamma$  and the points of  $X_\gamma$  not covered by  $R \cup S_B$ .

▷ **Claim 7.**  $\sum_\gamma \text{OPT}_\gamma$  approximates  $\text{OPT}$  with additive error  $\tilde{O}(\frac{n}{b} + \frac{n}{\sqrt{g}})$  w.h.p.

*Proof.* A feasible solution can be formed by taking the union of the solutions corresponding to  $\text{OPT}_\gamma$ , together with  $R$  (to cover points not covered by  $\text{VD}(R)$ ) and  $S_B$  (to cover points inside boundary cells). As  $|R| = \frac{n}{b}$  and  $|S_B| = \tilde{O}(\frac{n}{\sqrt{g}})$  w.h.p., this proves that  $\text{OPT} \leq \sum_\gamma \text{OPT}_\gamma + \tilde{O}(\frac{n}{b} + \frac{n}{\sqrt{g}})$ .

In the other direction, observe that if a halfspace  $h$  crosses two different clusters  $\gamma_i$  and  $\gamma_j$ , it must also cross some boundary cell in  $X$  by convexity: pick points  $p \in h \cap \gamma_i$  and  $q \in h \cap \gamma_j$ ; then the line segment  $\overline{pq}$  must hit the wall of some boundary cell. So, after removing  $R \cup S_B$  from the global optimal solution, we get disjoint local solutions in the clusters. This proves that  $\text{OPT} \geq \sum_\gamma \text{OPT}_\gamma$ . ◁

We use the following known fact about approximating a sum via random sampling (which is of course a standard trick):

► **Lemma 8.** *Suppose  $a_1 + \dots + a_m = T$  where  $a_i \in [0, U]$ . Take a random subset  $R$  of  $r = \frac{(c_0/\varepsilon^2)mU \log n}{T}$  elements from  $a_1, \dots, a_m$ , for a sufficiently large constant  $c_0$ . Then  $\sum_{a_i \in R} a_i \cdot \frac{r}{m}$  is a  $(1 + \varepsilon)$ -approximation to  $T$  w.h.p.*

**Proof.** By rescaling the  $a_i$ ’s and  $T$  by a factor  $U$ , we may assume that  $U = 1$ . Define a random variable  $Y_i$ , which is  $a_i$  with probability  $\frac{r}{m}$ , and 0 otherwise. Then  $E[\sum_i Y_i] = T \cdot \frac{r}{m} = (c_0/\varepsilon^2) \log n$ . The result follows from a standard Chernoff bound on the  $Y_i$ ’s. ◀

By applying the above lemma with  $m = (n/b)/g$ ,  $T = \Theta(t)$ , and  $U = \tilde{O}(bg)$  (assuming  $\text{OPT}$  is finite), we can  $O(1)$ -approximate  $\text{OPT}$  by summing  $\text{OPT}_\gamma$  over a random sample of  $r = \tilde{O}(\frac{mU}{T}) = \tilde{O}(\frac{n}{t})$  clusters  $\gamma$ .

We can generate the set  $S_B$  by finding the halfspaces of  $S$  that contain the  $O((n/b)/\sqrt{g})$  vertices of the cells in  $B$  – this corresponds to  $O((n/b)/\sqrt{g})$  halfspace range reporting queries for the dual 3D point set  $S^*$ , each with output size  $\tilde{O}(b)$  w.h.p. and each taking  $\tilde{O}(b)$  time [9]. Thus,  $S_B$  can be found in  $\tilde{O}(\frac{n}{\sqrt{g}})$  time. We compute the union of  $R \cup S_B$ , which is the complement of an intersection of halfspaces, by the dual of 3D convex hull algorithm [17]. This takes  $\tilde{O}(\frac{n}{b} + \frac{n}{\sqrt{g}})$  time.

For each chosen cluster  $\gamma$ , we can generate  $S_\gamma$  similarly by  $O(g)$  halfspace range reporting queries for  $S^*$ , each with output size  $\tilde{O}(b)$  w.h.p. Thus,  $S_\gamma$  can be found in  $\tilde{O}(bg)$  time. We can generate  $X_\gamma$  by performing  $O(g)$  triangle range reporting queries for the 2D  $xy$ -projection of the point set  $X$ . Thus,  $X_\gamma$  can be found in  $\tilde{O}(g \frac{n^{1/2+\delta}}{z^{1/2}} + |X_\gamma|)$  time. We filter points of  $X_\gamma$  covered by  $R \cup S_B$ , by performing  $|X_\gamma|$  planar point location queries [17] in the  $xy$ -projection of the boundary of the union of  $R \cup S_B$ . This takes  $\tilde{O}(|X_\gamma|)$  time. We can then compute an  $O(1)$ -approximation to  $\text{OPT}_\gamma$  by running a known static set cover algorithm [4, 11] in  $\tilde{O}(bg + |X_\gamma|)$  time.

The expected sum of  $X_\gamma$  over all chosen clusters  $\gamma$  is  $O(n \cdot \frac{r}{m}) = O(rbg)$ . The total expected time over  $r$  clusters is  $\tilde{O}(rbg + rg \frac{n^{1/2+\delta}}{z^{1/2}}) = \tilde{O}(\frac{bgn}{t} + \frac{gn^{3/2+\delta}}{tz^{1/2}})$ . The overall expected running time is  $\tilde{O}(\frac{n}{b} + \frac{n}{\sqrt{g}} + \frac{bgn}{t} + \frac{gn^{3/2+\delta}}{tz^{1/2}})$ , and we obtain an  $(O(1), \tilde{O}(\frac{n}{b} + \frac{n}{\sqrt{g}}))$ -approximation. (The expected bound can be converted to worst-case by placing a time limit and re-running logarithmically many times.) Choosing  $g = b^2$  yields the following result:

► **Lemma 9.** *Given parameters  $b$  and  $z \in [1, n]$  and any constant  $\varepsilon > 0$ , there exists a data structure for the dynamic set cover problem for  $O(n)$  upper halfspaces and  $O(n)$  points in 3D that supports insertions and deletions in  $\tilde{O}(z)$  amortized time and can find the value of an  $(O(1), \tilde{O}(\frac{n}{b}))$ -approximation w.h.p. in  $\tilde{O}(\frac{n}{b} + \frac{b^3 n}{\text{OPT}} + \frac{b^2 n^{3/2+\delta}}{\text{OPT} \cdot z^{1/2}})$  time for any constant  $\delta > 0$ .*

A minor technicality is that when applying Lemma 8, we have assumed that the optimal value is finite. The problem of checking whether a solution exists, i.e., whether a point set is covered by a set of halfspaces (or more generally, maintaining the lowest-depth point), subject to insertions and deletions of points and halfspaces, has already been solved before by Chan [8, Theorem 4.1], who gave a fully dynamic algorithm with  $\tilde{O}(n^{2/3})$  time per operation (based on augmenting partition trees with counters, similar to what we have done here).

**Combining the algorithms.** Finally, we combine all three algorithms:

1. When  $\text{OPT} \leq n^{2/9}$ , we use the algorithm for small OPT; the running time is  $\tilde{O}(\text{OPT} \cdot n^{2/3+\delta}) \leq \tilde{O}(n^{8/9+\delta})$ .
2. When  $n^{2/9} < \text{OPT} \leq n^{10/13}$ , we use the algorithm for medium OPT; the running time is  $\tilde{O}(\frac{n^{7/8}}{\text{OPT}^{1/8}} + \frac{n}{\sqrt{\text{OPT}}} + \text{OPT}^{1/3} n^{2/3+O(\delta)}) \leq O(n^{12/13+O(\delta)})$ .
3. When  $\text{OPT} > n^{10/13}$ , we use the algorithm for large OPT with  $b = \tilde{\Theta}(n^{3/13})$  and  $z = n^{7/13}$ , so that an  $(O(1), \tilde{O}(\frac{n}{b}))$ -approximation is indeed an  $O(1)$ -approximation; the running time is  $\tilde{O}(\frac{n}{b} + \frac{b^3 n}{\text{OPT}} + \frac{b^2 n^{3/2+\delta}}{\text{OPT} \cdot z^{1/2}}) \leq O(n^{12/13+O(\delta)})$ .

► **Theorem 10.** *There exists a data structure for the dynamic set cover problem for  $O(n)$  upper halfspaces and  $O(n)$  points in 3D that maintains the value of an  $O(1)$ -approximate solution w.h.p. with  $O(n^{12/13+\delta})$  amortized insertion and deletion time for any constant  $\delta > 0$ .*

---

## References

- 1 Anna Adamaszek, Sariel Har-Peled, and Andreas Wiese. Approximation schemes for independent set and sparse subsets of polygons. *Journal of the ACM*, 66(4):29:1–29:40, 2019. doi:10.1145/3326122.
- 2 Pankaj K. Agarwal, Hsien-Chih Chang, Subhash Suri, Allen Xiao, and Jie Xue. Dynamic geometric set cover and hitting set. In *Proceedings of the 36th Symposium on Computational Geometry (SoCG)*, volume 164, pages 2:1–2:15, 2020. doi:10.4230/LIPICs.SoCG.2020.2.

- 3 Pankaj K. Agarwal and Jeff Erickson. Geometric range searching and its relatives. In B. Chazelle, J. E. Goodman, and R. Pollack, editors, *Advances in Discrete and Computational Geometry*, pages 1–56. AMS Press, 1999. URL: <http://jeffe.cs.illinois.edu/pubs/survey.html>.
- 4 Pankaj K. Agarwal and Jiangwei Pan. Near-linear algorithms for geometric hitting sets and set covers. *Discrete & Computational Geometry*, 63(2):460–482, 2020. Preliminary version in SoCG’14. doi:10.1007/s00454-019-00099-6.
- 5 Sunil Arya, David M. Mount, Nathan S. Netanyahu, Ruth Silverman, and Angela Y. Wu. An optimal algorithm for approximate nearest neighbor searching in fixed dimensions. *Journal of the ACM*, 45(6):891–923, 1998. doi:10.1145/293347.293348.
- 6 Hervé Brönnimann and Michael T. Goodrich. Almost optimal set covers in finite VC-dimension. *Discrete & Computational Geometry*, 14(4):463–479, 1995.
- 7 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.
- 8 Timothy M. Chan. Semi-online maintenance of geometric optima and measures. *SIAM Journal on Computing*, 32(3):700–716, 2003. doi:10.1137/S0097539702404389.
- 9 Timothy M. Chan. A dynamic data structure for 3-d convex hulls and 2-d nearest neighbor queries. *Journal of the ACM*, 57(3):16:1–16:15, 2010. doi:10.1145/1706591.1706596.
- 10 Timothy M. Chan. Optimal partition trees. *Discrete & Computational Geometry*, 47(4):661–690, 2012. doi:10.1007/s00454-012-9410-z.
- 11 Timothy M. Chan and Qizheng He. Faster approximation algorithms for geometric set cover. In *Proceedings of the 36th Symposium on Computational Geometry (SoCG)*, volume 164, pages 27:1–27:14, 2020. doi:10.4230/LIPIcs.SocG.2020.27.
- 12 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.
- 13 Kenneth L. Clarkson. New applications of random sampling in computational geometry. *Discrete & Computational Geometry*, 2:195–222, 1987. doi:10.1007/BF02187879.
- 14 Kenneth L. Clarkson. Algorithms for polytope covering and approximation. In *Workshop on Algorithms and Data Structures*, pages 246–252, 1993.
- 15 Kenneth L. Clarkson and Kasturi Varadarajan. Improved approximation algorithms for geometric set cover. *Discrete & Computational Geometry*, 37(1):43–58, 2007.
- 16 Artur Czumaj, Funda Ergün, Lance Fortnow, Avner Magen, Ilan Newman, Ronitt Rubinfeld, and Christian Sohler. Approximating the weight of the Euclidean minimum spanning tree in sublinear time. *SIAM Journal on Computing*, 35(1):91–109, 2005. doi:10.1137/S0097539703435297.
- 17 Mark de Berg, Otfried Cheong, Marc J. van Kreveld, and Mark H. Overmars. *Computational Geometry: Algorithms and Applications*. Springer, 3rd edition, 2008. URL: <https://www.worldcat.org/oclc/227584184>.
- 18 Greg N. Frederickson. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM Journal on Computing*, 16(6):1004–1022, 1987.
- 19 Monika Henzinger, Stefan Neumann, and Andreas Wiese. Dynamic approximate maximum independent set of intervals, hypercubes and hyperrectangles. In *Proceedings of the 36th Symposium on Computational Geometry (SoCG)*, volume 164, pages 51:1–51:14, 2020. doi:10.4230/LIPIcs.SocG.2020.51.
- 20 Richard J. Lipton and Robert Endre Tarjan. Applications of a planar separator theorem. *SIAM Journal on Computing*, 9(3):615–627, 1980. doi:10.1137/0209046.
- 21 Chih-Hung Liu, Evanthia Papadopoulou, and D. T. Lee. An output-sensitive approach for the  $L_1/L_\infty$   $k$ -nearest-neighbor Voronoi diagram. In *Proceedings of the 19th Annual European Symposium on Algorithms (ESA)*, volume 6942 of *Lecture Notes in Computer Science*, pages 70–81. Springer, 2011. doi:10.1007/978-3-642-23719-5\_7.

## 25:14 More Dynamic Data Structures for Geometric Set Cover

- 22 Jiří Matoušek. Efficient partition trees. *Discrete & Computational Geometry*, 8(3):315–334, 1992.
- 23 Jiří Matoušek. Reporting points in halfspaces. *Computational Geometry*, 2(3):169–186, 1992.
- 24 Jiří Matoušek. Range searching with efficient hierarchical cutting. *Discrete & Computational Geometry*, 10:157–182, 1993. doi:10.1007/BF02573972.
- 25 Nabil H. Mustafa, Rajiv Raman, and Saurabh Ray. Quasi-polynomial time approximation scheme for weighted geometric set cover on pseudodisks and halfspaces. *SIAM Journal on Computing*, 44(6):1650–1669, 2015. doi:10.1137/14099317X.
- 26 Nabil H. Mustafa and Saurabh Ray. Improved results on geometric hitting set problems. *Discrete & Computational Geometry*, 44(4):883–895, 2010. doi:10.1007/s00454-010-9285-9.
- 27 Micha Sharir. On  $k$ -sets in arrangement of curves and surfaces. *Discrete & Computational Geometry*, 6:593–613, 1991. doi:10.1007/BF02574706.