

Dynamic geometric set cover and hitting set

Pankaj K. Agarwal *
pankaj@cs.duke.edu

Hsien-Chih Chang *
hsienchih.chang@duke.edu

Subhash Suri †
suri@cs.ucsb.edu

Allen Xiao *
axiao@cs.duke.edu

Jie Xue †
jiexue@ucsb.edu

Abstract

We investigate dynamic versions of geometric set cover and hitting set where points and ranges may be inserted or deleted, and we want to efficiently maintain an (approximately) optimal solution for the current problem instance. While their static versions have been extensively studied in the past, surprisingly little is known about dynamic geometric set cover and hitting set. For instance, even for the most basic case of one-dimensional interval set cover and hitting set, no nontrivial results were known. The main contribution of our paper are two frameworks that lead to efficient data structures for dynamically maintaining set covers and hitting sets in \mathbb{R}^1 and \mathbb{R}^2 . The first framework uses bootstrapping and gives a $(1 + \varepsilon)$ -approximate data structure for dynamic interval set cover in \mathbb{R}^1 with $O(n^\alpha/\varepsilon)$ amortized update time for any constant $\alpha > 0$; in \mathbb{R}^2 , this method gives $O(1)$ -approximate data structures for unit-square (and quadrant) set cover and hitting set with $O(n^{1/2+\alpha})$ amortized update time. The second framework uses local modification, and leads to a $(1 + \varepsilon)$ -approximate data structure for dynamic interval hitting set in \mathbb{R}^1 with $\tilde{O}(1/\varepsilon)$ amortized update time; in \mathbb{R}^2 , it gives $O(1)$ -approximate data structures for unit-square (and quadrant) set cover and hitting set in the *partially* dynamic settings with $\tilde{O}(1)$ amortized update time.

1 Introduction

Given a pair (S, \mathcal{R}) where S is a set of points and \mathcal{R} is a collection of geometric ranges in a Euclidean space, the *geometric set cover* (resp., *hitting set*) problem is to find the smallest number of ranges in \mathcal{R} (resp., points in S) that cover all points in S (resp., hit all ranges in \mathcal{R}). Geometric set cover and hitting set are classical geometric optimization problems, with numerous applications in databases, sensor networks, VLSI design, etc.

In many applications, the problem instance can change over time and re-computing a new solution after each change is too costly. In these situations, a dynamic algorithm that can update the solution after a change more efficiently than constructing the entire new solution from scratch is highly desirable. This motivates the main problem studied in our paper: dynamically maintaining geometric set covers and hitting sets under insertion and deletion of points and ranges.

Although (static) geometric set cover and hitting set have been extensively studied over the years, their dynamic variants are surprisingly open. For example, even for the most fundamental case, dynamic *interval* set cover and hitting set in one dimension, no nontrivial results were previously known. In this paper, we propose two algorithmic frameworks for the problems, which lead to efficient data structures for dynamic set cover and hitting set for intervals in \mathbb{R}^1 and unit squares and quadrants in \mathbb{R}^2 . We believe that our approaches can be extended to solve dynamic set cover and hitting set in other geometric settings, or more generally, other dynamic problems in computational geometry.

*Department of Computer Science, Duke University, USA

†Department of Computer Science, University of California at Santa Barbara, USA

1.1 Related work

The set cover and hitting set problems in general setting are well-known to be NP-complete [18]. A simple greedy algorithm achieves an $O(\log n)$ -approximation [11, 19, 21], which is tight under appropriate complexity-theoretic assumptions [13, 20]. In many geometric settings, the problems remain NP-hard or even hard to approximate [5, 22, 23]. However, by exploiting the geometric nature of the problems, efficient algorithms with better approximation factors can be obtained. For example, Mustafa and Ray [25] showed the existence of polynomial-time approximation schemes (PTAS) for halfspace hitting set in \mathbb{R}^3 and disk hitting set. There is also a PTAS for unit-square set cover given by Erlebach and van Leeuwen [14]. Agarwal and Pan [2] proposed approximation algorithms with near-linear running time for the set cover and hitting set problems for halfspaces in \mathbb{R}^3 , disks in \mathbb{R}^2 , and orthogonal rectangles.

Dynamic problems have received a considerable attention in recent years [3, 6, 7, 9, 10, 17, 26, 27]. In particular, dynamic set cover in general setting has been studied in [1, 8, 16]. All the results were achieved in the partially dynamic setting where ranges are fixed and only points are dynamic. Gupta et al. [16] showed that an $O(\log n)$ -approximation can be maintained using $O(f \log n)$ amortized update time and an $O(f^3)$ -approximation can be maintained using $O(f^2)$ amortized update time, where f is the maximum number of ranges that a point belongs to. Bhattacharya et al. [8] gave an $O(f^2)$ -approximation data structure for dynamic set cover with $O(f \log n)$ amortized update time. Abboud et al. [1] proved that one can maintain a $(1 + \epsilon)f$ -approximation using $O(f^2 \log n / \epsilon^5)$ amortized update time.

In geometric settings, only the dynamic hitting set problem has been considered [15]. Ganjugunte [15] studied two different dynamic settings: **(i)** only the range set \mathcal{R} is dynamic and **(ii)** \mathcal{R} is dynamic and S is semi-dynamic (i.e., insertion-only). Ganjugunte [15] showed that, for pseudo-disks in \mathbb{R}^2 , dynamic hitting set in setting **(i)** can be solved using $O(\gamma(n) \log^4 n)$ amortized update time with approximation factor $O(\log^2 n)$, and that in setting **(ii)** can be solved using $O(\gamma(n) \sqrt{n} \log^4 n)$ amortized update time with approximation factor $O(\log^6 n / \log \log n)$, where $\gamma(n)$ is the time for finding a point in X contained in a query pseudo-trapezoid (see [15] for details). Dynamic geometric hitting set in the fully dynamic setting (where both points and ranges can be inserted or deleted) as well as dynamic geometric set cover has not yet been studied before, to the best of our knowledge.

1.2 Our results

Let (S, \mathcal{R}) be a dynamic geometric set cover (resp., hitting set) instance. We are interested in proposing efficient data structures to *maintain* an (approximately) optimal solution for (S, \mathcal{R}) . This may have various definitions, resulting in different variants of the problem. A natural variant is to maintain the number opt , which is the size of an optimal set cover (resp., hitting set) for (S, \mathcal{R}) , or an approximation of opt . However, in many applications, only maintaining the optimum is not satisfactory and one may hope to maintain a “real” set cover (resp., hitting set) for the dynamic instance. Therefore, in this paper, we formulate the problem as follows. We require the data structure to, after each update, store (implicitly) a solution for the current problem instance satisfying some quality requirement such that certain information about the solution can be queried efficiently. For example, one can ask how large the solution is, whether a specific element in \mathcal{R} (resp., S) is used in the solution, what the entire solution is, etc. We will make this more precise shortly.

In dynamic settings, it is usually more natural and convenient to consider a *multiset* solution for the set cover (resp., hitting set) instance. That is, we allow the solution to be a *multiset* of elements in \mathcal{R} (resp., S) that cover all points in S (resp., hit all ranges in \mathcal{R}), and the quality of the solution is also evaluated in terms of the multiset cardinality. In the static problems, one can always efficiently remove the duplicates in a multiset solution to obtain a (ordinary) set cover or hitting set with even better quality (i.e., smaller cardinality), hence computing a multiset solution is essentially equivalent to computing an ordinary solution. However, in dynamic settings, the update time has to be sublinear, and in general this is not sufficient for detecting and removing duplicates. Therefore, in this paper, we mainly focus on multiset solutions (though some of our data structures can maintain an ordinary set cover or hitting set). Unless explicitly mentioned otherwise, solutions for set cover and hitting set always refer to multiset solutions hereafter.

Precisely, we require a dynamic set cover (resp., hitting set) data structure to store implicitly, after each update, a set cover \mathcal{R}' (resp., a hitting set S') for the current instance (S, \mathcal{R}) such that the following queries are supported.

- **Size query:** reporting the (multiset) size of \mathcal{R}' (resp., S').
- **Membership query:** reporting, for a given range $R \in \mathcal{R}$ (resp., a given point $a \in S$), the number of copies of R (resp., a) contained in \mathcal{R}' (resp., S').
- **Reporting query:** reporting all the elements in \mathcal{R}' (resp., S').

We require the size query to be answered in $O(1)$ time, a membership query to be answered in $O(\log |\mathcal{R}'|)$ time (resp., $O(\log |S'|)$ time), and the reporting query to be answered in $O(|\mathcal{R}'|)$ time (resp., $O(|S'|)$ time); this is the best one can expect in the pointer machine model.

We say that a set cover (resp., hitting set) instance is *fully dynamic* if insertions and deletions on both points and ranges are allowed, and *partially dynamic* if only the points (resp., ranges) can be inserted and deleted. This paper mainly focuses on the fully dynamic setting, while some results are achieved in the partially dynamic setting. Thus, unless explicitly mentioned otherwise, problems are always considered in the fully dynamic setting.

The main contribution of this paper are two frameworks for designing dynamic geometric set cover and hitting set data structures, leading to efficient data structures in \mathbb{R}^1 and \mathbb{R}^2 (see Table 1). The first framework is based on bootstrapping, which results in efficient (approximate) dynamic data structures for interval set cover, quadrant/unit-square set cover and hitting set (see the first three rows of Table 1 for detailed bounds). The second framework is based on local modification, which results in efficient (approximate) dynamic data structures for interval hitting set, quadrant/unit-square set cover and hitting set in the *partially* dynamic setting (see the last three rows of Table 1 for detailed bounds).

Framework	Problem	Range	Approx. factor	Update time	Setting
Bootstrapping	SC	Interval	$1 + \varepsilon$	$\tilde{O}(n^\alpha/\varepsilon)$	Fully dynamic
	SC & HS	Quadrant	$O(1)$	$\tilde{O}(n^{1/2+\alpha})$	Fully dynamic
	SC & HS	Unit square	$O(1)$	$\tilde{O}(n^{1/2+\alpha})$	Fully dynamic
Local modification	HS	Interval	$1 + \varepsilon$	$\tilde{O}(1/\varepsilon)$	Fully dynamic
	SC & HS	Quadrant	$O(1)$	$\tilde{O}(1)$	Part. dynamic
	SC & HS	Unit square	$O(1)$	$\tilde{O}(1)$	Part. dynamic

Table 1: Summary of our results for dynamic geometric set cover and hitting set (SC = set cover and HS = hitting set). All update times are amortized. The notation $\tilde{O}(\cdot)$ hides logarithmic factors, n is the size of the current instance, and $\alpha > 0$ is any small constant. All data structures can be constructed in $\tilde{O}(n_0)$ time where n_0 is the size of the initial instance.

Organization. The rest of the paper is organized as follows. Section 2 gives the preliminaries required for the paper and Section 3 gives an overview of our two frameworks. Our first framework (bootstrapping) and second framework (local modification) for designing efficient dynamic geometric set cover and hitting set data structures are presented in Section 4 and Section 5, respectively. To make the paper more readable, the proofs of the technical lemmas and some details are deferred to the appendix.

2 Preliminaries

In this section, we introduce the basic notions used throughout the paper.

Multi-sets and disjoint union. A *multi-set* is a set in which elements can have multiple copies. The *multiplicity* of an element a in a multi-set A is the number of the copies of a in A . For two multi-sets A and B , we use $A \sqcup B$ to denote the *disjoint union* of A and B , in which the multiplicity of an element a is the sum of the multiplicities of a in A and B .

Basic data structures. A data structure built on a dataset (e.g., point set, range set, set cover or hitting set instances) of size n is *basic* if it can be constructed in $\tilde{O}(n)$ time and can be dynamized with $\tilde{O}(1)$ update time (with a bit abuse of terminology, sometimes we also use “basic data structures” to denote the dynamized version of such data structures).

Output-sensitive algorithms. In some set cover and hitting set problems, if the problem instance is properly stored in some data structure, it is possible to compute an (approximate) optimal solution in sub-linear time. An *output-sensitive* algorithm for a set cover or hitting set problem refers to an algorithm that can compute an (approximate) optimal solution in $\tilde{O}(\text{out})$ time (where out is the size of the output solution), by using some *basic* data structure built on the problem instance.

3 An overview of our two frameworks

The basic idea of our first framework is *bootstrapping*. Namely, we begin from a simple inefficient dynamic set cover or hitting set data structure (e.g., a data structure that re-computes a solution after each update), and repeatedly use the current data structure to obtain an improved one. The main challenge here is to design the bootstrapping procedure: how to use a given data structure to construct a new data structure with improved update time. We achieve this by using output-sensitive algorithms and carefully partitioning the problem instances to sub-instances.

Our second framework is much simpler, which is based on *local modification*. Namely, we construct a new solution by slightly modifying the previous one after each update, and re-compute a new solution periodically using an output-sensitive algorithm. This framework applies to the problems which are *stable*, in the sense that the optimum of a dynamic instance does not change significantly.

4 First framework: Bootstrapping

In this section, we present our first framework for dynamic geometric set cover and hitting set, which is based on bootstrapping and results in sub-linear data structures for dynamic interval set cover and dynamic quadrant and unit-square set cover (resp., hitting set).

4.1 Warm-up: 1D set cover for intervals

As a warm up, we first study the 1D problem: dynamic interval set cover. First, we observe that interval set cover admits a simple *exact* output-sensitive algorithm. Indeed, interval set cover can be solved using the greedy algorithm that repeatedly picks the leftmost uncovered point and covers it using the interval with the rightmost right endpoint, and the algorithm can be easily made output-sensitive if we store the points and intervals in binary search trees.

Lemma 1. *Interval set cover admits an exact output-sensitive algorithm.*

This algorithm will serve an important role in the design of our data structure.

4.1.1 Bootstrapping

As mentioned before, our data structure is designed using bootstrapping. Specifically, we prove the following bootstrapping theorem, which is the technical heart of our result. The theorem roughly states that given a

dynamic interval set cover data structure, one can obtain another dynamic interval set cover data structure with improved update time.

Theorem 2. *Let $\alpha \in [0, 1]$ be a number. If there exists a $(1 + \varepsilon)$ -approximate dynamic interval set cover data structure \mathcal{D}_{old} with $\tilde{O}(n^\alpha/\varepsilon^{1-\alpha})$ amortized update time and $\tilde{O}(n_0)$ construction time for any $\varepsilon > 0$, then there exists a $(1 + \varepsilon)$ -approximate dynamic interval set cover data structure \mathcal{D}_{new} with $\tilde{O}(n^{\alpha'}/\varepsilon^{1-\alpha'})$ amortized update time and $\tilde{O}(n_0)$ construction time for any $\varepsilon > 0$, where $\alpha' = \alpha/(1 + \alpha)$. Here n (resp., n_0) denotes the size of the current (resp., initial) problem instance.*

Assuming the existence of \mathcal{D}_{old} as in the theorem, we are going to design the improved data structure \mathcal{D}_{new} . Let (S, \mathcal{I}) be a dynamic interval set cover instance, and $\varepsilon > 0$ be the approximation factor. We denote by n (resp., n_0) the size of the current (resp., initial) (S, \mathcal{I}) .

The construction of \mathcal{D}_{new} . Initially, $|S| + |\mathcal{I}| = n_0$. Essentially, our data structure \mathcal{D}_{new} consists of two parts¹. The first part is the basic data structure \mathcal{A} required for the output-sensitive algorithm of Lemma 1. The second part is a family of \mathcal{D}_{old} data structures defined as follows. Let f be a function to be determined shortly. We partition the real line \mathbb{R} into $r = \lceil n_0/f(n_0, \varepsilon) \rceil$ connected portions (i.e., intervals) J_1, \dots, J_r such that each portion J_i contains $O(f(n_0, \varepsilon))$ points in S and $O(f(n_0, \varepsilon))$ endpoints of the intervals in \mathcal{I} . Define $S_i = S \cap J_i$ and define $\mathcal{I}_i \subseteq \mathcal{I}$ as the sub-collections consisting of the intervals that “partially intersect” J_i , i.e., $\mathcal{I}_i = \{I \in \mathcal{I} : J_i \cap I \neq \emptyset \text{ and } J_i \not\subseteq I\}$. When the instance (S, \mathcal{I}) is updated, the partition J_1, \dots, J_r will remain unchanged, but the S_i ’s and \mathcal{I}_i ’s will change along with S and \mathcal{I} . We view each (S_i, \mathcal{I}_i) as a dynamic interval set cover instance, and let $\mathcal{D}_{\text{old}}^{(i)}$ be the data structure \mathcal{D}_{old} built on (S_i, \mathcal{I}_i) for the approximation parameter $\tilde{\varepsilon} = \varepsilon/2$. Thus, $\mathcal{D}_{\text{old}}^{(i)}$ maintains a $(1 + \tilde{\varepsilon})$ -approximate optimal set cover for (S_i, \mathcal{I}_i) . The second part of \mathcal{D}_{new} consists of the data structures $\mathcal{D}_{\text{old}}^{(1)}, \dots, \mathcal{D}_{\text{old}}^{(r)}$.

Update and reconstruction. After an operation on (S, \mathcal{I}) , we update the basic data structure \mathcal{A} . Also, we update the data structure $\mathcal{D}_{\text{old}}^{(i)}$ if the instance (S_i, \mathcal{I}_i) changes due to the operation. Note that an operation on S changes exactly one S_i and an operation on \mathcal{I} changes at most two \mathcal{I}_i ’s (because an interval can belong to at most two \mathcal{I}_i ’s). Thus, we in fact only need to update at most two $\mathcal{D}_{\text{old}}^{(i)}$ ’s. Besides the update, we also reconstruct the entire data structure \mathcal{D}_{new} periodically (as is the case for many dynamic data structures). Specifically, the first reconstruction of \mathcal{D}_{new} happens after processing $f(n_0, \varepsilon)$ operations. The reconstruction is the same as the initial construction of \mathcal{D}_{new} , except that n_0 is replaced with n_1 , the size of (S, \mathcal{I}) at the time of reconstruction. Then the second reconstruction happens after processing $f(n_1, \varepsilon)$ operations since the first reconstruction, and so forth.

Maintaining a solution. We now discuss how to maintain a $(1 + \varepsilon)$ -approximate optimal set cover $\mathcal{I}_{\text{appx}}$ for (S, \mathcal{I}) . Let opt denote the optimum (i.e., the size of an optimal set cover) of the current (S, \mathcal{I}) . Set $\delta = \min\{(6 + 2\varepsilon) \cdot r/\varepsilon, n\}$. If $\text{opt} \leq \delta$, then the output-sensitive algorithm can compute an optimal set cover for (S, \mathcal{I}) in $\tilde{O}(\delta)$ time. Thus, we simulate the output-sensitive algorithm within that amount of time. If the algorithm successfully computes a solution, we use it as our $\mathcal{I}_{\text{appx}}$. Otherwise, we construct $\mathcal{I}_{\text{appx}}$ as follows. For $i \in \{1, \dots, r\}$, we say J_i is *coverable* if there exists $I \in \mathcal{I}$ such that $J_i \subseteq I$ and *uncoverable* otherwise. Let $P = \{i : J_i \text{ is coverable}\}$ and $P' = \{i : J_i \text{ is uncoverable}\}$. We try to use the intervals in \mathcal{I} to “cover” all coverable portions. That is, for each $i \in P$, we find an interval in \mathcal{I} that contains J_i , and denote by \mathcal{I}^* the collection of these intervals. Then we consider the uncoverable portions. If for some $i \in P'$, the data structure $\mathcal{D}_{\text{old}}^{(i)}$ tells us that the current (S_i, \mathcal{I}_i) does not have a set cover, then we immediately make a no-solution decision, i.e., decide that the current (S, \mathcal{I}) has no feasible set cover, and continue to the next operation. Otherwise, for every $i \in P'$, the data structure $\mathcal{D}_{\text{old}}^{(i)}$ maintains a $(1 + \tilde{\varepsilon})$ -approximate optimal set cover \mathcal{I}_i^* for (S_i, \mathcal{I}_i) . We then define $\mathcal{I}_{\text{appx}} = \mathcal{I}^* \sqcup (\bigsqcup_{i \in P'} \mathcal{I}_i^*)$.

¹In implementation level, we may need some additional support data structures (which are very simple). For simplicity of exposition, we shall mention them when discussing the implementation details.

Later we will prove that $\mathcal{I}_{\text{appx}}$ is always a $(1 + \varepsilon)$ -approximate optimal set cover for (S, \mathcal{I}) . Before this, let us consider how to store $\mathcal{I}_{\text{appx}}$ properly to support the size, membership, and reporting queries in the required query times. If $\mathcal{I}_{\text{appx}}$ is computed by the output-sensitive algorithm, then the size of $\mathcal{I}_{\text{appx}}$ is at most δ , and we have all the elements of $\mathcal{I}_{\text{appx}}$ in hand. In this case, it is not difficult to build a data structure on $\mathcal{I}_{\text{appx}}$ to support the desired queries. On the other hand, if $\mathcal{I}_{\text{appx}}$ is defined as the disjoint union of \mathcal{I}^* and \mathcal{I}_i^* 's, the size of $\mathcal{I}_{\text{appx}}$ might be very large and thus we are not able to explicitly extract all elements of $\mathcal{I}_{\text{appx}}$. Fortunately, in this case, each \mathcal{I}_i^* is already maintained in the data structure $\mathcal{D}_{\text{old}}^{(i)}$. Therefore, we actually only need to compute P , P' , and \mathcal{I}^* ; with these in hand, one can already build a data structure to support the desired queries for $\mathcal{I}_{\text{appx}}$. We defer the detailed discussion to Appendix B.

Correctness. We now prove the correctness of our data structure \mathcal{D}_{new} . We first show the correctness of the no-solution decision.

Lemma 3. *\mathcal{D}_{new} makes a no-solution decision iff the current (S, \mathcal{I}) has no set cover.*

Next, we show that the solution $\mathcal{I}_{\text{appx}}$ maintained by \mathcal{D}_{new} is truly a $(1 + \varepsilon)$ -approximate optimal set cover for (S, \mathcal{I}) . If $\mathcal{I}_{\text{appx}}$ is computed by the output-sensitive algorithm, then it is an optimal set cover for (S, \mathcal{I}) . Otherwise, $\text{opt} > \delta = \min\{(6 + 2\varepsilon) \cdot r/\varepsilon, n\}$, i.e., either $\text{opt} > (6 + 2\varepsilon) \cdot r/\varepsilon$ or $\text{opt} > n$. If $\text{opt} > n$, then the current (S, \mathcal{I}) has no set cover (i.e., $\text{opt} = \infty$) and thus \mathcal{D}_{new} makes a no-solution decision by Lemma 3. So assume $\text{opt} > (6 + 2\varepsilon) \cdot r/\varepsilon$. In this case, $\mathcal{I}_{\text{appx}} = \mathcal{I}^* \sqcup (\bigsqcup_{i \in P'} \mathcal{I}_i^*)$. For each $i \in P'$, let opt_i be the optimum of the instance (S_i, \mathcal{I}_i) . Then we have $|\mathcal{I}_i^*| \leq (1 + \tilde{\varepsilon}) \cdot \text{opt}_i$ for all $i \in P'$ where $\tilde{\varepsilon} = \varepsilon/2$. Since $|\mathcal{I}^*| \leq r$, we have

$$|\mathcal{I}_{\text{appx}}| = |\mathcal{I}^*| + \sum_{i \in P'} |\mathcal{I}_i^*| \leq r + \left(1 + \frac{\varepsilon}{2}\right) \sum_{i \in P'} \text{opt}_i. \quad (1)$$

Let \mathcal{I}_{opt} be an optimal set cover for (S, \mathcal{I}) . We observe that for $i \in P'$, $\mathcal{I}_{\text{opt}} \cap \mathcal{I}_i$ is a set cover for (S_i, \mathcal{I}_i) , because J_i is uncovered (so the points in S_i cannot be covered by any interval in $\mathcal{I} \setminus \mathcal{I}_i$). It immediately follows that $\text{opt}_i \leq |\mathcal{I}_{\text{opt}} \cap \mathcal{I}_i|$ for all $i \in P'$. Therefore, we have

$$\sum_{i \in P'} \text{opt}_i \leq \sum_{i \in P'} |\mathcal{I}_{\text{opt}} \cap \mathcal{I}_i|. \quad (2)$$

The right-hand side of the above inequality can be larger than $|\mathcal{I}_{\text{opt}}|$ as some intervals in \mathcal{I}_{opt} can belong to two \mathcal{I}_i 's. The following lemma bounds the number of such intervals.

Lemma 4. *There are at most $2r$ intervals in \mathcal{I}_{opt} that belong to exactly two \mathcal{I}_i 's.*

The above lemma immediately implies

$$\sum_{i \in P'} |\mathcal{I}_{\text{opt}} \cap \mathcal{I}_i| \leq |\mathcal{I}_{\text{opt}}| + 2r = \text{opt} + 2r. \quad (3)$$

Combining Inequalities 1, 2, and 3, we deduce that

$$\begin{aligned} |\mathcal{I}_{\text{appx}}| &\leq r + \left(1 + \frac{\varepsilon}{2}\right) \sum_{i \in P'} \text{opt}_i \\ &\leq r + \left(1 + \frac{\varepsilon}{2}\right) \sum_{i \in P'} |\mathcal{I}_{\text{opt}} \cap \mathcal{I}_i| \\ &\leq r + \left(1 + \frac{\varepsilon}{2}\right) \cdot (\text{opt} + 2r) = (3 + \varepsilon) \cdot r + \left(1 + \frac{\varepsilon}{2}\right) \cdot \text{opt} \\ &< \frac{\varepsilon}{2} \cdot \text{opt} + \left(1 + \frac{\varepsilon}{2}\right) \cdot \text{opt} = (1 + \varepsilon) \cdot \text{opt}, \end{aligned}$$

where the last inequality follows from the assumption $\text{opt} > (6 + 2\varepsilon) \cdot r/\varepsilon$.

Time analysis. We briefly discuss the update and construction time of \mathcal{D}_{new} ; a detailed analysis can be found in Appendix B. Since \mathcal{D}_{new} is reconstructed periodically, it suffices to consider the first period (i.e., the period before the first reconstruction). The construction of \mathcal{D}_{new} can be easily done in $\tilde{O}(n_0)$ time. The update time of \mathcal{D}_{new} consists of the time for updating the data structures \mathcal{A} and $\mathcal{D}_{\text{old}}^{(1)}, \dots, \mathcal{D}_{\text{old}}^{(r)}$, the time for maintaining the solution, and the time for reconstruction. Since the period consists of $f(n_0, \varepsilon)$ operations, the size of each (S_i, \mathcal{L}_i) is always bounded by $O(f(n_0, \varepsilon))$ during the period. As argued before, we only need to update at most two $\mathcal{D}_{\text{old}}^{(i)}$'s after each operation. Thus, updating the \mathcal{D}_{old} data structures takes $\tilde{O}(f(n_0, \varepsilon)^\alpha / \varepsilon^{1-\alpha})$ amortized time. Maintaining the solution can be done in $\tilde{O}(\delta + r)$ time, with a careful implementation. The time for reconstruction is bounded by $\tilde{O}(n_0 + f(n_0, \varepsilon))$; we amortize it over the $f(n_0, \varepsilon)$ operations in the period and the amortized time cost is then $\tilde{O}(n_0 / f(n_0, \varepsilon))$, i.e., $\tilde{O}(r)$. In total, the amortized update time of \mathcal{D}_{new} (during the first period) is $\tilde{O}(f(n_0, \varepsilon)^\alpha / \varepsilon^{1-\alpha} + \delta + r)$. If we set $f(n, \varepsilon) = \min\{n^{1-\alpha'} / \varepsilon^{\alpha'}, n/2\}$ where α' is as defined in Theorem 2, a careful calculation (see Appendix B) shows that the amortized update time becomes $\tilde{O}(n^{\alpha'} / \varepsilon^{1-\alpha'})$.

4.1.2 Putting everything together

With the bootstrapping theorem in hand, we are now able to design our dynamic interval set cover data structure. The starting point is a “trivial” data structure, which simply uses the output-sensitive algorithm of Lemma 1 to re-compute an optimal interval set cover after each update. Clearly, the update time of this data structure is $\tilde{O}(n)$ and the construction time is $\tilde{O}(n_0)$. Thus, there exists a $(1 + \varepsilon)$ -approximate dynamic interval set cover data structure with $\tilde{O}(n^{\alpha_0} / \varepsilon^{1-\alpha_0})$ amortized update time for $\alpha_0 = 1$ and $\tilde{O}(n_0)$ construction time. Define $\alpha_i = \alpha_{i-1} / (1 + \alpha_{i-1})$ for $i \geq 1$. By applying Theorem 2 i times for a constant $i \geq 1$, we see the existence of a $(1 + \varepsilon)$ -approximate dynamic interval set cover data structure with $\tilde{O}(n^{\alpha_i} / \varepsilon^{1-\alpha_i})$ amortized update time and $\tilde{O}(n_0)$ construction time. One can easily verify that $\alpha_i = 1/(i + 1)$ for all $i \geq 0$. Therefore, for any constant $\alpha > 0$, we have an index $i \geq 0$ satisfying $\alpha_i < \alpha$ and hence $\tilde{O}(n^{\alpha_i} / \varepsilon^{1-\alpha_i}) = O(n^\alpha / \varepsilon)$. We finally conclude the following.

Theorem 5. *For a given approximation factor $\varepsilon > 0$ and any constant $\alpha > 0$, there exists a $(1 + \varepsilon)$ -approximate dynamic interval set cover data structure \mathcal{D} with $O(n^\alpha / \varepsilon)$ amortized update time and $\tilde{O}(n_0)$ construction time.*

4.2 2D set cover and hitting set for quadrants and unit squares

In this section, we present our bootstrapping framework for 2D dynamic set cover and hitting set. Our framework works for quadrants and unit squares.

We first show that dynamic unit-square set cover, dynamic unit-square hitting set, and dynamic quadrant hitting set can all be reduced to dynamic quadrant set cover.

Lemma 6. *Suppose there exists a c -approximate dynamic quadrant set cover data structure with $f(n)$ amortized update time and $\tilde{O}(n_0)$ construction time, where f is an increasing function. Then there exist $O(c)$ -approximate dynamic unit-square set cover, dynamic unit-square hitting set, and dynamic quadrant hitting set data structures with $\tilde{O}(f(n))$ amortized update time and $\tilde{O}(n_0)$ construction time.*

Now it suffices to consider dynamic quadrant set cover. In order to do bootstrapping, we need an output-sensitive algorithm for quadrant set cover, analog to the one in Lemma 1 for intervals. To design such an algorithm is considerably more difficult compared to the 1D case, and we defer it to Section 4.2.2. Before this, let us first discuss the bootstrapping procedure, assuming the existence of a μ -approximate output-sensitive algorithm for quadrant set cover.

4.2.1 Bootstrapping

We prove the following bootstrapping theorem, which is the technical heart of our result.

Theorem 7. Assume quadrant set cover admits a μ -approximate output-sensitive algorithm for some constant $\mu \geq 1$. Then we have the following result.

(*) Let $\alpha \in [0, 1]$ be a number. If there exists a $(\mu + \varepsilon)$ -approximate dynamic quadrant set cover data structure \mathcal{D}_{old} with $\tilde{O}(n^\alpha/\varepsilon^{1-\alpha})$ amortized update time and $\tilde{O}(n_0)$ construction time for any $\varepsilon > 0$, then there exists a $(\mu + \varepsilon)$ -approximate dynamic quadrant set cover data structure \mathcal{D}_{new} with $\tilde{O}(n^{\alpha'}/\varepsilon^{1-\alpha'})$ amortized update time and $\tilde{O}(n_0)$ construction time for any $\varepsilon > 0$, where $\alpha' = 2\alpha/(1+2\alpha)$. Here n (resp., n_0) denotes the size of the current (resp., initial) problem instance.

Assuming the existence of \mathcal{D}_{old} as in the theorem, we are going to design the improved data structure \mathcal{D}_{new} . Let (S, \mathcal{Q}) be a dynamic quadrant set cover instance. As before, we denote by n (resp., n_0) the size of the current (resp., initial) (S, \mathcal{Q}) .

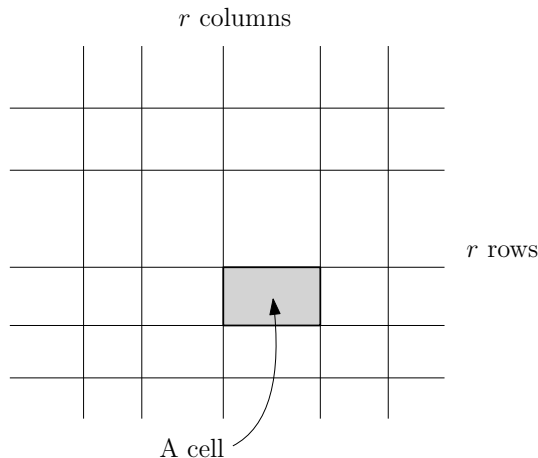


Figure 1: The $r \times r$ grid. Note that the cells may have different sizes.




Figure 2: A quadrant Q that left intersects $\square_{i,j}$.

The construction of \mathcal{D}_{new} . Initially, $|S| + |\mathcal{Q}| = n_0$. Essentially, our data structure \mathcal{D}_{new} consists of two parts. The first part is the data structure \mathcal{A} required for the μ -approximate output-sensitive algorithm. The second part is a family of \mathcal{D}_{old} data structures defined as follows. Let f be a function to be determined shortly. We use an orthogonal grid to partition the plane \mathbb{R}^2 into $r \times r$ cells for $r = \lceil n_0/f(n_0, \varepsilon) \rceil$ such that each row (resp., column) of the grid contains $O(f(n_0, \varepsilon))$ points in S and $O(f(n_0, \varepsilon))$ vertices of the quadrants in \mathcal{Q} (see Figure 1 for an illustration). Denote by $\square_{i,j}$ the cell in the i -th row and j -th column. Define $S_{i,j} = S \cap \square_{i,j}$. Also, we need to define a sub-collection $\mathcal{Q}_{i,j} \subseteq \mathcal{Q}$. Recall that in the 1D case, we define \mathcal{I}_i as the sub-collection of intervals in \mathcal{I} that partially intersect the portion J_i . However, for technical reasons, here we cannot simply define $\mathcal{Q}_{i,j}$ as the sub-collection of quadrants in \mathcal{Q} that partially intersects

$\square_{i,j}$. Instead, we define $\mathcal{Q}_{i,j}$ as follows. We include in $\mathcal{Q}_{i,j}$ all the quadrants in \mathcal{Q} whose vertices lie in $\square_{i,j}$. Besides, we also include in $\mathcal{Q}_{i,j}$ the following (at most) four *special* quadrants. We say a quadrant Q *left intersects* $\square_{i,j}$ if Q partially intersects $\square_{i,j}$ and contains the left edge of $\square_{i,j}$ (see Figure 2 for an illustration); similarly, we define “right intersects”, “top intersects”, and “bottom intersects”. Among a collection of quadrants, the *leftmost/rightmost/topmost/bottommost* quadrant refers to the quadrant whose vertex is the leftmost/rightmost/topmost/bottommost. We include in $\mathcal{Q}_{i,j}$ the rightmost quadrant in \mathcal{Q} that left intersects $\square_{i,j}$, the leftmost quadrant in \mathcal{Q} that right intersects $\square_{i,j}$, the bottommost quadrant in \mathcal{Q} that top intersects $\square_{i,j}$, and the topmost quadrant in \mathcal{Q} that bottom intersects $\square_{i,j}$ (if these quadrants exist). When the instance (S, \mathcal{Q}) is updated, the grid keeps unchanged, but the $S_{i,j}$ ’s and $\mathcal{Q}_{i,j}$ ’s change along with S and \mathcal{Q} . We view each $(S_{i,j}, \mathcal{Q}_{i,j})$ as a dynamic quadrant set cover instance, and let $\mathcal{D}_{\text{old}}^{(i,j)}$ be the data structure \mathcal{D}_{old} built on $(S_{i,j}, \mathcal{Q}_{i,j})$ for the approximation factor $\tilde{\varepsilon} = \varepsilon/2$. The second part of \mathcal{D}_{new} consists of the data structures $\mathcal{D}_{\text{old}}^{(i,j)}$ for $i, j \in \{1, \dots, r\}$.

Update and reconstruction. After each operation on (S, \mathcal{Q}) , we update the data structure \mathcal{A} . Also, if some $(S_{i,j}, \mathcal{Q}_{i,j})$ changes, we update the data structure $\mathcal{D}_{\text{old}}^{(i,j)}$. Note that an operation on S changes exactly one $S_{i,j}$, and an operation on \mathcal{Q} may only change the $\mathcal{Q}_{i,j}$ ’s in one row and one column (specifically, if the vertex of the inserted/deleted quadrant lies in $\square_{i,j}$, then only $\mathcal{Q}_{i,1}, \dots, \mathcal{Q}_{i,r}, \mathcal{Q}_{1,j}, \dots, \mathcal{Q}_{r,j}$ may change). Thus, we in fact only need to update the $\mathcal{D}_{\text{old}}^{(i,j)}$ ’s in one row and one column. Besides the update, we also reconstruct the entire data structure \mathcal{D}_{new} periodically, where the (first) reconstruction happens after processing $f(n_0, \varepsilon)$ operations. This part is totally the same as in our 1D data structure.

Maintaining a solution. We now discuss how to maintain a $(\mu + \varepsilon)$ -approximate optimal set cover $\mathcal{Q}_{\text{appx}}$ for (S, \mathcal{Q}) . Let opt denote the optimum of the current (S, \mathcal{Q}) . Set $\delta = \min\{(8\mu + 4\varepsilon + 2) \cdot r^2/\varepsilon, n\}$. If $\text{opt} \leq \delta$, then the output-sensitive algorithm can compute a μ -approximate optimal set cover for (S, \mathcal{Q}) in $\tilde{O}(\mu\delta)$ time. Thus, we simulate the output-sensitive algorithm within that amount of time. If the algorithm successfully computes a solution, we use it as our $\mathcal{Q}_{\text{appx}}$. Otherwise, we construct $\mathcal{Q}_{\text{appx}}$ as follows. We say the cell $\square_{i,j}$ is *coverable* if there exists $Q \in \mathcal{Q}$ that contains $\square_{i,j}$ and *uncoverable* otherwise. Let $P = \{(i, j) : \square_{i,j} \text{ is coverable}\}$ and $P' = \{(i, j) : \square_{i,j} \text{ is uncoverable}\}$. We try to use the quadrants in \mathcal{I} to “cover” all coverable cells. That is, for each $(i, j) \in P$, we find a quadrant in \mathcal{Q} that contains $\square_{i,j}$, and denote by \mathcal{Q}^* the set of all these quadrants. Then we consider the uncoverable cells. If for some $(i, j) \in P'$, the data structure $\mathcal{D}_{\text{old}}^{(i,j)}$ tells us that the instance $(S_{i,j}, \mathcal{Q}_{i,j})$ has no set cover, then we immediately make a no-solution decision, i.e., decide that the current (S, \mathcal{Q}) has no feasible set cover, and continue to the next operation. Otherwise, for each $(i, j) \in P'$, the data structure $\mathcal{D}_{\text{old}}^{(i,j)}$ maintains a $(\mu + \tilde{\varepsilon})$ -approximate optimal set cover $\mathcal{Q}_{i,j}^*$ for $(S_{i,j}, \mathcal{Q}_{i,j})$. We then define $\mathcal{Q}_{\text{appx}} = \mathcal{Q}^* \sqcup \left(\bigsqcup_{(i,j) \in P'} \mathcal{Q}_{i,j}^* \right)$.

We will see later that $\mathcal{Q}_{\text{appx}}$ is always a $(\mu + \varepsilon)$ -approximate optimal set cover for (S, \mathcal{Q}) . Before this, let us briefly discuss how to store $\mathcal{Q}_{\text{appx}}$ to support the desired queries. If $\mathcal{Q}_{\text{appx}}$ is computed by the output-sensitive algorithm, then we have all the elements of $\mathcal{Q}_{\text{appx}}$ in hand and can easily store them in a data structure to support the queries. Otherwise, $\mathcal{Q}_{\text{appx}}$ is defined as the disjoint union of \mathcal{Q}^* and $\mathcal{Q}_{i,j}^*$ ’s. In this case, the size and reporting queries can be handled in the same way as that in the 1D problem, by taking advantage of the fact that $\mathcal{Q}_{i,j}^*$ is maintained in $\mathcal{D}_{\text{old}}^{(i,j)}$. However, the situation for the membership query is more complicated, because now a quadrant in \mathcal{Q} may belong to many $\mathcal{Q}_{i,j}^*$ ’s. This issue can be handled by collecting all special quadrants in $\mathcal{Q}_{\text{appx}}$ and building on them a data structure that supports the membership query. We defer the detailed discussion to Appendix C.

Correctness. We now prove the correctness of our data structure \mathcal{D}_{new} . First, we show that the no-solution decision made by our data structure is correct.

Lemma 8. \mathcal{D}_{new} makes a no-solution decision iff the current (S, \mathcal{Q}) has no set cover.

Next, we show that the solution $\mathcal{Q}_{\text{appx}}$ maintained by \mathcal{D}_{new} is truly a $(\mu + \varepsilon)$ -approximate optimal set cover for (S, \mathcal{Q}) . If $\mathcal{Q}_{\text{appx}}$ is computed by the output-sensitive algorithm, then it is a μ -approximate optimal set

cover for (S, \mathcal{Q}) . Otherwise, $\text{opt} > \delta = \min\{(8\mu + 4\varepsilon + 2) \cdot r^2/\varepsilon, n\}$, i.e., either $\text{opt} > (8\mu + 4\varepsilon + 2) \cdot r^2/\varepsilon$ or $\text{opt} > n$. If $\text{opt} > n$, then (S, \mathcal{Q}) has no set cover (i.e., $\text{opt} = \infty$) and \mathcal{D}_{new} makes a no-solution decision by Lemma 8. So assume $\text{opt} > (8\mu + 4\varepsilon + 2)r^2/\varepsilon$. In this case, $\mathcal{Q}_{\text{appx}} = \mathcal{Q}^* \sqcup (\bigsqcup_{(i,j) \in P'} \mathcal{Q}_{i,j}^*)$. For each $(i, j) \in P'$, let $\text{opt}_{i,j}$ be the optimum of $(S_{i,j}, \mathcal{Q}_{i,j})$. Then we have $|\mathcal{Q}_{i,j}^*| \leq (\mu + \tilde{\varepsilon}) \cdot \text{opt}_{i,j}$ for all $(i, j) \in P'$ where $\tilde{\varepsilon} = \varepsilon/2$. Since $|\mathcal{Q}^*| \leq r^2$, we have

$$|\mathcal{Q}_{\text{appx}}| = |\mathcal{Q}^*| + \sum_{(i,j) \in P'} |\mathcal{Q}_{i,j}^*| \leq r^2 + \left(\mu + \frac{\varepsilon}{2}\right) \sum_{(i,j) \in P'} \text{opt}_{i,j}. \quad (4)$$

Let $\mathcal{Q}'_{i,j} \subseteq \mathcal{Q}_{i,j}$ consist of the non-special quadrants, i.e., those whose vertices are in $\square_{i,j}$.

Lemma 9. *We have $|\mathcal{Q}_{\text{opt}} \cap \mathcal{Q}'_{i,j}| + 4 \geq \text{opt}_{i,j}$ for all $(i, j) \in P'$, and in particular,*

$$\text{opt} + 4r^2 = |\mathcal{Q}_{\text{opt}}| + 4r^2 \geq \sum_{(i,j) \in P'} \text{opt}_{i,j}. \quad (5)$$

Using Equations 4 and 5, we deduce that

$$\begin{aligned} |\mathcal{Q}_{\text{appx}}| &\leq r^2 + \left(\mu + \frac{\varepsilon}{2}\right) \sum_{(i,j) \in P'} \text{opt}_{i,j} \\ &\leq r^2 + \left(\mu + \frac{\varepsilon}{2}\right) (\text{opt} + 4r^2) \\ &\leq (4\mu + 2\varepsilon + 1) \cdot r^2 + \left(\mu + \frac{\varepsilon}{2}\right) \cdot \text{opt} \\ &< \frac{\varepsilon}{2} \cdot \text{opt} + \left(\mu + \frac{\varepsilon}{2}\right) \cdot \text{opt} = (\mu + \varepsilon) \cdot \text{opt}, \end{aligned}$$

where the last inequality follows from the fact that $\text{opt} > (8\mu + 4\varepsilon + 2) \cdot r^2/\varepsilon$.

Time analysis. We briefly discuss the update and construction time of \mathcal{D}_{new} ; a detailed analysis can be found in Appendix C. It suffices to consider the first period (i.e., the period before the first reconstruction). We first observe the following fact.

Lemma 10. *At any time in the first period, we have $\sum_{k=1}^r (|S_{i,k}| + |\mathcal{Q}_{i,k}|) = O(f(n_0, \varepsilon) + r)$ for all $i \in \{1, \dots, r\}$ and $\sum_{k=1}^r (|S_{k,j}| + |\mathcal{Q}_{k,j}|) = O(f(n_0, \varepsilon) + r)$ for all $j \in \{1, \dots, r\}$.*

The above lemma implies that the sum of the sizes of all $(S_{i,j}, \mathcal{Q}_{i,j})$ is $O(n_0 + r^2)$ at any time in the first period. Therefore, constructing \mathcal{D}_{new} can be easily done in $\tilde{O}(n_0 + r^2)$ time. The update time of \mathcal{D}_{new} consists of the time for reconstruction, the time for updating \mathcal{A} and $\mathcal{D}_{\text{old}}^{(i,j)}$'s, and the time for maintaining the solution. Using almost the same analysis as in the 1D problem, we can show that the reconstruction takes $\tilde{O}(r + r^2/f(n_0, \varepsilon))$ amortized time and maintaining the solution can be done in $\tilde{O}(\delta + r^2)$ time, with a careful implementation. The time for updating the \mathcal{D}_{old} data structures requires a different analysis. Let $m_{i,j}$ denote the current size of $(S_{i,j}, \mathcal{Q}_{i,j})$. As argued before, we in fact only need to update the \mathcal{D}_{old} data structures in one row and one column (say the i -th row and j -th column). Hence, updating the \mathcal{D}_{old} data structures takes $\tilde{O}(\sum_{k=1}^r m_{i,k}^\alpha/\varepsilon^{1-\alpha} + \sum_{k=1}^r m_{k,j}^\alpha/\varepsilon^{1-\alpha})$ amortized time. Lemma 10 implies that $\sum_{k=1}^r m_{i,k} = O(f(n_0, \varepsilon) + r)$ and $\sum_{k=1}^r m_{k,j} = O(f(n_0, \varepsilon) + r)$. Since $\alpha \leq 1$, by Hölder's inequality and Lemma 10,

$$\sum_{k=1}^r m_{i,k}^\alpha \leq \left(\frac{\sum_{k=1}^r m_{i,k}}{r}\right)^\alpha \cdot r = O(r^{1-\alpha} \cdot (f(n_0, \varepsilon) + r)^\alpha)$$

and similarly $\sum_{k=1}^r m_{k,j}^\alpha = O(r^{1-\alpha} \cdot (f(n_0, \varepsilon) + r)^\alpha)$. It follows that updating the \mathcal{D}_{old} data structures takes $\tilde{O}(r^{1-\alpha} \cdot (f(n_0, \varepsilon) + r)^\alpha/\varepsilon^{1-\alpha})$ amortized time. In total, the amortized update time of \mathcal{D}_{new} (during the first period) is $\tilde{O}(r^{1-\alpha} \cdot (f(n_0, \varepsilon) + r)^\alpha + \delta + r^2)$. If we set $f(n, \varepsilon) = \min\{n^{1-\alpha'}/(\sqrt{\varepsilon})^{\alpha'}, n/2\}$ where α' is as defined in Theorem 7, a careful calculation (see Appendix C) shows that the amortized update time becomes $\tilde{O}(n^{\alpha'}/\varepsilon^{1-\alpha'})$.

4.2.2 An output-sensitive quadrant set cover algorithm

We propose an $O(1)$ -approximate output-sensitive algorithm for quadrant set cover, which is needed for applying Theorem 7. Let (S, \mathcal{Q}) be a quadrant set cover instance of size n , and opt be its optimum. Our goal is to compute an $O(1)$ -approximate optimal set cover for (S, \mathcal{Q}) in $\tilde{O}(\text{opt})$ time, using some basic data structure built on (S, \mathcal{Q}) .

For simplicity, let us assume that (S, \mathcal{Q}) has a set cover; how to handle the no-solution case is discussed in Appendix D.1. There are four types of quadrants in \mathcal{Q} , southeast, southwest, northeast, northwest; we denote by $\mathcal{Q}^{\text{SE}}, \mathcal{Q}^{\text{SW}}, \mathcal{Q}^{\text{NE}}, \mathcal{Q}^{\text{NW}} \subseteq \mathcal{Q}$ the sub-collections of these types of quadrants, respectively. Let U^{SE} denote the union of the quadrants in \mathcal{Q}^{SE} , and define $U^{\text{SW}}, U^{\text{NE}}, U^{\text{NW}}$ similarly. Since (S, \mathcal{Q}) has a set cover, we have $S = (S \cap U^{\text{SE}}) \cup (S \cap U^{\text{SW}}) \cup (S \cap U^{\text{NE}}) \cup (S \cap U^{\text{NW}})$. Therefore, if we can compute $O(1)$ -approximate optimal set covers for $(S \cap U^{\text{SE}}, \mathcal{Q})$, $(S \cap U^{\text{SW}}, \mathcal{Q})$, $(S \cap U^{\text{NE}}, \mathcal{Q})$, and $(S \cap U^{\text{NW}}, \mathcal{Q})$, then the union of these four set covers is an $O(1)$ -approximate optimal set cover for (S, \mathcal{Q}) .

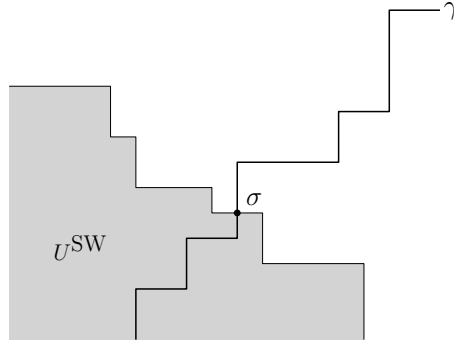


Figure 3: Illustrating the curve γ and the point σ .

With this observation, it now suffices to show how to compute an $O(1)$ -approximate optimal set cover for $(S \cap U^{\text{SE}}, \mathcal{Q})$ in $\tilde{O}(\text{opt}^{\text{SE}})$ time, where opt^{SE} is the optimum of $(S \cap U^{\text{SE}}, \mathcal{Q})$. The main challenge is to guarantee the running time and approximation ratio simultaneously. We begin by introducing some notation. Let γ denote the boundary of U^{SE} , which is an orthogonal staircase curve from bottom-left to top-right. If $\gamma \cap U^{\text{SW}} \neq \emptyset$, then $\gamma \cap U^{\text{SW}}$ is a connected portion of γ that contains the bottom-left end of γ . Define σ as the “endpoint” of $\gamma \cap U^{\text{SW}}$, i.e., the point on $\gamma \cap U^{\text{SW}}$ that is closest to the top-right end of γ . See Figure 3 for an illustration. If $\gamma \cap U^{\text{SW}} = \emptyset$, we define σ as the bottom-left end of γ (which is a point whose y -coordinate equals to $-\infty$). For a number $\tilde{y} \in \mathbb{R}$, we define $\phi(\tilde{y})$ as the *leftmost* point in $S \cap U^{\text{SE}}$ whose y -coordinate is greater than \tilde{y} ; we say $\phi(\tilde{y})$ does not exist if no point in $S \cap U^{\text{SE}}$ has y -coordinate greater than \tilde{y} . For a point $a \in \mathbb{R}^2$ and a collection \mathcal{P} of quadrants, we define $\Phi_{\rightarrow}(a, \mathcal{P})$ and $\Phi_{\uparrow}(a, \mathcal{P})$ as the rightmost and topmost quadrants in \mathcal{P} that contains a , respectively. For a quadrant Q , we denote by $x(Q)$ and $y(Q)$ the x - and y -coordinates of the vertex of Q , respectively.

To get some intuition, let us consider a very simple case, where \mathcal{Q} only consists of southeast quadrants. In this case, one can compute an optimal set cover for $(S \cap U^{\text{SE}}, \mathcal{Q})$ using a greedy algorithm similar to the 1D interval set cover algorithm: repeatedly pick the leftmost uncovered point in $S \cap U^{\text{SE}}$ and cover it using the topmost (southeast) quadrant in \mathcal{Q} . Using the notations defined above, we can describe this algorithm as follows. Set $\mathcal{Q}_{\text{ans}} \leftarrow \emptyset$ and $\tilde{y} \leftarrow -\infty$ initially, and repeatedly do $a \leftarrow \phi(\tilde{y})$, $Q \leftarrow \Phi_{\uparrow}(\sigma, \mathcal{Q}^{\text{SE}})$, $\mathcal{Q}_{\text{ans}} \leftarrow \mathcal{Q}_{\text{ans}} \cup \{Q\}$, $\tilde{y} \leftarrow y(Q)$ until $\phi(\tilde{y})$ does not exist. Eventually, \mathcal{Q}_{ans} is the set cover we want.

Now we try to extend this algorithm to the general case. However, the situation here becomes much more complicated, since we may have three other types of quadrants in \mathcal{Q} , which have to be carefully dealt with in order to guarantee the correctness. But the intuition remains the same: we still construct the solution in a greedy manner. The following procedure describes our algorithm.

1. $\mathcal{Q}_{\text{ans}} \leftarrow \emptyset$. $\tilde{y} \leftarrow -\infty$. If $\phi(\tilde{y})$ does not exist, then go to Step 6.

2. $\mathcal{Q}_{\text{ans}} \leftarrow \{\Phi_{\rightarrow}(\sigma, \mathcal{Q}^{\text{SW}}), \Phi_{\uparrow}(\sigma, \mathcal{Q}^{\text{SE}})\}$. $\tilde{y} \leftarrow y(\Phi_{\uparrow}(\sigma, \mathcal{Q}^{\text{SE}}))$. If $\phi(\tilde{y})$ exists, then $a \leftarrow \phi(\tilde{y})$, else go to Step 6.
3. If $a \in U^{\text{NE}}$, then $\mathcal{Q}_{\text{ans}} \leftarrow \mathcal{Q}_{\text{ans}} \cup \{\Phi_{\uparrow}(a, \mathcal{Q}^{\text{NE}}), \Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}})\}$ and go to Step 6.
4. If $a \in U^{\text{NW}}$, then $\mathcal{Q}_{\text{ans}} \leftarrow \mathcal{Q}_{\text{ans}} \cup \{\Phi_{\rightarrow}(a, \mathcal{Q}^{\text{NW}}), \Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}})\}$ and $Q \leftarrow \Phi_{\uparrow}(v, \mathcal{Q}^{\text{SE}})$ where v is the vertex of $\Phi_{\rightarrow}(a, \mathcal{Q}^{\text{NW}})$, otherwise $Q \leftarrow \Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}})$.
5. $\mathcal{Q}_{\text{ans}} \leftarrow \mathcal{Q}_{\text{ans}} \cup \{Q\}$. $\tilde{y} \leftarrow y(Q)$. If $\phi(\tilde{y})$ exists, then $a \leftarrow \phi(\tilde{y})$ and go to Step 3.
6. Output \mathcal{Q}_{ans} .

The following lemma proves the correctness of our algorithm.

Lemma 11. \mathcal{Q}_{ans} covers all points in $S \cap U^{\text{SE}}$, and $|\mathcal{Q}_{\text{ans}}| = O(\text{opt}^{\text{SE}})$.

The remaining task is to show how to perform our algorithm in $\tilde{O}(\text{opt}^{\text{SE}})$ time using basic data structures. It is clear that our algorithm terminates in $O(\text{opt}^{\text{SE}})$ steps, since we include at least one quadrant to \mathcal{Q}_{ans} in each iteration of the loop Step 3–5 and eventually $|\mathcal{Q}_{\text{ans}}| = O(\text{opt}^{\text{SE}})$ by Lemma 11. Thus, it suffices to show that each step can be done in $\tilde{O}(1)$ time. In every step of our algorithm, all work can be done in constant time except the tasks of computing the point σ , testing whether $a \in U^{\text{NE}}$ and $a \in U^{\text{NW}}$ for a given point a , computing the quadrants $\Phi_{\rightarrow}(a, \mathcal{Q}^{\text{SW}}), \Phi_{\rightarrow}(a, \mathcal{Q}^{\text{NW}}), \Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}}), \Phi_{\uparrow}(a, \mathcal{Q}^{\text{NE}})$ for a given point a , and computing $\phi(\tilde{y})$ for a given number \tilde{y} . All these tasks except the computation of $\phi(\cdot)$ can be easily done in $\tilde{O}(1)$ time by storing the quadrants in binary search trees. To compute $\phi(\cdot)$ in $\tilde{O}(1)$ time is more difficult, and we achieve this by properly using range trees built on both S and \mathcal{Q}^{SE} . The details are presented in Appendix D.2.

Using the above algorithm, we can compute $O(1)$ -approximate optimal set covers for $(S \cap U^{\text{SE}}, \mathcal{Q})$, $(S \cap U^{\text{SW}}, \mathcal{Q})$, $(S \cap U^{\text{NE}}, \mathcal{Q})$, and $(S \cap U^{\text{NW}}, \mathcal{Q})$. As argued before, the union of these four set covers, denoted by \mathcal{Q}^* , is an $O(1)$ -approximate optimal set covers for (S, \mathcal{Q}) .

Theorem 12. *Quadrant set cover admits an $O(1)$ -approximate output-sensitive algorithm.*

4.2.3 Putting everything together

With the bootstrapping theorem in hand, we are now able to design our dynamic quadrant set cover data structure. Again, the starting point is a “trivial” data structure which uses the output-sensitive algorithm of Theorem 12 to re-compute an optimal quadrant set cover after each update. Clearly, the update time of this data structure is $\tilde{O}(n)$ and the construction time is $\tilde{O}(n_0)$. Let $\mu = O(1)$ be the approximation ratio of the output-sensitive algorithm. The trivial data structure implies the existence of a $(\mu + \varepsilon)$ -approximate dynamic quadrant set cover data structure with $\tilde{O}(n^{\alpha_0}/\varepsilon^{1-\alpha_0})$ amortized update time for $\alpha_0 = 1$ and $\tilde{O}(n_0)$ construction time. Define $\alpha_i = 2\alpha_{i-1}/(1 + 2\alpha_{i-1})$ for $i \geq 1$. By applying Theorem 7 i times for a constant $i \geq 1$, we see the existence of a $(\mu + \varepsilon)$ -approximate dynamic quadrant set cover data structure with $\tilde{O}(n^{\alpha_i}/\varepsilon^{1-\alpha_i})$ amortized update time and $\tilde{O}(n_0)$ construction time. One can easily verify that $\alpha_i = 2^i/(2^{i+1} - 1)$ for all $i \geq 0$. Therefore, for any constant $\alpha > 0$, we have an index $i \geq 0$ satisfying $\alpha_i < 1/2 + \alpha$ and hence $\tilde{O}(n^{\alpha_i}/\varepsilon^{1-\alpha_i}) = O(n^{1/2+\alpha}/\varepsilon)$. Setting ε to be any constant, we finally conclude the following.

Theorem 13. *For any constant $\alpha > 0$, there exists an $O(1)$ -approximate dynamic quadrant set cover data structure with $O(n^{1/2+\alpha})$ amortized update time and $\tilde{O}(n_0)$ construction time.*

By the reduction of Lemma 6, we have the following corollary.

Corollary 14. *For any constant $\alpha > 0$, there exist $O(1)$ -approximate dynamic unit-square set cover, dynamic unit-square hitting set, and dynamic quadrant hitting set data structures with $O(n^{1/2+\alpha})$ amortized update time and $\tilde{O}(n_0)$ construction time.*

5 Second framework: Local modification

In this section, we give our second framework for dynamic geometric set cover and hitting set, which results in poly-logarithmic update time data structures for dynamic interval hitting set, as well as dynamic quadrant and unit-square set cover (resp., hitting set) in the partially dynamic setting. This framework applies to *stable* dynamic geometric set cover and hitting set problems, in which each operation cannot change the optimum of the problem instance significantly. One can easily see that any dynamic set cover (resp., hitting set) problem in the partially dynamic setting is stable. Interestingly, we will see that the dynamic interval hitting set problem, even in the fully dynamic setting, is stable.

The basic idea behind this framework is quite simple, and we call it *local modification*. Namely, at the beginning, we compute a good solution S^* for the problem instance using an output-sensitive algorithm, and after each operation, we slightly modify S^* , by adding/removing a constant number of elements, to guarantee that S^* is a feasible solution for the current instance. After processing many operations, we re-compute a new S^* and repeat this procedure. The stability of the problem can guarantee that S^* is always a good approximation for the optimal solution.

Next, we formally define the notion of *stability*. The *quasi-optimum* of a set cover instance (S, \mathcal{R}) is the size of an optimal set cover for (S', \mathcal{R}) , where $S' \subseteq S$ consists of the points that can be covered by \mathcal{R} . The quasi-optimum of a hitting set instance is defined in a similar way. For a dynamic set cover (resp., hitting set) instance Π , we use $\text{opt}_i(\Pi)$ to denote the quasi-optimum of Π at the time i (i.e., after the i -th operation).

Definition 15. A dynamic set cover (resp., hitting set) instance Π is **stable** if $|\text{opt}_i(\Pi) - \text{opt}_{i+1}(\Pi)| \leq 1$ for all $i \geq 0$. A dynamic set cover (resp., hitting set) problem is **stable** if any instance of the problem is stable.

Lemma 16. Any dynamic set cover (resp., hitting set) problem in the partially dynamic setting is stable. The dynamic interval hitting set problem is stable.

5.1 Dynamic interval hitting set

We show how to use the idea of local modification to obtain a $(1 + \varepsilon)$ -approximate dynamic interval hitting set data structure \mathcal{D} with $\tilde{O}(1/\varepsilon)$ amortized update time.

Similarly to interval set cover, the interval hitting set problem can also be solved using a simple greedy algorithm, which can be performed in $\tilde{O}(\text{opt})$ time if we store (properly) the points and intervals in some basic data structure (e.g., binary search trees). We omit the proof of the following lemma, as it is almost the same as the proof of Lemma 1.

Lemma 17. The interval hitting set problem yields an exact output-sensitive algorithm.

Let (S, \mathcal{I}) be a dynamic interval hitting set instance. We denote by n (resp., n_0) the size of the current (resp., initial) (S, \mathcal{I}) , and by opt the optimum of the current (S, \mathcal{I}) .

The data structure. Our data structure \mathcal{D} basically consists of three parts. The first part is the basic data structure \mathcal{A} built on (S, \mathcal{I}) required for the output-sensitive algorithm. The second part is the following dynamic data structure \mathcal{B} which can tell us whether the current instance has a hitting set.

Lemma 18. There exists a dynamic data structure \mathcal{B} built on (S, \mathcal{I}) with $\tilde{O}(1)$ update time and $\tilde{O}(n_0)$ construction time that can indicate whether the current (S, \mathcal{I}) has a hitting set or not.

The third part consists of the following two basic data structures \mathcal{C}_1 and \mathcal{C}_2 .

Lemma 19. One can store S in a basic data structure \mathcal{C}_1 such that given a point a , the rightmost (resp., leftmost) point in S to the left (resp., right) of a can be reported in $\tilde{O}(1)$ time (if they exist).

Lemma 20. One can store S in a basic data structure \mathcal{C}_2 such that given an interval I , a point $a \in S$ that is contained in I can be reported in $\tilde{O}(1)$ time (if it exists).

Maintaining a solution. Let (S_0, \mathcal{I}_0) be the initial (S, \mathcal{I}) , and define $\mathcal{I}'_0 \subseteq \mathcal{I}_0$ as the sub-collection consisting of all intervals that are hit by S_0 . In the construction of \mathcal{D} , after building the data structures \mathcal{A} and \mathcal{B} , we compute an optimal hitting set S^* for (S_0, \mathcal{I}'_0) . Set $\text{cnt} = 0$ and $\text{opt}^\sim = |S^*|$ initially. The data structure \mathcal{D} handles the operations as follows. After each operation, \mathcal{D} first increment cnt by 1 and query the data structure \mathcal{B} to see whether the current (S, \mathcal{I}) has a hitting set. Whenever $\text{cnt} \geq \varepsilon \cdot \text{opt}^\sim / (2 + \varepsilon)$ and the current (S, \mathcal{I}) has a hitting set, we re-compute an optimal hitting set S^* for the current (S, \mathcal{I}) using the output-sensitive algorithm and reset $\text{cnt} = 0$ and $\text{opt}^\sim = |S^*|$. Otherwise, we apply local modification on S^* (i.e., “slightly” modify S^*) as follows.

- If the operation inserts a point a to S , then include a in S^* . If the operation deletes a point a from S , then remove a from S^* (if $a \in S^*$) and include in S^* the rightmost point in S to the left of a and the leftmost point in S to the right of a (if they exist), which can be found using the data structure \mathcal{C}_1 of Lemma 19.
- If the operation inserts an interval I to \mathcal{I} , then include in S^* an arbitrary point in S that hits I (if it exists), which can be found using the data structure \mathcal{C}_2 of Lemma 20. If the operation deletes an interval from \mathcal{I} , keep S^* unchanged.

Note that we do local modification on S^* no matter whether the current (S, \mathcal{I}) has a hitting set or not. After this, our data structure \mathcal{D} uses S^* as the solution for the current (S, \mathcal{I}) , if the current (S, \mathcal{I}) has a hitting set. In order to support the desired queries for S^* , we can simply store S^* in a (dynamic) binary search tree. When we re-compute S^* , we re-build the binary search tree that stores S^* . For each local modification on S^* , we do insertion/deletion on the binary search tree according to the change of S^* .

Correctness. To see the correctness of our data structure, we have to show that S^* is always a $(1 + \varepsilon)$ -approximate optimal hitting set for the current (S, \mathcal{I}) , whenever (S, \mathcal{I}) has a hitting set. To this end, we observe some simple facts. First, the counter cnt is exactly the number of the operations processed after the last re-computation of S^* , and the number opt^\sim is the optimum at the last re-computation of S^* . Second, after each local modification, the size of S^* either increases by 1 or keeps unchanged. Based on these two facts and the stability of dynamic interval hitting set (Lemma 16), we can easily bound the size of S^* .

Lemma 21. *We have $|S^*| \leq (1 + \varepsilon) \cdot \text{opt}$ at any time.*

With the above lemma in hand, it now suffices to show that S^* is a feasible hitting set.

Lemma 22. *Whenever the current (S, \mathcal{I}) has a hitting set, S^* is a hitting set for (S, \mathcal{I}) .*

Time analysis. The construction time of \mathcal{D} is clearly $\tilde{O}(n_0)$. Specifically, when constructing \mathcal{D} , we need to build the data structures $\mathcal{A}, \mathcal{B}, \mathcal{C}_1, \mathcal{C}_2$, all of which can be built in $\tilde{O}(n_0)$ time. Also, we need to compute an optimal hitting set for (S_0, \mathcal{I}'_0) using the output-sensitive algorithm, which can also be done in $\tilde{O}(n_0)$ time. We show that the amortized update time of \mathcal{D} is $\tilde{O}(1/\varepsilon)$. After each operation, we need to update the data structures $\mathcal{A}, \mathcal{B}, \mathcal{C}_1, \mathcal{C}_2$, which takes $\tilde{O}(1)$ time. If we do local modification on S^* , then it takes $\tilde{O}(1)$ time. Indeed, in each local modification, there are $O(1)$ elements to be added to or removed from S^* . These elements can be found in $\tilde{O}(1)$ time using the data structures \mathcal{C}_1 and \mathcal{C}_2 . After these elements are found, we need to modify the binary search tree that stores S^* for supporting the desired queries, which can also be done in $\tilde{O}(1)$ time. In the case where we re-compute a new S^* , it takes $\tilde{O}(\text{opt})$ time. However, we can amortize this time over the operations in between the last re-computation of S^* and the current one. Let opt^\sim be the quasi-optimum of (S, \mathcal{I}) at the last re-computation. Suppose there are t operations in between the two re-computations. Then $t \geq \varepsilon \cdot \text{opt}^\sim / (2 + \varepsilon)$. By the stability of dynamic interval hitting set (Lemma 16), we have $\text{opt} \leq \text{opt}^\sim + t$, because opt is also the quasi-optimum of the current (S, \mathcal{I}) . It follows that

$$\frac{\text{opt}}{t} \leq \frac{\text{opt}^\sim + t}{t} \leq \frac{2 + \varepsilon}{\varepsilon} + 1 = O(1/\varepsilon).$$

Thus, the amortized time for the current re-computation is $\tilde{O}(1/\varepsilon)$. We conclude the following.

Theorem 23. *For a given approximation factor $\varepsilon > 0$, there is a $(1 + \varepsilon)$ -approximate dynamic interval hitting data structure \mathcal{D} with $\tilde{O}(1/\varepsilon)$ amortized update time and $\tilde{O}(n_0)$ construction time.*

5.2 Dynamic set cover and hitting set in the partially dynamic settings

The idea of local modification can also be applied to dynamic geometric set cover (resp., hitting set) problems in the partially dynamic setting (which are stable by Lemma 16), as long as we have an output-sensitive algorithm for the problem and an efficient dynamic data structure that can indicate whether the current instance has a feasible solution or not.

For an example, let us consider dynamic quadrant set cover in the partially dynamic setting. Let (S, \mathcal{Q}) be a dynamic quadrant set cover instance with only point updates. As before, we denote by n (resp., n_0) the size of the current (resp., initial) (S, \mathcal{Q}) , and by opt the optimum of the current (S, \mathcal{Q}) .

The data structure. Our data structure \mathcal{D} consists of three parts. The first part is the basic data structure \mathcal{A} required for performing the output-sensitive algorithm of Theorem 12. The second part is the following dynamic data structure \mathcal{B} which can tell us whether the current instance has a hitting set.

Lemma 24. *There exists a dynamic data structure \mathcal{B} built on (S, \mathcal{Q}) with $\tilde{O}(1)$ update time and $\tilde{O}(n_0)$ construction time that can indicate whether the current (S, \mathcal{Q}) has a set cover or not.*

The third part is a simple (static) data structure \mathcal{C} built on \mathcal{Q} that can report in $\tilde{O}(1)$ time, for a given point $a \in \mathbb{R}^2$, a quadrant in \mathcal{Q} that contains a (if it exists); this can be done using a standard orthogonal range-search data structure, e.g., range tree, whose construction time is $\tilde{O}(|\mathcal{Q}|)$.

Maintaining a solution. Let S_0 be the initial S , and define $S'_0 \subseteq S_0$ as the subset consisting of the points covered by \mathcal{Q} . At the beginning, we compute a μ -approximate optimal set cover \mathcal{Q}^* for (S'_0, \mathcal{Q}) , where μ is the approximation ratio of the output-sensitive algorithm. This can be done as follows. We first compute S'_0 by testing whether each point in S is covered by some range in \mathcal{Q} using the data structure \mathcal{C} , and then build on (S'_0, \mathcal{Q}) the basic data structure required for the output-sensitive algorithm and perform the output-sensitive algorithm on (S'_0, \mathcal{Q}) . Set $\text{cnt} = 0$ and $\text{opt}^\sim = |\mathcal{Q}^*|$. After each iteration, \mathcal{D} first increments cnt by 1 and checks whether the current (S, \mathcal{Q}) has a set cover or not using the data structure \mathcal{B} . Whenever $\text{cnt} \geq (\varepsilon/\mu) \cdot \text{opt}^\sim / (2 + \varepsilon)$ and the current (S, \mathcal{Q}) has a set cover, we re-compute a μ -approximate optimal set cover \mathcal{Q}^* for the current (S, \mathcal{Q}) using the output-sensitive algorithm, and reset $\text{cnt} = 0$ and $\text{opt}^\sim = |\mathcal{Q}^*|$. Otherwise, we apply local modification on \mathcal{Q}^* as follows. If the operation inserts a point a to S , then include in \mathcal{Q}^* an arbitrary range in \mathcal{Q} that contains a , which can be found using the data structure \mathcal{C} (if it exists). If the operation deletes a point a from S , then keep \mathcal{Q}^* unchanged. After this, our data structure \mathcal{D} uses \mathcal{Q}^* as the solution for the current (S, \mathcal{Q}) . As in the last section, we can store \mathcal{Q}^* in a (dynamic) binary search tree to support the desired queries (by fixing some order on the family of quadrants).

Correctness. Using the same argument as in the last section, we can show that $|\mathcal{Q}^*| \leq (\mu + \varepsilon) \cdot \text{opt}$ at any time. Furthermore, it is also easy to see that \mathcal{Q}^* is always a feasible set cover.

Lemma 25. *Whenever the current (S, \mathcal{Q}) has a set cover, \mathcal{Q}^* is a set cover for (S, \mathcal{Q}) .*

Time analysis. Using the same analysis as in the last section, we can show that the construction time of the data structure \mathcal{D} is $\tilde{O}(n)$, and the amortized update time is $\tilde{O}(1/\varepsilon)$. Setting ε to be any constant, we conclude the following.

Theorem 26. *There exists an $O(1)$ -approximate partially dynamic quadrant set cover data structure with $\tilde{O}(1)$ amortized update time and $\tilde{O}(n_0)$ construction time.*

Using the reduction of Lemma 6, we also have the following result. Note that although the reduction of Lemma 6 is described for the fully dynamic setting, it actually applies to reduce dynamic unit-square set cover (resp., hitting set) in the partially dynamic setting and dynamic quadrant hitting set in the partially dynamic setting to dynamic quadrant set cover data structure in the partially dynamic setting.

Corollary 27. *There exist an $O(1)$ -approximate partially dynamic unit-square set cover, partially dynamic unit-square hitting set, and partially dynamic quadrant hitting set data structures with $\tilde{O}(1)$ amortized update time and $\tilde{O}(n_0)$ construction time.*

References

- [1] Amir Abboud, Raghavendra Addanki, Fabrizio Grandoni, Debmalya Panigrahi, and Barna Saha. Dynamic set cover: improved algorithms and lower bounds. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 114–125. ACM, 2019.
- [2] Pankaj K Agarwal and Jiangwei Pan. Near-linear algorithms for geometric hitting sets and set covers. In *Proceedings of the thirtieth annual symposium on Computational geometry*, page 271. ACM, 2014.
- [3] Surender Baswana, Manoj Gupta, and Sandeep Sen. Fully dynamic maximal matching in $o(\log n)$ update time. *SIAM Journal on Computing*, 44(1):88–113, 2015.
- [4] Jon Louis Bentley. Decomposable searching problems. Technical report, Carnegie-Mellon University, Department of Computer Science, 1978.
- [5] Piotr Berman and Bhaskar DasGupta. Complexities of efficient solutions of rectilinear polygon cover problems. *Algorithmica*, 17(4):331–356, 1997.
- [6] Aaron Bernstein and Cliff Stein. Faster fully dynamic matchings with small approximation ratios. In *Proceedings of the twenty-seventh annual ACM-SIAM Symposium on Discrete Algorithms*, pages 692–711. Society for Industrial and Applied Mathematics, 2016.
- [7] Sayan Bhattacharya, Deeparnab Chakrabarty, and Monika Henzinger. Deterministic fully dynamic approximate vertex cover and fractional matching in $o(1)$ amortized update time. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 86–98. Springer, 2017.
- [8] Sayan Bhattacharya, Monika Henzinger, and Giuseppe F Italiano. Design of dynamic algorithms via primal-dual method. In *International Colloquium on Automata, Languages, and Programming*, pages 206–218. Springer, 2015.
- [9] Sayan Bhattacharya, Monika Henzinger, and Giuseppe F Italiano. Deterministic fully dynamic data structures for vertex cover and matching. *SIAM Journal on Computing*, 47(3):859–887, 2018.
- [10] Sayan Bhattacharya, Monika Henzinger, and Danupon Nanongkai. New deterministic approximation algorithms for fully dynamic matching. In *Proceedings of the forty-eighth annual ACM Symposium on Theory of Computing*, pages 398–411. ACM, 2016.
- [11] Vasek Chvatal. A greedy heuristic for the set-covering problem. *Mathematics of operations research*, 4(3):233–235, 1979.
- [12] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. MIT press, 2009.
- [13] Irit Dinur and David Steurer. Analytical approach to parallel repetition. In *Proceedings of the forty-sixth annual ACM Symposium on Theory of Computing*, pages 624–633. ACM, 2014.

- [14] Thomas Erlebach and Erik Jan Van Leeuwen. Ptas for weighted set cover on unit squares. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 166–177. Springer, 2010.
- [15] Shashidhara K Ganjugunte. *Geometric hitting sets and their variants*. PhD thesis, Duke University, 2011.
- [16] Anupam Gupta, Ravishankar Krishnaswamy, Amit Kumar, and Debmalya Panigrahi. Online and dynamic algorithms for set cover. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 537–550. ACM, 2017.
- [17] Manoj Gupta and Richard Peng. Fully dynamic (1+ ϵ)-approximate matchings. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 548–557. IEEE, 2013.
- [18] Juris Hartmanis. Computers and intractability: a guide to the theory of np-completeness. *Siam Review*, 24(1):90, 1982.
- [19] David S Johnson. Approximation algorithms for combinatorial problems. *Journal of computer and system sciences*, 9(3):256–278, 1974.
- [20] Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within $2-\epsilon$. *Journal of Computer and System Sciences*, 74(3):335–349, 2008.
- [21] László Lovász. On the ratio of optimal integral and fractional covers. *Discrete mathematics*, 13(4):383–390, 1975.
- [22] Nimrod Megiddo and Kenneth J Supowit. On the complexity of some common geometric location problems. *SIAM journal on computing*, 13(1):182–196, 1984.
- [23] Nimrod Megiddo and Arie Tamir. On the complexity of locating linear facilities in the plane. *Operations research letters*, 1(5):194–197, 1982.
- [24] Christian Worm Mortensen. Fully dynamic orthogonal range reporting on ram. *SIAM Journal on Computing*, 35(6):1494–1525, 2006.
- [25] Nabil H Mustafa and Saurabh Ray. Improved results on geometric hitting set problems. *Discrete & Computational Geometry*, 44(4):883–895, 2010.
- [26] Ofer Neiman and Shay Solomon. Simple deterministic algorithms for fully dynamic maximal matching. *ACM Transactions on Algorithms (TALG)*, 12(1):7, 2016.
- [27] Shay Solomon. Fully dynamic maximal matching in constant update time. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 325–334. IEEE, 2016.

A Missing proofs

In this section, we give the proofs of the technical lemmas in Section 4.1 and Section 4.2, which are missing in the main body of the paper.

A.1 Proof of Lemma 1

Consider an interval set cover instance (S, \mathcal{I}) of size n .

Lemma 28. *One can store S in a basic data structure that can report in $\tilde{O}(1)$ time, for a query point $q \in \mathbb{R}$, the leftmost point in S that is to the right of q (if it exists).*

Proof. We simply store S in a binary search tree. Given a query point $q \in \mathbb{R}$, one can find the leftmost point in S to the right of q by searching in the tree in $\tilde{O}(1)$ time. Clearly, the tree can be built in $\tilde{O}(|S|)$ time and dynamized with $\tilde{O}(1)$ update time, and hence it is basic. \square

Lemma 29. *One can store \mathcal{I} in a basic data structure that can report in $\tilde{O}(1)$ time, for a query point $q \in \mathbb{R}$, the interval in \mathcal{I} with the rightmost right endpoint that contains q (if it exists).*

Proof. We store \mathcal{I} in a binary search tree \mathcal{T} where the key of each interval is its left endpoint. We augment each node $\mathbf{u} \in \mathcal{T}$ with an additional field which stores the interval in the subtree rooted at \mathbf{u} that has the rightmost right endpoint. Given a query point $q \in \mathbb{R}$, we first look for the interval in \mathcal{T} with the rightmost right endpoint whose key is smaller than or equal to q . With the augmented fields, this interval can be found in $\tilde{O}(1)$ time. If this interval contains q , then we report it; otherwise, no interval in \mathcal{I} contains q . Clearly, \mathcal{T} can be built in $\tilde{O}(|\mathcal{I}|)$ time and dynamized with $\tilde{O}(1)$ update time, and hence it is basic. \square

We store S in the data structure of Lemma 28 and \mathcal{I} in the data structure of Lemma 29.

Recall the greedy algorithm for the static interval set cover problem. The algorithm computes an optimal set cover \mathcal{I}_{opt} for (S, \mathcal{I}) by repeatedly picking the leftmost uncovered point and covering it using the rightmost interval. Formally, the procedure is presented below.

1. $q \leftarrow -\infty$ and $\mathcal{I}_{\text{opt}} \leftarrow \emptyset$.
2. Find the leftmost point $a \in S$ to the right of q .
3. $\mathcal{I}_{\text{opt}} \leftarrow \mathcal{I}_{\text{opt}} \cup \{I_a\}$, where I_a is the right-rightmost interval in \mathcal{I} covering a .
4. $q \leftarrow$ the right endpoint of I_a , then go to Step 2.

With the data structures of Lemma 28 and Lemma 29 in hand, the greedy algorithm can be performed in $\tilde{O}(\text{opt})$ time, where opt is the optimum of (S, \mathcal{I}) . Indeed, the data structure of Lemma 28 can be used to implement Step 2 and the data structure of Lemma 29 can be used to implement Step 3, both in $\tilde{O}(1)$ time. Since the algorithm terminates after $O(\text{opt})$ iterations, the total time cost is $\tilde{O}(\text{opt})$.

A.2 Proof of Lemma 3

Our algorithm makes a no-solution decision iff (S_i, \mathcal{I}_i) has no set cover for some $i \in P'$. If (S_i, \mathcal{I}_i) has a set cover for all $i \in P'$, then (S, \mathcal{I}) clearly has a set cover, because the portions J_i for $i \in P$ are coverable. If (S_i, \mathcal{I}_i) has no set cover for some $i \in P'$, then (S, \mathcal{I}) has no set cover, because J_i is uncoverable and thus the points in S_i can only be covered by the intervals in \mathcal{I}_i . Therefore, the no-solution decision of our data structure is correct.

A.3 Proof of Lemma 4

Suppose the portions J_1, \dots, J_r are sorted from left to right. Let s_i be the separation point of J_i and J_{i+1} . Observe that an interval $I \in \mathcal{I}_{\text{opt}}$ belongs to exactly two \mathcal{I}_i 's only if I contains one of the separation points s_1, \dots, s_{r-1} . We claim that for each s_i , at most two intervals in \mathcal{I}_{opt} contain s_i . Assume there are three intervals I^-, I, I^+ that contain s_i . Without loss of generality, assume that I^- (resp., I^+) has the leftmost left endpoint (resp., the rightmost right endpoint) among I^-, I, I^+ . Then one can easily see that $I \subseteq I^- \cup I^+$. Therefore, $\mathcal{I}_{\text{opt}} \setminus \{I\}$ is also a set cover for (S, \mathcal{I}) , contradicting the optimality of \mathcal{I}_{opt} . Thus, at most two intervals in \mathcal{I}_{opt} contain s_i . It follows that there are at most $2(r-1)$ intervals in \mathcal{I}_{opt} that contain some separation point, and only these intervals can belong to exactly two \mathcal{I}_i 's, which proves the lemma.

A.4 Proof of Lemma 6

We first observe that unit-square set cover and hitting set are equivalent. A unit square Z contains a point $a \in \mathbb{R}^2$ iff the unit square centered at a contains the center of Z . Therefore, given a unit-square set cover instance (S, \mathcal{Z}) , if we replace each point $a \in S$ with a unit square Z_a centered at a and replace each unit square $Z \in \mathcal{Z}$ with the center c_Z of Z , then (S, \mathcal{Z}) is reduced to an equivalent unit-square hitting set instance $(\{c_Z : Z \in \mathcal{Z}\}, \{Z_a : a \in S\})$. In the same way, one can also reduce a unit-square hitting set instance to an equivalent unit-square set cover instance. Thus, we only need to consider unit-square set cover.

In order to reduce dynamic unit-square set cover to dynamic quadrant set cover, we apply a grid technique. To this end, let us first consider a static unit-square set cover instance (S, \mathcal{Z}) . We build a grid Γ on the plane consisting of square cells of side-length 1. For each cell \square of the grid, we define $S_\square = S \cap \square$ and $\mathcal{Z}_\square = \{Z \in \mathcal{Z} : Z \cap \square \neq \emptyset\}$. We call a cell *nonempty* if $S_\square \neq \emptyset$ or $\mathcal{Z}_\square \neq \emptyset$, and *truly nonempty* if $S_\square \neq \emptyset$. Let Ψ be the set of all truly nonempty cells. It is clear that (S, \mathcal{Z}) has a set cover iff $(S_\square, \mathcal{Z}_\square)$ has a set cover for all $\square \in \Psi$, since the points in S_\square can only be covered using the unit squares in \mathcal{Z}_\square . Let $\mathcal{Z}_\square^* \subseteq \mathcal{Z}_\square$ be a c -approximate optimal set cover for $(S_\square, \mathcal{Z}_\square)$. We claim that $\bigsqcup_{\square \in \Psi} \mathcal{Z}_\square^*$ is a $4c$ -approximate optimal set cover for (S, \mathcal{Z}) . Consider an optimal set cover \mathcal{Z}^* for (S, \mathcal{Z}) . We have $|\mathcal{Z}_\square^*| \leq c \cdot |\mathcal{Z}^* \cap \mathcal{Z}_\square|$, because $\mathcal{Z}^* \cap \mathcal{Z}_\square$ is a set cover for $(S_\square, \mathcal{Z}_\square)$. On the other hand, we have $\sum_{\square \in \Psi} |\mathcal{Z}^* \cap \mathcal{Z}_\square| \leq 4|\mathcal{Z}^*|$, since any unit square can intersect at most four cells of Γ . It follows that

$$\left| \bigsqcup_{\square \in \Psi} \mathcal{Z}_\square^* \right| = \sum_{\square \in \Psi} |\mathcal{Z}_\square^*| \leq \sum_{\square \in \Psi} c \cdot |\mathcal{Z}^* \cap \mathcal{Z}_\square| \leq 4c \cdot |\mathcal{Z}^*|.$$

In this way, we reduce the instance (S, \mathcal{Z}) to the instances $(S_\square, \mathcal{Z}_\square)$ for $\square \in \Psi$. It seems that solving each $(S_\square, \mathcal{Z}_\square)$ is still a unit-square set cover problem. However, it is in fact a quadrant set cover problem. Indeed, a cell \square is a square of side-length 1, hence the intersection of any unit square Z and \square is equal to the intersection of a quadrant Q_Z and \square , where Q_Z is the quadrant obtained by removing the two edges of Z outside \square . Since the points in S_\square are all contained in \square , the quadrant $Q_Z \in \mathcal{Z}_\square$ covers the same set of points in S_\square as the unit square Z . Therefore, $(S_\square, \mathcal{Z}_\square)$ is equivalent to the quadrant set cover instance $(S_\square, \{Q_Z : Z \in \mathcal{Z}_\square\})$.

The above observation gives us a reduction from dynamic unit-square set cover to dynamic quadrant set cover. Suppose we have a c -approximate dynamic quadrant set cover data structure \mathcal{D} with $f(n)$ amortized update time and $\tilde{O}(n_0)$ construction time, where f is an increasing function. We are going to design a $4c$ -approximate dynamic unit-square set cover data structure \mathcal{D}' . Let (S, \mathcal{Z}) be a dynamic unit-square set cover instance. When constructing \mathcal{D}' , we build the grid Γ described above and compute $(S_\square, \mathcal{Z}_\square)$ for all nonempty cells. The grid Γ is always fixed, while the instances $(S_\square, \mathcal{Z}_\square)$ change along with the change of S and \mathcal{Z} . As argued before, we can regard each $(S_\square, \mathcal{Z}_\square)$ as a dynamic quadrant set cover instance. We then construct a data structure \mathcal{D}_\square for each $(S_\square, \mathcal{Z}_\square)$, which is the dynamic quadrant set cover data structure \mathcal{D} built on $(S_\square, \mathcal{Z}_\square)$. If $(S_\square, \mathcal{Z}_\square)$ has a set cover, then \mathcal{D}_\square maintains a c -approximate optimal set cover for $(S_\square, \mathcal{Z}_\square)$, which we denote by \mathcal{Z}_\square^* . Furthermore, we also need two data structures to maintain the nonempty and truly nonempty cells of Γ , respectively. We store the nonempty cells in a (dynamic) binary search tree \mathcal{T}_1 (by fixing some total order of the cells of Γ). Each node $\mathbf{u} \in \mathcal{T}_1$ has a pointer pointing to the data structure \mathcal{D}_\square where \square is the cell corresponding to \mathbf{u} . Similarly, we store all the truly nonempty cells in a (dynamic) binary search tree \mathcal{T}_2 , with a pointer at each node pointing to the data structure \mathcal{D}_\square for the corresponding cell \square . The data structures \mathcal{D}_\square and the binary search trees $\mathcal{T}_1, \mathcal{T}_2$ form our dynamic unit-square set cover data structure \mathcal{D}' . Besides this, \mathcal{D}' also maintains two numbers m and s , where m is the number of the truly nonempty cells \square such that $(S_\square, \mathcal{Z}_\square)$ has no set cover and s is the sum of $|\mathcal{Z}_\square^*|$ over all truly nonempty cells \square such that $(S_\square, \mathcal{Z}_\square)$ has a set cover. The construction time of \mathcal{D}' is $\tilde{O}(n_0)$. Indeed, finding the nonempty (and truly nonempty) cells and building the binary search trees $\mathcal{T}_1, \mathcal{T}_2$ can be easily done in $\tilde{O}(n_0)$ time. Furthermore, the sum of the sizes of all $(S_\square, \mathcal{Z}_\square)$ is $\tilde{O}(n_0)$, since each point in S lies in exactly one cell and each unit-square in \mathcal{Z} intersects at most four cells. Therefore, computing the instances $(S_\square, \mathcal{Z}_\square)$ can be done in $\tilde{O}(n_0)$ time. The time for building each data structure \mathcal{D}_\square is near-linear in the size of $(S_\square, \mathcal{Z}_\square)$, thus building all \mathcal{D}_\square takes $\tilde{O}(n_0)$ time.

Next, we consider how to handle updates on (S, \mathcal{Z}) . For an update on (S, \mathcal{Z}) , there are (at most) four cells \square for which $(S_\square, \mathcal{Z}_\square)$ changes, and we call them *involved* cells. For each involved cell \square , we do the following. First, if \square is previously empty (resp., not truly nonempty) and becomes nonempty (resp., truly nonempty) after the update, we insert a new node to the binary search tree \mathcal{T}_1 (resp., \mathcal{T}_2) corresponding to \square . Conversely, if \square is previously nonempty (resp., truly nonempty) and becomes empty (resp., not truly nonempty) after the update, we delete the node corresponding to \square from \mathcal{T}_1 (resp., \mathcal{T}_2). After this, we need to update the data structure \mathcal{D}_\square . If \square is previously empty and becomes nonempty after the update, we create a new data structure \mathcal{D}_\square for \square , otherwise we find the data structure \mathcal{D}_\square (which can be done by searching the node \mathbf{u} corresponding to \square in \mathcal{T}_1 and using the pointer at \mathbf{u}) and update it. We observe that all the above work can be done in $\tilde{O}(f(n))$ amortized time. Indeed, updating the binary search trees \mathcal{T}_1 and \mathcal{T}_2 takes logarithmic time. By assumption, the amortized update time of \mathcal{D}_\square is $f(n_\square)$, where n_\square is the size of the current $(S_\square, \mathcal{Z}_\square)$, which is smaller than or equal to $f(n)$ as f is an increasing function. Since there are at most four involved cells, the amortized update time of \mathcal{D}' is $\tilde{O}(f(n))$.

At any point, if $m > 0$, then \mathcal{D}' directly decides that the current (S, \mathcal{Z}) has no set cover. Otherwise, \mathcal{D}' uses $\mathcal{Z}_{\text{appx}} = \bigsqcup_{\square \in \Psi} \mathcal{Z}_\square^*$ as the set cover solution for the current (S, \mathcal{Z}) , where Ψ is the set of the truly nonempty cells. By our previous observation, $\mathcal{Z}_{\text{appx}}$ is a $4c$ -approximate optimal set cover for (S, \mathcal{Z}) . To support the desired queries for $\mathcal{Z}_{\text{appx}}$ is quite easy. First, the size of $\mathcal{Z}_{\text{appx}}$ is just the number s maintained by \mathcal{D}' , hence the size query for $\mathcal{Z}_{\text{appx}}$ can be answered in $O(1)$ time. To answer a membership query, let $Z \in \mathcal{Z}$ be the query element. We want to know the multiplicity of Z in $\mathcal{Z}_{\text{appx}}$. There are at most four cells of Γ that intersect Z . For each such cell \square , we search in the binary search tree \mathcal{T}_2 . If we cannot find a node of \mathcal{T}_2 corresponding to \square , then $\square \notin \Psi$. Otherwise, the node $\mathbf{u} \in \mathcal{T}_2$ corresponding to \square gives us a pointer to the data structure \mathcal{D}_\square , which can report the multiplicity of Z in \mathcal{Z}_\square^* in $O(\log |\mathcal{Z}_\square^*|)$ time. The sum of all these multiplicities is just the multiplicity of Z in $\mathcal{Z}_{\text{appx}}$. The time cost for answer the membership query is $O(\log |\Psi| + \log |\mathcal{Z}_{\text{appx}}|)$, since the size of \mathcal{T}_2 is $|\Psi|$. Note that $|\Psi| \leq |\mathcal{Z}_{\text{appx}}|$, since $\mathcal{Z}_\square^* \neq \emptyset$ for all $\square \in \Psi$ (for otherwise \mathcal{Z}_\square^* is not a set cover for S_\square). Thus, a membership query takes $O(\log |\mathcal{Z}_{\text{appx}}|)$ time. To answer the reporting query, we do a traversal in \mathcal{T}_2 . For each node $\mathbf{u} \in \mathcal{T}_2$, we use the data structure \mathcal{D}_\square to report the elements in \mathcal{Z}_\square^* in $O(|\mathcal{Z}_\square^*|)$ time, where $\square \in \Psi$ is the truly nonempty cell corresponding to \mathbf{u} . The time cost is $O(|\Psi| + |\mathcal{Z}_{\text{appx}}|)$. Again, since $|\Psi| \leq |\mathcal{Z}_{\text{appx}}|$, the reporting query can be answered in $O(|\mathcal{Z}_{\text{appx}}|)$ time. To summarize, \mathcal{D}' is a $4c$ -approximate dynamic unit-square set cover data structure with $\tilde{O}(f(n))$ amortized update time and $\tilde{O}(n_0)$ construction time.

Finally, we show how to reduce dynamic quadrant hitting set to dynamic quadrant set cover. Let us first consider a static quadrant hitting set instance (S, \mathcal{Q}) . There are four types of quadrants in \mathcal{Q} , southeast, southwest, northeast, and northwest; we denote by $\mathcal{Q}^{\text{SE}}, \mathcal{Q}^{\text{SW}}, \mathcal{Q}^{\text{NE}}, \mathcal{Q}^{\text{NW}} \subseteq \mathcal{Q}$ the sub-collections consisting of these types of quadrants, respectively. Let S_1, S_2, S_3, S_4 be c -approximate optimal hitting sets for the instances $(S, \mathcal{Q}^{\text{SE}}), (S, \mathcal{Q}^{\text{SE}}), (S, \mathcal{Q}^{\text{SE}}), (S, \mathcal{Q}^{\text{SE}})$, respectively. We claim that $\bigcup_{i=1}^4 S_i$ is a $4c$ -approximate hitting set for (S, \mathcal{Q}) . Indeed, we have $|S_1| \leq c \cdot \text{opt}^{\text{SE}} \leq c \cdot \text{opt}$, where opt^{SE} is the optimum of $(S, \mathcal{Q}^{\text{SE}})$ and opt is the optimum of (S, \mathcal{Q}) . Similarly, we can show that $|S_i| \leq c \cdot \text{opt}$ for all $i \in \{1, \dots, 4\}$. Thus, $|\bigcup_{i=1}^4 S_i| = \sum_{i=1}^4 |S_i| \leq 4c \cdot \text{opt}$. Therefore, if we can solve dynamic quadrant hitting set for a single type of quadrants with an approximation factor c , we are able to solve the general dynamic quadrant hitting set problem with an approximation factor $4c$.

It is easy to see that dynamic quadrant hitting set for a single type of quadrants is equivalent to dynamic quadrant set cover for a single type of quadrants. To see this, let us consider southeast quadrants. Note that a point a is contained in a southeast quadrant Q iff the northwest quadrant whose vertex is a (denoted by Q_a) contains the vertex v_Q of Q . Therefore, given a quadrant hitting set instance $(S, \mathcal{Q}^{\text{SE}})$ where \mathcal{Q}^{SE} only contains southeast quadrants, if we replace each point $a \in S$ with the northwest quadrant Q_a and replace each southeast quadrant $Q \in \mathcal{Q}^{\text{SE}}$ with its vertex v_Q , then $(S, \mathcal{Q}^{\text{SE}})$ is reduced to an equivalent quadrant set cover instance $(\{v_Q : Q \in \mathcal{Q}^{\text{SE}}\}, \{Q_a : a \in S\})$. To summarize, if we can solve dynamic quadrant set cover with an approximation factor c , then we can also solve dynamic quadrant hitting set for a single type of quadrants with an approximation factor c and in turn solve the general dynamic quadrant hitting set problem with an approximation factor $4c$. This completes the proof.

A.5 Proof of Lemma 8

Our algorithm makes a no-solution decision iff $(S_{i,j}, \mathcal{Q}_{i,j})$ has no set cover for some $(i, j) \in P'$. If $(S_{i,j}, \mathcal{Q}_{i,j})$ has a set cover for all $(i, j) \in P'$, then (S, \mathcal{Q}) clearly has a set cover, because the cells $\square_{i,j}$ for $(i, j) \in P$ are coverable. Suppose $(S_{i,j}, \mathcal{Q}_{i,j})$ has no set cover for some $(i, j) \in P'$. Then there exists a point $a \in S_{i,j}$ that is not covered by any quadrant in $\mathcal{Q}_{i,j}$. We claim that a is not covered by any quadrant in \mathcal{Q} . Consider a quadrant $Q \in \mathcal{Q}$. If Q does not intersect $\square_{i,j}$, then $a \notin Q$. Otherwise, Q must partially intersect $\square_{i,j}$, because $\square_{i,j}$ is uncoverable. If the vertex of Q lies in $\square_{i,j}$, then $Q \in \mathcal{Q}_{i,j}$ and thus $a \notin Q$. The remaining case is that Q partially intersects $\square_{i,j}$ and contains one edge of $\square_{i,j}$. Without loss of generality, assume Q left intersects $\square_{i,j}$. The rightmost quadrant Q' that left intersects $\square_{i,j}$ is contained in $\mathcal{Q}_{i,j}$. Since $Q \cap \square_{i,j} \subseteq Q' \cap \square_{i,j}$ and $a \notin Q'$, we have $a \notin Q$. It follows that a is not covered by \mathcal{Q} and hence (S, \mathcal{Q}) has no set cover. Therefore, the no-solution decision of our data structure is correct.

A.6 Proof of Lemma 9

Fix $(i, j) \in P'$. Let $\mathcal{Q}^\sim \subseteq \mathcal{Q}_{\text{opt}}$ consist of all quadrants in \mathcal{Q}_{opt} that intersect $\square_{i,j}$. Then \mathcal{Q}^\sim covers all points in $S_{i,j}$, because the points in $S_{i,j}$ can only be covered by quadrants that intersect $\square_{i,j}$. Since $\square_{i,j}$ is uncoverable by \mathcal{Q} , all quadrants in \mathcal{Q}^\sim must partially intersect $\square_{i,j}$. It follows that a quadrant in \mathcal{Q}^\sim either has its vertex in $\square_{i,j}$ or contains one edge of $\square_{i,j}$. We claim that if a point $a \in \square_{i,j}$ is covered by \mathcal{Q}^\sim , then a is covered by either $\mathcal{Q}^\sim \cap \mathcal{Q}'_{i,j}$ or a special quadrant in $\mathcal{Q}_{i,j}$. Let $Q \in \mathcal{Q}^\sim$ be a quadrant such that $a \in Q$. If $Q \in \mathcal{Q}'_{i,j}$, we are done. Otherwise, Q must contain one edge of $\square_{i,j}$, say the left edge. In this case, $\mathcal{Q}_{i,j}$ should contain a special quadrant Q_0 , which is the rightmost quadrant that left intersects $\square_{i,j}$. It is clear that $Q \cap \square_{i,j} \subseteq Q_0 \cap \square_{i,j}$, which implies that $a \in Q_0$. From this claim, it directly follows that all points in $S_{i,j}$ are covered by $(\mathcal{Q}^\sim \cap \mathcal{Q}'_{i,j}) \cup \mathcal{Q}''_{i,j}$, where $\mathcal{Q}''_{i,j} \subseteq \mathcal{Q}_{i,j}$ consists of the special quadrants. Since $(\mathcal{Q}^\sim \cap \mathcal{Q}'_{i,j}) \cup \mathcal{Q}''_{i,j} \subseteq \mathcal{Q}_{i,j}$ and $(\mathcal{Q}^\sim \cap \mathcal{Q}'_{i,j}) \cup \mathcal{Q}''_{i,j}$ covers $S_{i,j}$, we have $|(\mathcal{Q}^\sim \cap \mathcal{Q}'_{i,j}) \cup \mathcal{Q}''_{i,j}| \leq \text{opt}_{i,j}$. Note that $\mathcal{Q}^\sim \cap \mathcal{Q}'_{i,j} = \mathcal{Q}_{\text{opt}} \cap \mathcal{Q}'_{i,j}$ and $|\mathcal{Q}''_{i,j}| \leq 4$. So we have

$$|\mathcal{Q}_{\text{opt}} \cap \mathcal{Q}'_{i,j}| + 4 = |\mathcal{Q}^\sim \cap \mathcal{Q}'_{i,j}| + 4 \geq |(\mathcal{Q}^\sim \cap \mathcal{Q}'_{i,j}) \cup \mathcal{Q}''_{i,j}| \geq \text{opt}_{i,j}.$$

Equation 5 then follows from the equation

$$\text{opt} = |\mathcal{Q}_{\text{opt}}| = \sum_{i=1}^r \sum_{j=1}^r |\mathcal{Q}_{\text{opt}} \cap \mathcal{Q}'_{i,j}| \geq \sum_{(i,j) \in P'} |\mathcal{Q}_{\text{opt}} \cap \mathcal{Q}'_{i,j}|.$$

A.7 Proof of Lemma 10

It suffices to show $\sum_{k=1}^r (|S_{i,k}| + |\mathcal{Q}_{i,k}|) = O(f(n_0, \varepsilon) + r)$ for all $i \in \{1, \dots, r\}$ at any time in the first period. Fix $i \in \{1, \dots, r\}$. Initially, the i -th row of the grid contains $O(f(n_0, \varepsilon))$ points in S and $O(f(n_0, \varepsilon))$ vertices of the quadrants in \mathcal{Q} . Since the period only consists of $f(n_0, \varepsilon)$ operations, the i -th row always contains $O(f(n_0, \varepsilon))$ points in S and $O(f(n_0, \varepsilon))$ vertices of the quadrants in \mathcal{Q} during the entire period. Thus, at any time in the period, $\sum_{k=1}^r |S_{i,k}| = O(f(n_0, \varepsilon))$ and the total number of the non-special quadrants in $\mathcal{Q}_{i,1}, \dots, \mathcal{Q}_{i,r}$ is bounded by $O(f(n_0, \varepsilon))$. The number of the special quadrants in $\mathcal{Q}_{i,1}, \dots, \mathcal{Q}_{i,r}$ is clearly $O(r)$, since each $\mathcal{Q}_{i,k}$ contains at most four special quadrants. Thus, the equation $\sum_{k=1}^r (|S_{i,k}| + |\mathcal{Q}_{i,k}|) = O(f(n_0, \varepsilon) + r)$ holds at any time in the period.

A.8 Proof of Lemma 11

We first prove an invariant of our algorithm: whenever a is updated in the algorithm, the following two properties hold.

1. a is set to be the leftmost point in $S \cap U^{\text{SE}}$ that is not covered by \mathcal{Q}_{ans} .
2. Any point in $S \cap U^{\text{SE}}$ above a is not covered by \mathcal{Q}_{ans} .

The first update of a happens in Step 2. At this time, $\mathcal{Q}_{\text{ans}} = \{\Phi_{\rightarrow}(\sigma, \mathcal{Q}^{\text{SW}}), \Phi_{\uparrow}(\sigma, \mathcal{Q}^{\text{SE}})\}$, $\tilde{y} = y(\Phi_{\uparrow}(\sigma, \mathcal{Q}^{\text{SE}}))$, and a is set to be $\phi(\tilde{y})$, i.e., the leftmost point in $S \cap U^{\text{SE}}$ whose y -coordinate is larger than \tilde{y} . In order to prove the invariant, we only need to show that **(1)** any point in $S \cap U^{\text{SE}}$ strictly to the left of a is covered by \mathcal{Q}_{ans} and **(2)** any point in $S \cap U^{\text{SE}}$ above a (including a itself) is not covered by \mathcal{Q}_{ans} . Let $b \in S \cap U^{\text{SE}}$ be a point strictly to the left of a . First, the y -coordinate of b is at most \tilde{y} , since a is the leftmost point in $S \cap U^{\text{SE}}$ whose y -coordinate is larger than \tilde{y} . Therefore, if b is to the right of σ , then $b \in \Phi_{\uparrow}(\sigma, \mathcal{Q}^{\text{SE}})$. Now assume b is strictly to the left of σ . Because σ is on the boundary γ of U^{SE} , b must be below σ . Thus, $b \in \Phi_{\rightarrow}(\sigma, \mathcal{Q}^{\text{SW}})$. We see that \mathcal{Q}_{ans} covers b . Let $a' \in S \cap U^{\text{SE}}$ be a point above a . Then the y -coordinate of a' is greater than \tilde{y} , which implies $a' \notin \Phi_{\uparrow}(\sigma, \mathcal{Q}^{\text{SE}})$. Furthermore, a' must be strictly to the right of σ . Indeed, if a' is to the left of σ , then a' is below σ (as argued above) and hence the y -coordinate of a' is at most \tilde{y} , resulting in a contradiction. Now we see a' is strictly above σ and strictly to the right of σ . Note that σ is on the boundary of U^{SW} , which implies $a' \notin U^{\text{SW}}$ and in particular $a' \notin \Phi_{\rightarrow}(\sigma, \mathcal{Q}^{\text{SW}})$. So a' is not covered by \mathcal{Q}_{ans} . The invariant holds when we update a in Step 2.

In the rest of the procedure, a is only updated in Step 5. Note that Step 3–5 form a loop in which we include a constant number of quadrants to cover a , see Figure 4. It suffices to show that if the two properties hold at the beginning of an iteration of the loop (before Step 3), then they also hold after we update a in Step 5 of that iteration. Suppose we are at the beginning of an iteration, and the two properties hold. If $a \in U^{\text{NE}}$, then we go directly from Step 3 to Step 6 and the algorithm terminates. Otherwise, we go to Step 4. Here we need to distinguish two cases: $a \in U^{\text{NW}}$ and $a \notin U^{\text{NW}}$.

[Case 1] We first consider the case $a \notin U^{\text{NW}}$. In this case, we only include in \mathcal{Q}_{ans} a new quadrant $Q = \Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}})$. Note that after including Q , \mathcal{Q}_{ans} covers all the points in $S \cap U^{\text{SE}}$ whose y -coordinates are at most $y(Q)$. To see this, consider a point $b \in S \cap U^{\text{SE}}$ whose y -coordinate is at most $y(Q)$. If b is strictly to the left of a , then b is covered by \mathcal{Q}_{ans} even before we include Q , since a is the leftmost point in $S \cap U^{\text{SE}}$ that is not covered by \mathcal{Q}_{ans} at the beginning of this iteration (by our assumption). Otherwise, b is to the right of a and is hence covered by Q . On the other hand, we also notice that after including Q , \mathcal{Q}_{ans} does not cover any the points in $S \cap U^{\text{SE}}$ whose y -coordinates are greater than $y(Q)$. To see this, consider a point $b \in S \cap U^{\text{SE}}$ whose y -coordinate is greater than $y(Q)$. Then $b \notin Q$ and b is strictly above a (since $a \in Q$). By our assumption, at the beginning of this iteration (when Q is not included in \mathcal{Q}_{ans}), \mathcal{Q}_{ans} does not cover any point above a , and in particular does not cover b . Since $b \notin Q$, we see that, after we include Q in \mathcal{Q}_{ans} , the new \mathcal{Q}_{ans} still does not cover b . To summarize, the new \mathcal{Q}_{ans} contains all points in $S \cap U^{\text{SE}}$ whose y -coordinates are at most $y(Q)$ and no points in $S \cap U^{\text{SE}}$ whose y -coordinates are greater than $y(Q)$. Therefore, when we set a to be $\phi(y(Q))$ in Step 5, a is the leftmost point in $S \cap U^{\text{SE}}$ that is not covered by \mathcal{Q}_{ans} and any point in $S \cap U^{\text{SE}}$ above a is not covered by \mathcal{Q}_{ans} .

[Case 2] Next, we consider the case $a \in U^{\text{NW}}$. In this case, we include in \mathcal{Q}_{ans} three new quadrants: $\Phi_{\rightarrow}(a, \mathcal{Q}^{\text{NW}})$, $\Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}})$, and $Q = \Phi_{\uparrow}(v, \mathcal{Q}^{\text{SE}})$ where v is the vertex of $\Phi_{\rightarrow}(a, \mathcal{Q}^{\text{NW}})$. We first show that after including these three new quadrants, \mathcal{Q}_{ans} covers all the points in $S \cap U^{\text{SE}}$ whose y -coordinates are at most $y(Q)$. Consider a point $b \in S \cap U^{\text{SE}}$ whose y -coordinate is at most $y(Q)$. If the x -coordinate of b is at least $x(Q)$, then $b \in Q$. Otherwise, b is strictly to the left of v (because $v \in Q$). Thus, if b is above v , then $b \in \Phi_{\rightarrow}(a, \mathcal{Q}^{\text{NW}})$. It now suffices to consider the case when b is strictly below v . In this case, b is strictly below a , since v is the vertex of $\Phi_{\rightarrow}(a, \mathcal{Q}^{\text{NW}})$. If b is to the right of a , then $b \in \Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}})$. If b is strictly to the left of a , then b is covered by \mathcal{Q}_{ans} even before we include the three new quadrants, since a is the leftmost point in $S \cap U^{\text{SE}}$ that is not covered by \mathcal{Q}_{ans} at the beginning of this iteration (by our assumption). In any case, b is covered by \mathcal{Q}_{ans} . On the other hand, we notice that after including these three new quadrants, \mathcal{Q}_{ans} does not cover any points in $S \cap U^{\text{SE}}$ whose y -coordinates are greater than $y(Q)$. To see this, consider a point $b \in S \cap U^{\text{SE}}$ whose y -coordinate is greater than $y(Q)$. We first establish some obvious facts. Since the NW quadrant $\Phi_{\rightarrow}(a, \mathcal{Q}^{\text{NW}})$ and the SE quadrant $\Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}})$ both contain a , we have $v \in \Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}})$. It follows that $y(\Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}})) \leq y(Q)$, by the definition of Q . Therefore, b is strictly above a , since the y -coordinate of a is at most $y(\Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}}))$. By our assumption, at the beginning of this iteration (when the three new quadrants are not included in \mathcal{Q}_{ans}), \mathcal{Q}_{ans} does not cover any point above a , and in particular does not cover b . Also, the fact $y(\Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}})) \leq y(Q)$ implies $b \notin \Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}})$. To see $b \notin \Phi_{\rightarrow}(a, \mathcal{Q}^{\text{NW}})$, let $Q' \in \mathcal{Q}^{\text{SE}}$ be a quadrant contains b (such a quadrant exists as $b \in U^{\text{SE}}$). Then $y(Q') > y(Q)$. By the

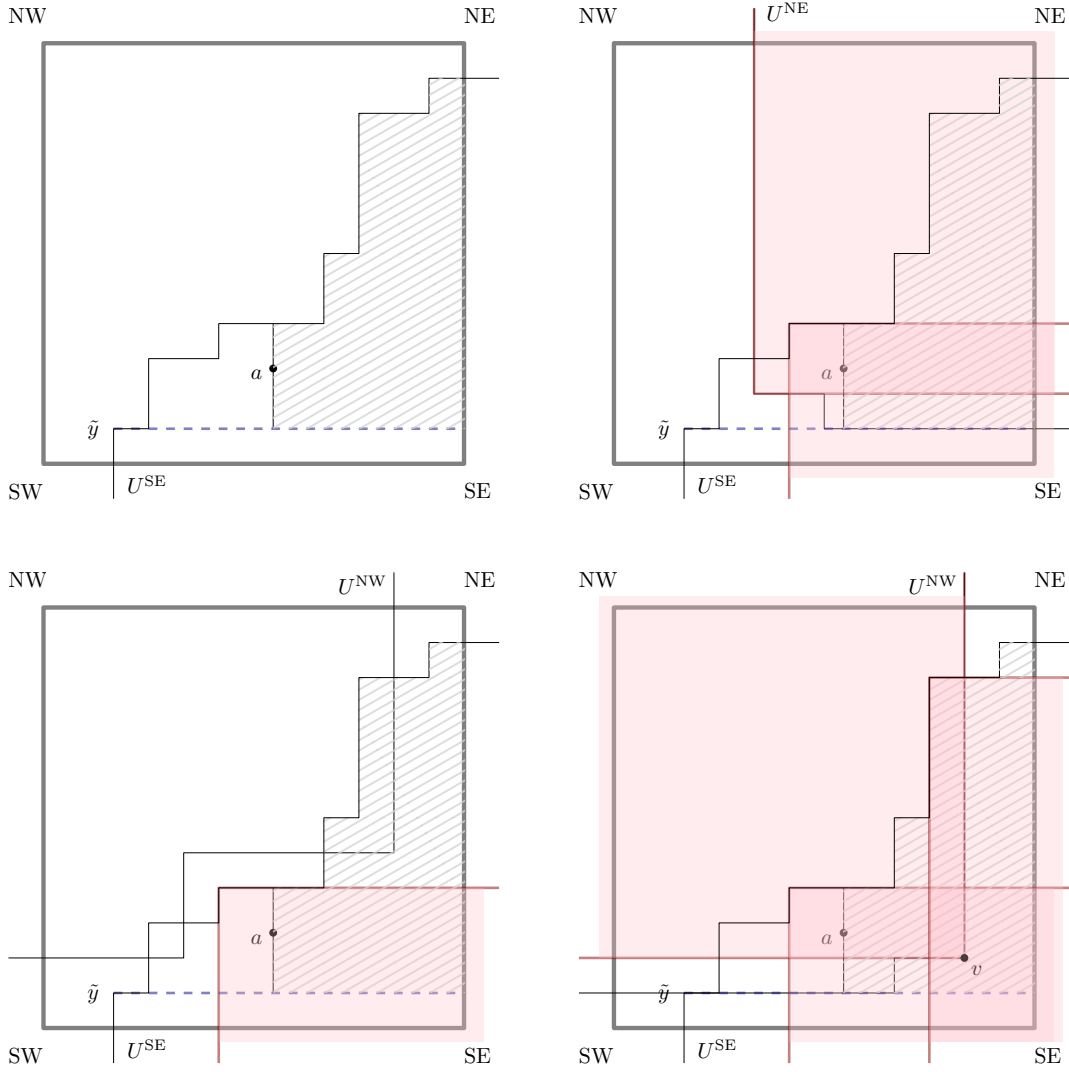


Figure 4: Different options for covering a , depending on what other quadrant types cover a . The shaded region contains all remaining uncovered points (top-left figure). If $a \in U^{\text{NE}}$, then the entire shaded region can be covered by two quadrants (top-right figure). Otherwise (bottom figures), some progress can be made using up to three quadrants from Q^{SE} and Q^{NW} . After choosing these quadrants, the remaining uncovered points of $S \cap U^{\text{SE}}$ are not covered by any quadrant of Q^{SE} , Q^{NE} , and Q^{NW} that also covers a . Note that $a \notin U^{\text{SW}}$.

definition of Q , we have $v \notin Q'$ and thus $Q' \cap \Phi_{\rightarrow}(a, \mathcal{Q}^{\text{NW}}) = \emptyset$, which implies $b \notin \Phi_{\rightarrow}(a, \mathcal{Q}^{\text{NW}})$. Finally, it is clear that $b \notin Q$. So after we include the three new quadrants in \mathcal{Q}_{ans} , the new \mathcal{Q}_{ans} still does not cover b . To summarize, the new \mathcal{Q}_{ans} contains all points in $S \cap U^{\text{SE}}$ whose y -coordinates are at most $y(Q)$ and no points in $S \cap U^{\text{SE}}$ whose y -coordinates are greater than $y(Q)$. Therefore, when we set a to be $\phi(y(Q))$ in Step 5, a is the leftmost point in $S \cap U^{\text{SE}}$ that is not covered by \mathcal{Q}_{ans} and any point in $S \cap U^{\text{SE}}$ above a is not covered by \mathcal{Q}_{ans} .

Combining the discussions for the two cases, the invariant is proved. With the invariant in hand, we now prove the lemma. For convenience, we denote by a_i the point a in the i -th iteration (before the update of Step 5) of the loop Step 3–5. The invariant of our algorithm guarantees that a_i is the leftmost point in $S \cap U^{\text{SE}}$ that is not covered by \mathcal{Q}_{ans} at the beginning of the i -th iteration. Furthermore, in the i -th iteration, we always add the quadrant $\Phi_{\uparrow}(a_i, \mathcal{Q}^{\text{SE}})$ to \mathcal{Q}_{ans} , hence a_i is covered by \mathcal{Q}_{ans} at the end of the i -th iteration. This implies that our algorithm always terminates, because \mathcal{Q}_{ans} covers at least one more point in $S \cap U^{\text{SE}}$ in each iteration.

We first show (the eventual) \mathcal{Q}_{ans} covers all points in $S \cap U^{\text{SE}}$. If we go to Step 6 from Step 1, then $S \cap U^{\text{SE}} = \emptyset$ and nothing needs to be proved. If we go to Step 6 from Step 2, we know that the y -coordinates of all points in $S \cap U^{\text{SE}}$ are at most $y(\Phi_{\uparrow}(\sigma, \mathcal{Q}^{\text{SE}}))$. Consider a point $a \in S \cap U^{\text{SE}}$. If a is to the right of σ , then $a \in \Phi_{\uparrow}(\sigma, \mathcal{Q}^{\text{SE}})$, as the y -coordinate of a is at most $y(\Phi_{\uparrow}(\sigma, \mathcal{Q}^{\text{SE}}))$. If a is to the left of σ , then $a \in \Phi_{\rightarrow}(\sigma, \mathcal{Q}^{\text{SW}})$, because any point in U^{SE} to the left of σ is covered by $\Phi_{\rightarrow}(\sigma, \mathcal{Q}^{\text{SW}})$. Thus, a is covered by \mathcal{Q}_{ans} , implying that all points in $S \cap U^{\text{SE}}$ are covered by \mathcal{Q}_{ans} . The last case is that we go to Step 6 from Step 3. By the invariant, we know that a is the leftmost point in $S \cap U^{\text{SE}}$ that is not covered by \mathcal{Q}_{ans} just before Step 3. When we include the two quadrants $\Phi_{\uparrow}(a, \mathcal{Q}^{\text{NE}})$ and $\Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}})$ in \mathcal{Q}_{ans} , any point to the right of a is covered by \mathcal{Q}_{ans} . Thus, all points in $S \cap U^{\text{SE}}$ are covered by \mathcal{Q}_{ans} when we go from Step 3 to Step 6. The remaining case is that we go to Step 6 from Step 5. This case happens only when $\phi(\tilde{y})$ does not exist, or equivalently, the y -coordinate of every point in $S \cap U^{\text{SE}}$ is at most $y(Q)$. Recall that when proving the invariant, we showed that after we include Q in \mathcal{Q}_{ans} in Step 5, \mathcal{Q}_{ans} covers all the points in $S \cap U^{\text{SE}}$ whose y -coordinates are at most $y(Q)$. Hence, all points in $S \cap U^{\text{SE}}$ are covered by \mathcal{Q}_{ans} when we go from Step 5 to Step 6.

We then show that $|\mathcal{Q}_{\text{ans}}| = O(\text{opt}^{\text{SE}})$. To this end, we notice that in each iteration of the loop Step 3–5, we include in \mathcal{Q}_{ans} a constant number of quadrants. Suppose we do k iterations in total during the algorithm. It suffices to show that $k = O(\text{opt}^{\text{SE}})$. We claim that any quadrant $Q \in \mathcal{Q}$ contains at most one of a_1, \dots, a_k . First, we observe that any quadrant in \mathcal{Q}^{NE} may only contain a_k . Indeed, if $a_i \in U^{\text{SE}}$ for some $i < k$, then in the i -th iteration, the algorithm goes from Step 3 to Step 6 and terminates. Next, we show that any quadrant in \mathcal{Q}^{SW} does not contain a_i (or equivalently, $a_i \notin U^{\text{SW}}$) for all $i \in \{1, \dots, k\}$. Let $i \in \{1, \dots, k\}$. By the invariant we proved before, a_i is the leftmost point in $S \cap U^{\text{SE}}$ that is not covered by \mathcal{Q}_{ans} at the beginning of the i -th iteration. Since we include $\Phi_{\rightarrow}(\sigma, \mathcal{Q}^{\text{SW}})$ and $\Phi_{\uparrow}(\sigma, \mathcal{Q}^{\text{SE}})$ in \mathcal{Q}_{ans} Step 2, we know that $a_i \notin \Phi_{\rightarrow}(\sigma, \mathcal{Q}^{\text{SW}})$ and $a_i \notin \Phi_{\uparrow}(\sigma, \mathcal{Q}^{\text{SE}})$. Thus, a_i must be strictly above σ , since any point below σ is covered by either $\Phi_{\rightarrow}(\sigma, \mathcal{Q}^{\text{SW}})$ or $\Phi_{\uparrow}(\sigma, \mathcal{Q}^{\text{SE}})$. It follows that a_i is to the right of σ , because $a_i \in U^{\text{SE}}$ and σ is on the boundary γ of U^{SE} . In fact, a_i is strictly to the right of σ , since any point in $S \cap U^{\text{SE}}$ that has the same x -coordinate as σ is covered by $\Phi_{\uparrow}(\sigma, \mathcal{Q}^{\text{SE}})$. Because σ is also on the boundary of U^{SW} , given the fact that a_i is strictly above σ and strictly to the right of σ , we see that $a_i \notin U^{\text{SW}}$.

Now it suffices to verify that any quadrant in \mathcal{Q}^{SE} or \mathcal{Q}^{NW} contains at most one of a_1, \dots, a_k . Let $i, j \in \{1, \dots, k\}$ such that $i < j$. We want to show that no quadrant in \mathcal{Q}^{SE} or \mathcal{Q}^{NW} contains both a_i and a_j . Since a_j is not covered by \mathcal{Q}_{ans} at the beginning of the j -th iteration, it is also not covered by \mathcal{Q}_{ans} at any point before the j -th iteration and in particular, at the beginning of the i -th iteration. This implies that a_j is to the right of a_i . In the i -th iteration, we always add the quadrant $\Phi_{\uparrow}(a_i, \mathcal{Q}^{\text{SE}})$ to \mathcal{Q}_{ans} . We have $a_j \notin \Phi_{\uparrow}(a_i, \mathcal{Q}^{\text{SE}})$, since a_j is not covered by \mathcal{Q}_{ans} at the end of the i -th iteration. Thus, a_j is strictly above a_i . Now let $Q \in \mathcal{Q}^{\text{SE}}$ be a quadrant that contains a_i . Note that a point to the right of a_i is contained in Q only if it is contained in $\Phi_{\uparrow}(a_i, \mathcal{Q}^{\text{SE}})$, simply by the definition of Φ_{\uparrow} . Because $a_j \notin \Phi_{\uparrow}(a_i, \mathcal{Q}^{\text{SE}})$, we have $a_j \notin Q$. Consequently, no quadrant in \mathcal{Q}^{SE} contains both a_i and a_j . Next, we show that no quadrant in \mathcal{Q}^{NW} contains both a_i and a_j . If $a_i \notin U^{\text{NW}}$, we are done. So assume $a_i \in U^{\text{NW}}$. In this case, we add the quadrant $\Phi_{\rightarrow}(a_i, \mathcal{Q}^{\text{NW}})$ to \mathcal{Q}_{ans} in Step 4 of the i -th iteration. Again, we have $a_j \notin \Phi_{\rightarrow}(a_i, \mathcal{Q}^{\text{NW}})$ for a_j is

not covered by \mathcal{Q}_{ans} at the end of the i -th iteration. Let $Q \in \mathcal{Q}^{\text{NW}}$ be a quadrant that contains a_i . Note that a point above a_i is contained in Q only if it is contained in $\Phi_{\rightarrow}(a_i, \mathcal{Q}^{\text{NW}})$, simply by the definition of Φ_{\rightarrow} . Because $a_j \notin \Phi_{\rightarrow}(a_i, \mathcal{Q}^{\text{NW}})$, we have $a_j \notin Q$. Consequently, no quadrant in \mathcal{Q}^{NW} contains both a_i and a_j .

To summarize, we showed that any quadrant $Q \in \mathcal{Q}$ contains at most one of a_1, \dots, a_k . Therefore, covering all of a_1, \dots, a_k requires at least k quadrants in \mathcal{Q} . This implies $k \leq \text{opt}^{\text{SE}}$, because $a_1, \dots, a_k \in S \cap U^{\text{SE}}$. As a result, $|\mathcal{Q}_{\text{ans}}| = O(k) = O(\text{opt}^{\text{SE}})$.

A.9 Proof of Lemma 16

Let $\Pi = (S, \mathcal{R})$ be a dynamic set cover instance with only point updates, and $i \geq 0$ be a number. We claim that $|\text{opt}_i(\Pi) - \text{opt}_{i+1}(\Pi)| \leq 1$. Denote by S_i and S_{i+1} be the point set S at the time i and $i+1$. Also, let $S'_i \subseteq S_i$ and $S'_{i+1} \subseteq S_{i+1}$ consist of the points that are covered by \mathcal{R} . Suppose the $(i+1)$ -th operation inserts a point a to S , so $S_{i+1} = S_i \cup \{a\}$. If a is not covered by \mathcal{R} , then $S'_{i+1} = S'_i$ and $\text{opt}_{i+1}(\Pi) = \text{opt}_i(\Pi)$. If a is covered by \mathcal{R} , then $S'_{i+1} = S'_i \cup \{a\}$. Let $R \in \mathcal{R}$ be a range that covers a . An optimal set cover for (S'_i, \mathcal{R}) together with R is a set cover for (S'_{i+1}, \mathcal{R}) . It follows that $\text{opt}_i(\Pi) \leq \text{opt}_{i+1}(\Pi) \leq \text{opt}_i(\Pi) + 1$. The case where the $(i+1)$ -th operation deletes a point from S is symmetric. This shows that dynamic set cover in the partially dynamic setting is stable. The same argument can also be applied to show that dynamic hitting set in the partially dynamic setting is stable.

Next, we consider the dynamic interval hitting set problem (in the fully dynamic setting). Let $\Pi = (S, \mathcal{I})$ be a dynamic interval hitting set instance, and $i \geq 0$ be a number. Denote by S_i and \mathcal{I}_i as the point set S and interval collection \mathcal{I} at the time i . Let $\mathcal{I}'_i \subseteq \mathcal{I}_i$ consist of the intervals that are hit by S_i . We want to show that $|\text{opt}_i(\Pi) - \text{opt}_{i+1}(\Pi)| \leq 1$. We distinguish two cases.

[Case 1] The $(i+1)$ -th operation happens on \mathcal{I} . Suppose the $(i+1)$ -th operation is an insertion on \mathcal{I} , and let I be the interval inserted. If I is not hit by S_i , then $(S_{i+1}, \mathcal{I}'_{i+1}) = (S_i, \mathcal{I}'_i)$ and $\text{opt}_{i+1}(\Pi) = \text{opt}_i(\Pi)$. If I is hit by S_i , then $(S_{i+1}, \mathcal{I}'_{i+1}) = (S_i, \mathcal{I}'_i \cup \{I\})$. In this case, we have $\text{opt}_i(\Pi) \leq \text{opt}_{i+1}(\Pi) \leq \text{opt}_i(\Pi) + 1$, since an optimal hitting set for (S_i, \mathcal{I}'_i) together with a point hitting I is a hitting set for $(S_{i+1}, \mathcal{I}'_{i+1})$. Thus, $|\text{opt}_i(\Pi) - \text{opt}_{i+1}(\Pi)| \leq 1$. The case where the $(i+1)$ -th operation is a deletion on \mathcal{I} is symmetric.

[Case 2] The $(i+1)$ -th operation happens on S . Suppose the $(i+1)$ -th operation is an insertion on S , and let a be the point inserted. Then $(S_{i+1}, \mathcal{I}'_{i+1}) = (S_i \cup \{a\}, \mathcal{I}'_i \cup \mathcal{J})$ where $\mathcal{J} = \{I \in \mathcal{I} : a \in I\}$. It is clear that $\text{opt}_{i+1}(\Pi) \leq \text{opt}_i(\Pi) + 1$, because an optimal hitting set for (S_i, \mathcal{I}'_i) together with the point a is a hitting set for $(S_{i+1}, \mathcal{I}'_{i+1})$. It suffices to show that $\text{opt}_i(\Pi) \leq \text{opt}_{i+1}(\Pi) + 1$. Let $S^* \subseteq S_{i+1}$ be an optimal hitting set for $(S_{i+1}, \mathcal{I}'_{i+1})$. If $a \notin S^*$, then S^* is also an optimal hitting set for (S_i, \mathcal{I}'_i) . Otherwise, let $a^- \in S_i$ (resp., $a^+ \in S_i$) be the rightmost (resp., leftmost) point to the left (resp., right) of a ; in the case where a is the leftmost (resp., rightmost) point in S_i , then let a^- (resp., a^+) be an arbitrary point in S_i . We claim that $(S^* \setminus \{a\}) \cup \{a^-, a^+\}$ is a hitting set for (S_i, \mathcal{I}'_i) . Consider an interval $I \in \mathcal{I}'_i$. If $I \notin \mathcal{J}$, then I is hit by some point in $S^* \setminus \{a\}$. Otherwise, $I \in \mathcal{J}$, and thus I is hit by a . Since $I \in \mathcal{I}'_i$, I is also hit by some point in S_i , say b . If b is to the left of a , then I must be hit by a^- . On the other hand, if b is to the right of a , I must be hit by a^+ . Thus, in any case, I is hit by $(S^* \setminus \{a\}) \cup \{a^-, a^+\}$. It follows that $|(S^* \setminus \{a\}) \cup \{a^-, a^+\}| = \text{opt}_{i+1}(\Pi) + 1 \geq \text{opt}_i(\Pi)$. The case where the $(i+1)$ -th operation is a deletion on S is symmetric.

A.10 Proof of Lemma 18

For convenience, we include in S two dummy points $x^- = -\infty$ and $x^+ = +\infty$. We assume these two dummy points are always in S . Since x^- and x^+ do not hit any interval, including them in S is safe.

First, we store \mathcal{I} in a binary search tree \mathcal{T}_1 where the key of an interval is its left endpoint. We augment each node $\mathbf{u} \in \mathcal{T}_1$ with an additional field which stores the interval in the subtree rooted at \mathbf{u} with the leftmost right endpoint. We notice that using \mathcal{T}_1 , we can decide in $\tilde{O}(1)$ time, for two given numbers $a, a^+ \in \mathbb{R}$, whether there is an interval in \mathcal{I} whose both endpoints are in the open interval (a^-, a^+) . Specifically, we first look for the interval in \mathcal{T}_1 with the leftmost right endpoint whose key is greater than a . With the augmented

fields, this interval can be found in $O(\log n)$ time. If the two endpoints of this interval are contained (a, a^+) , then we return “yes”; otherwise, no interval in \mathcal{I} has two endpoints in (a, a^+) and we return “no”. Clearly, \mathcal{T}_1 can be constructed in $\tilde{O}(\mathcal{I})$ time and dynamized with $\tilde{O}(1)$ update time, and hence it is basic.

We then store S in a standard range tree \mathcal{T}_2 . The points in S are stored at the leaves of \mathcal{T}_2 . Each node $\mathbf{u} \in \mathcal{T}_2$ corresponds to a *canonical subset* $S(\mathbf{u})$ of S consisting of the points stored in the subtree rooted at \mathbf{u} . At each internal node $\mathbf{u} \in \mathcal{T}_2$, we store the leftmost and rightmost points in $S(\mathbf{u})$ as well as a separation point $s_{\mathbf{u}} \in \mathbb{R}$ such that all points in the left (resp., right) subtree of \mathbf{u} are to the left (resp., right) of $s_{\mathbf{u}}$. At each leaf \mathbf{u} of \mathcal{T} that does not correspond to the point x^+ , we maintain an additional field $\sigma(\mathbf{u})$ indicating whether there exists an interval in \mathcal{I} whose both endpoints are in the open interval $(a_{\mathbf{u}}, a_{\mathbf{u}}^+)$ where $a_{\mathbf{u}} \in S$ is the point corresponding to \mathbf{u} and $a_{\mathbf{u}}^+ \in S$ is the leftmost point that is to the right of $a_{\mathbf{u}}$; we set $\sigma(\mathbf{u}) = 1$ if such an interval exists and set $\sigma(\mathbf{u}) = 0$ otherwise. Note that using the binary search tree \mathcal{T}_1 , $\sigma(\mathbf{u})$ can be computed in $\tilde{O}(1)$ time for each leaf \mathbf{u} of \mathcal{T}_2 . Indeed, we only need to find $a_{\mathbf{u}}^+$, which can be done in $\tilde{O}(1)$ time by a walk in \mathcal{T}_2 , and query \mathcal{T}_1 to see whether there is an interval in \mathcal{I} whose both endpoints are in the open interval $(a_{\mathbf{u}}, a_{\mathbf{u}}^+)$. Besides, we also maintain a counter σ^* that is the total number of the leaves whose σ -values are equal to 1. It is easy to see that (S, \mathcal{I}) has a hitting set iff $\sigma^* = 0$.

It is clear that \mathcal{T}_2 can be constructed in $\tilde{O}(|S|)$ time. We show that when (S, \mathcal{I}) changes, we can maintain \mathcal{T}_2 with the σ -fields in $\tilde{O}(1)$ time, by using \mathcal{T}_1 . Suppose a point $a \in \mathbb{R}$ is inserted to S . We need to first compute the σ -value of the leaf corresponding to a using \mathcal{T}_1 . Let $b \in S$ be the rightmost point to the left of a . Due to the insertion of a , the σ -value of the leaf \mathbf{u} of \mathcal{T}_2 corresponding to b may also change. So we need to find \mathbf{u} , which can be done in $\tilde{O}(1)$ time by a walk in \mathcal{T}_2 , and update $\sigma(\mathbf{u})$ using \mathcal{T}_1 . The σ -values of all the other leaves remain unchanged. A deletion of a point from S is handled similarly. Now suppose an interval I is inserted to or deleted from \mathcal{I} . We find in $\tilde{O}(1)$ the leaf \mathbf{u} of \mathcal{T}_2 corresponding to the rightmost point in S that is to the left of the left endpoint of I . Note that the insertion/deletion of I does not change the σ -values of the leaves other than \mathbf{u} . So it suffices to re-compute $\sigma(\mathbf{u})$ using \mathcal{T}_1 . The time cost for all the cases above is $\tilde{O}(1)$. The rotations of \mathcal{T}_2 (for self-balancing) do not change the σ -fields. It follows that the counter σ^* can also be maintained in $\tilde{O}(1)$ time.

The dynamic data structure \mathcal{B} in the lemma just consists of (the dynamic versions of) the binary search tree \mathcal{T}_1 and the range tree \mathcal{T}_2 . The update time of \mathcal{B} is $\tilde{O}(1)$ time and the construction time is $\tilde{O}(n_0)$. To indicate whether the current (S, \mathcal{I}) has a hitting set or not, we simply check whether $\sigma^* = 0$ or not.

A.11 Proof of Lemma 19

We can simply store S in a binary search tree. Then the rightmost (resp., leftmost) point in S to the left (resp., right) of a given point a can be reported in $\tilde{O}(1)$ time by searching in the tree. Clearly, the binary search tree can be constructed in $\tilde{O}(|S|)$ time and dynamized with $\tilde{O}(1)$ update time, and thus it is basic.

A.12 Proof of Lemma 20

We can simply store S in a binary search tree. Then a point $a \in S$ contained in a given interval I can be reported in $\tilde{O}(1)$ time by searching in the tree. Clearly, the binary search tree can be constructed in $\tilde{O}(|S|)$ time and dynamized with $\tilde{O}(1)$ update time, and thus it is basic.

A.13 Proof of Lemma 21

We use opt' to denote the quasi-optimum of the current (S, \mathcal{I}) . It suffices to show $|S^*| \leq (1 + \varepsilon) \cdot \text{opt}'$ at any time, since we always have $\text{opt}' \leq \text{opt}$.

Initially, S^* is an optimal hitting set for (S_0, \mathcal{I}'_0) , so we have $|S^*| = \text{opt}'$ at that time. If the current S^* is obtained by re-computing using the output-sensitive algorithm, then $|S^*| = \text{opt}' = \text{opt}$, as we only do re-computation when the current (S, \mathcal{I}) has a hitting set. Suppose S^* is obtained by local modification. Consider the last re-computation of S^* , and we use S_1^* to denote the S^* at that point and use opt'_1 to denote the quasi-optimum of (S, \mathcal{I}) at that point. Then we have $|S_1^*| = \text{opt}'_1$. As argued before, cnt is the number

of the operations processed after the last re-computation of S^* . By the stability of dynamic interval hitting set (Lemma 16), we have $|\text{opt}' - \text{opt}'_1| \leq \text{cnt}$, implying $\text{opt}'_1 \leq \text{opt}' + \text{cnt}$. Furthermore, by the fact that the size of S^* either increases by 1 or keeps unchanged after each local modification, we have $|S^*| - |S^*_1| \leq \text{cnt}$. It follows that

$$|S^*| \leq |S^*_1| + \text{cnt} = \text{opt}'_1 + \text{cnt} \leq \text{opt}' + 2 \cdot \text{cnt}.$$

Note that $\text{cnt} \leq \varepsilon \cdot \text{opt}'_1 / (2 + \varepsilon)$, for otherwise we should re-compute S^* . This implies $\text{cnt} \leq (\varepsilon/2) \cdot \text{opt}'$ and hence $|S^*| \leq (1 + \varepsilon) \cdot \text{opt}'$.

A.14 Proof of Lemma 22

Let (S_i, \mathcal{I}_i) denote the instance (S, \mathcal{I}) at the time i , and S_i^* be the point set S^* at the time i . Define $\mathcal{I}'_i \subseteq \mathcal{I}_i$ as the sub-collection consisting of the intervals that are hit by S_i . Then (S_i, \mathcal{I}'_i) always has a hitting set. Furthermore, if (S_i, \mathcal{I}_i) has a hitting set, then $(S_i, \mathcal{I}'_i) = (S_i, \mathcal{I}_i)$. So it suffices to show that S_i^* is always a hitting set for (S_i, \mathcal{I}'_i) . We prove this by induction. It is clear that S_0^* is a hitting set for (S_0, \mathcal{I}'_0) . Assume S_{i-1}^* is a hitting set for $(S_{i-1}, \mathcal{I}'_{i-1})$ and we show that S_i^* is a hitting set for (S_i, \mathcal{I}'_i) . If S_i^* is obtained by re-computing, then $(S_i, \mathcal{I}'_i) = (S_i, \mathcal{I}_i)$, since we only re-compute S^* when the current (S, \mathcal{I}) has a hitting set. In this case, S_i^* is clearly a hitting set for both (S_i, \mathcal{I}'_i) and (S_i, \mathcal{I}_i) . So suppose S_i^* is obtained by local modification.

We consider different cases separately according to the i -th operation. If the i -th operation inserts a point a to S , then $S_i = S_{i-1} \cup \{a\}$ and $\mathcal{I}_i = \mathcal{I}_{i-1}$. In this case, $S_i^* = S_{i-1}^* \cup \{a\}$ and $\mathcal{I}'_i = \mathcal{I}'_{i-1} \cup \mathcal{J}$, where $\mathcal{J} = \{I \in \mathcal{I}_i : a \in I\}$. The intervals in \mathcal{I}'_{i-1} are hit by S_{i-1}^* and the intervals in \mathcal{J} are hit by the point a . Hence, S_i^* is a hitting set for (S_i, \mathcal{I}'_i) . If the i -th operation deletes a point a from S , then $S_i = S_{i-1} \setminus \{a\}$ and $\mathcal{I}_{i-1} = \mathcal{I}_i$. In this case, S_i^* is obtained from $S_{i-1}^* \setminus \{a\}$ by adding the rightmost point a^- to the left of a and the leftmost point a^+ in S_i to the right of a (if they exist). Consider an interval $I \in \mathcal{I}'_i$. We want to show that I is hit by S_i^* . Note that $\mathcal{I}'_i \subseteq \mathcal{I}'_{i-1}$. Thus, $I \in \mathcal{I}'_{i-1}$ and I is hit by S_{i-1}^* . If $a \notin I$, then I is hit by $S_{i-1}^* \setminus \{a\}$ and hence hit by S_i^* . Otherwise, $a \in I$. Since $I \in \mathcal{I}'_i$ and $a \notin S_i$, I must be hit by some point $b \in S_i$ different from a . If b is to the left of a , then I must be hit by a^- (as a^- is in between b and a). On the other hand, if b is to the right of a , I must be hit by a^+ (as a^+ is in between a and b). Therefore, I is hit by S_i^* . If the i -th operation inserts an interval I to \mathcal{I} , then $S_i = S_{i-1}$ and $\mathcal{I}_i = \mathcal{I}_{i-1} \cup \{I\}$. In this case, S_i^* is obtained from S_{i-1}^* by adding an arbitrary point $a \in S_i$ that hits I (if I is hit by S_i). If I is not hit by S_i , then $(S_i, \mathcal{I}'_i) = (S_{i-1}, \mathcal{I}'_{i-1})$ and $S_i^* = S_{i-1}^*$, thus S_i^* is a hitting set for (S_i, \mathcal{I}'_i) . If I is hit by S_i , then $\mathcal{I}'_i = \mathcal{I}'_{i-1} \cup \{I\}$ and $S_i^* = S_{i-1}^* \cup \{a\}$. The intervals in \mathcal{I}'_{i-1} are hit by S_{i-1}^* by our induction hypothesis and I is hit by the point a . Hence, S_i^* is a hitting set for (S_i, \mathcal{I}'_i) . If the i -th operation deletes an interval I from \mathcal{I} , then $S_i = S_{i-1}$ and $\mathcal{I}_i = \mathcal{I}_{i-1} \setminus \{I\}$. In this case, $S_i^* = S_{i-1}^*$. Note that $\mathcal{I}'_i \subseteq \mathcal{I}'_{i-1}$, which implies that S_i^* is a hitting set for (S_i, \mathcal{I}'_i) .

A.15 Proof of Lemma 24

Note that \mathcal{Q} is static. We simply store \mathcal{Q} in a static data structure \mathcal{B}_0 that can decide in $\tilde{O}(1)$ time, for a given point $a \in \mathbb{R}^2$, whether there exists a quadrant in \mathcal{Q} that covers a ; this can be done using a standard orthogonal stabbing data structure with $\tilde{O}(|\mathcal{Q}|)$ construction time. Then our data structure \mathcal{B} simply maintains the number \tilde{n} of the points in S that are not covered by \mathcal{Q} . Initially, \tilde{n} can be computed in $\tilde{O}(n_0)$ time by considering every point in S and use \mathcal{B}_0 to check if it is covered by \mathcal{Q} in $\tilde{O}(1)$ time. After an operation on S , we can update \tilde{n} in $\tilde{O}(1)$ time by checking whether the inserted/deleted point is covered by \mathcal{Q} using \mathcal{B}_0 .

A.16 Proof of Lemma 25

Let S_i denote the set S at the time i and \mathcal{Q}_i^* be the \mathcal{Q}^* at the time i . Define $S'_i \subseteq S_i$ as the sub-collection consisting of the points that are covered by \mathcal{Q} . Then (S'_i, \mathcal{Q}) always has a hitting set. Furthermore, if (S_i, \mathcal{Q}) has a set cover, then $(S'_i, \mathcal{Q}) = (S_i, \mathcal{Q})$. So it suffices to show that \mathcal{Q}_i^* is always a set cover for (S'_i, \mathcal{Q}) . We

prove this by induction. It is clear that \mathcal{Q}_0^* is a set cover for (S'_0, \mathcal{Q}) . Assume \mathcal{Q}_{i-1}^* is a set cover for (S'_{i-1}, \mathcal{Q}) and we show that \mathcal{Q}_i^* is a set cover for (S'_i, \mathcal{Q}) . If \mathcal{Q}_i^* is obtained by re-computing, then $(S'_i, \mathcal{Q}) = (S_i, \mathcal{Q})$, since we only re-compute \mathcal{Q}^* when the current (S, \mathcal{Q}) has a set cover. In this case, \mathcal{Q}_i^* is clearly a set cover for both (S'_i, \mathcal{Q}) and (S_i, \mathcal{Q}) . So suppose \mathcal{Q}_i^* is obtained by local modification. If the i -th iteration inserts a point a to S , then $S_i = S_{i-1} \cup \{a\}$. In this case, \mathcal{Q}_i^* is obtained by including in \mathcal{Q}_{i-1}^* an arbitrary quadrant $Q \in \mathcal{Q}$ that contains a (if such a quadrant exists). If a is not covered by \mathcal{R} , then $S'_i = S'_{i-1}$ and $\mathcal{Q}_i^* = \mathcal{Q}_{i-1}^*$, thus \mathcal{Q}_i^* is a set cover for (S'_i, \mathcal{Q}) . If a is covered by \mathcal{R} , then $S'_i = S'_{i-1} \cup \{a\}$ and $\mathcal{Q}_i^* = \mathcal{Q}_{i-1}^* \cup \{Q\}$. The points in S'_{i-1} are covered by \mathcal{Q}_{i-1}^* by our induction hypothesis and the point a is covered by Q . Hence, \mathcal{Q}_i^* is a set cover for (S'_i, \mathcal{Q}) . If the i -th iteration delete a point a from S , then $S_i = S_{i-1} \setminus \{a\}$. In this case, $\mathcal{Q}_i^* = \mathcal{Q}_{i-1}^*$. Note that $S'_i \subseteq S'_{i-1}$, which implies that \mathcal{Q}_i^* is a set cover for (S'_i, \mathcal{Q}) .

B Implementation details and detailed time analysis for the dynamic interval set cover data structure

We present the implementation details of our dynamic interval set cover data structure \mathcal{D}_{new} in Section 4.1 as well as a detailed time analysis. Assume the function f we choose satisfies two properties: **(1)** $f(m, \varepsilon) \leq m/2$ for any m , and **(2)** $f(\Theta(m), \varepsilon) = \Theta(f(m, \varepsilon))$.

First, we discuss how to construct \mathcal{D}_{new} . Constructing the data structure \mathcal{A} takes $\tilde{O}(n_0)$ time, as it is basic. The portions J_1, \dots, J_r can be computed in $\tilde{O}(n_0)$ time by sorting the points in S and the endpoints of the intervals in \mathcal{I} . Once J_1, \dots, J_r are computed, we build a (static) point location data structure \mathcal{B}_1 which can report in $O(\log r)$ time, for a given point $a \in \mathbb{R}$, the portion J_i that contains a . Clearly, \mathcal{B}_1 can be constructed in $\tilde{O}(r)$ time. With \mathcal{B}_1 in hand, we can determine in $\tilde{O}(1)$ time, for each point $a \in S$ (resp., each interval $I \in \mathcal{I}$), the portion J_i that contains a (resp., the two J_i 's that contain the endpoints of I). By doing this for all points in S and all intervals in \mathcal{I} , we obtain all S_i and all \mathcal{I}_i in $\tilde{O}(n_0 + r)$ time. After this, we can build the data structures $\mathcal{D}_{\text{old}}^{(i)}$'s. Constructing each $\mathcal{D}_{\text{old}}^{(i)}$ takes $\tilde{O}(f(n_0, \varepsilon))$ time since $|S_i| + |\mathcal{I}_i| = O(f(n_0, \varepsilon))$. Hence, the time for constructing all $\mathcal{D}_{\text{old}}^{(1)}, \dots, \mathcal{D}_{\text{old}}^{(r)}$ is $\tilde{O}(n_0)$.

The support data structure \mathcal{B}_1 will be used later in the implementation of the update procedure of \mathcal{D}_{new} (we do not need to update \mathcal{B}_1 since it is static). Besides, we need another support data structure \mathcal{B}_2 defined as follows.

Lemma 30. *One can store \mathcal{I} in a basic data structure \mathcal{B}_2 such that given an interval J , an interval in \mathcal{I} that contains J can be reported in $\tilde{O}(1)$ time (if it exists).*

Proof. We store \mathcal{I} in a binary search tree \mathcal{T} where the key of each interval is its left endpoint. We augment each node $\mathbf{u} \in \mathcal{T}$ with an additional field which stores the interval in the subtree rooted at \mathbf{u} that has the rightmost right endpoint. Given a query interval J , we first look for the interval in \mathcal{T} with the rightmost right endpoint whose key is smaller than or equal to the left endpoint of J . With the augmented fields, this interval can be found in $\tilde{O}(1)$ time. If this interval contains J , then we report it; otherwise, no interval in \mathcal{I} contains J . Clearly, \mathcal{T} can be built in $\tilde{O}(n_0)$ time and dynamized with $\tilde{O}(1)$ update time, and thus \mathcal{T} is basic. \square

Since \mathcal{B}_2 is basic, it can be constructed in $\tilde{O}(n_0)$ time and updated in $\tilde{O}(1)$ time. We conclude that the construction time of \mathcal{D}_{new} is $\tilde{O}(n_0)$.

Next, we consider how to implement the update procedure of \mathcal{D}_{new} . After each operation, we need to update the data structure \mathcal{A} and the support data structure \mathcal{B}_2 , which can be done in $\tilde{O}(1)$ time since they are basic. Also, if some (S_i, \mathcal{I}_i) changes, we need to update the data structure $\mathcal{D}_{\text{old}}^{(i)}$. An operation on S changes exactly one S_i and an operation on \mathcal{I} changes at most two \mathcal{I}_i 's. Thus, we only need to update at most two $\mathcal{D}_{\text{old}}^{(i)}$'s, and by using \mathcal{B}_1 we can find these $\mathcal{D}_{\text{old}}^{(i)}$'s in $\tilde{O}(1)$ time. Note that the size of each (S_i, \mathcal{I}_i) is bounded by $O(f(n_0, \varepsilon))$ during the first period (i.e., before the first reconstruction), because the period only consists of $f(n_0, \varepsilon)$ operations. Thus, updating the \mathcal{D}_{old} data structures takes $O(f^\alpha(n_0, \varepsilon)/\varepsilon^{1-\alpha})$ amortized time.

Then we discuss the maintenance of the solution. The time for simulating the output-sensitive algorithm is $\tilde{O}(\delta)$, i.e., $\tilde{O}(\min\{r/\varepsilon, n\})$. If the algorithm gives the solution $\mathcal{I}_{\text{appx}}$, we compute $|\mathcal{I}_{\text{appx}}|$ and store $\mathcal{I}_{\text{appx}}$ in a binary search tree; by doing this, we can answer the size, membership, and reporting queries for $\mathcal{Q}_{\text{appx}}$ in the required query times. This step takes $\tilde{O}(\delta)$ time, i.e., $\tilde{O}(\min\{r/\varepsilon, n\})$ time, since $|\mathcal{I}_{\text{appx}}| \leq \delta$ in this case. If the output-sensitive algorithm fails, we compute the sets P and P' . This can be done in $\tilde{O}(r)$ time by using the support data structure \mathcal{B}_2 . After this, we compute \mathcal{I}^* , which again takes $\tilde{O}(r)$ time by using \mathcal{B}_2 ; specifically, we consider each $i \in P$ and use \mathcal{B}_2 to find an interval in \mathcal{I} that contains J_i . We have $\mathcal{I}_{\text{appx}} = \mathcal{I}^* \sqcup (\bigsqcup_{i \in P'} \mathcal{I}_i^*)$. To support the size query for $\mathcal{I}_{\text{appx}}$ in $O(1)$ time, we need to compute $|\mathcal{I}_{\text{appx}}| = |\mathcal{I}^*| + \sum_{i \in P'} |\mathcal{I}_i^*|$. This can be done in $O(r)$ time, because we can query $\mathcal{D}_{\text{old}}^{(i)}$ to obtain $|\mathcal{I}_i^*|$ in $O(1)$ time. In order to support the membership and reporting queries, we store \mathcal{I}^* in a binary search trees \mathcal{T} . Also, we store the set P' , and for each $i \in P'$ we store a pointer pointing to the data structure $\mathcal{D}_{\text{old}}^{(i)}$. Consider a membership query $I \in \mathcal{I}$. Using the binary search tree \mathcal{T} , we can obtain the multiplicity of I in \mathcal{I}^* in $O(\log |\mathcal{I}^*|)$ time. Then we use \mathcal{B}_1 to find in $O(\log r)$ time the (at most) two \mathcal{I}_i 's that contain I (say $I \in \mathcal{I}_i$ and $I \in \mathcal{I}_j$). By querying $\mathcal{D}_{\text{old}}^{(i)}$ and $\mathcal{D}_{\text{old}}^{(j)}$, we know the multiplicities of I in \mathcal{I}_i^* and \mathcal{I}_j^* , which takes $O(\log |\mathcal{I}_i^*| + \log |\mathcal{I}_j^*|)$. The sum of these multiplicities is just the multiplicity of I in $\mathcal{I}_{\text{appx}}$. The time for answering the query is $O(\log |\mathcal{I}^*| + \log |\mathcal{I}_i^*| + \log |\mathcal{I}_j^*| + \log r)$. Note that $\log |\mathcal{I}^*| + \log |\mathcal{I}_i^*| + \log |\mathcal{I}_j^*| = O(\log |\mathcal{I}_{\text{appx}}|)$. Furthermore, because $|\mathcal{I}_{\text{appx}}| > \delta$ (as the output-sensitive algorithm fails), we have $|\mathcal{I}_{\text{appx}}| = \Omega(r)$. Thus, the time cost for a membership query is $O(\log |\mathcal{I}_{\text{appx}}|)$. Finally, consider the reporting query. We first use the binary search tree \mathcal{T} to report the elements in \mathcal{I}^* , using $O(|\mathcal{I}^*|)$ time. Then for each $i \in P'$, we query the data structure $\mathcal{D}_{\text{old}}^{(i)}$ to report the elements in \mathcal{I}_i^* . The total time cost is $O(|\mathcal{I}^*| + \sum_{i \in P'} |\mathcal{I}_i^*| + |P'|)$. Since $|P'| \leq r$ and $|\mathcal{I}_{\text{appx}}| = \Omega(r)$ as argued before, the time for answering the reporting query is $O(|\mathcal{I}_{\text{appx}}|)$. The above work for storing $\mathcal{I}_{\text{appx}}$ takes $\tilde{O}(r)$ time, since $|\mathcal{I}^*| = O(r)$ and $|P'| = O(r)$. To summarize, maintaining the solution takes $\tilde{O}(\min\{r/\varepsilon, n\} + r)$ time.

After processing $f(n_0, \varepsilon)$ operations, we need to reconstruct the entire data structure \mathcal{D}_{new} . The reconstruction is the same as the initial construction, except that n_0 is replaced with n_1 , the size of (S, \mathcal{I}) at the time of reconstruction. Thus, the reconstruction takes $\tilde{O}(n_1)$ time. We amortize the time cost over all the $f(n_0, \varepsilon)$ operations in the period. Since $n_1 \leq n_0 + f(n_0, \varepsilon)$, the amortized time for reconstruction is $\tilde{O}(n_0/f(n_0, \varepsilon))$, i.e., $\tilde{O}(r)$.

Combining the time for updating the \mathcal{D}_{old} data structures, the time for maintaining the solution, and the time for reconstruction, we see that the amortized update time of \mathcal{D}_{new} is $\tilde{O}(f^\alpha(n_0, \varepsilon)/\varepsilon^{1-\alpha} + \min\{r/\varepsilon, n\} + r)$ during the first period (since \mathcal{D}_{new} is reconstructed periodically, it suffices to analyze the update time in the first period). By property **(1)** of f , we have $n = \Theta(n_0)$ at any time in the period, i.e., the size of (S, \mathcal{I}) is $\Theta(n_0)$ at any time in the period. By property **(2)** of f , we further have $f(n, \varepsilon) = \Theta(f(n_0, \varepsilon))$ at any time in the period. It follows that the amortized update time of \mathcal{D}_{new} is $\tilde{O}(f^\alpha(n, \varepsilon)/\varepsilon^{1-\alpha} + \min\{n/(f(n, \varepsilon) \cdot \varepsilon), n\} + n/f(n, \varepsilon))$ during the period. To minimize the time complexity while guaranteeing the two conditions of f , we set $f(n, \varepsilon) = \min\{n^{1-\alpha'}/\varepsilon^{\alpha'}, n/2\}$ where α' is as defined in Theorem 2, i.e., $\alpha' = \alpha/(1+\alpha)$. The following lemma shows that our choice of f makes the time bound $\tilde{O}(n^{\alpha'}/\varepsilon^{1-\alpha'})$.

Lemma 31. *When $f(n, \varepsilon) = \min\{n^{1-\alpha'}/\varepsilon^{\alpha'}, n/2\}$, we have*

$$\frac{f^\alpha(n, \varepsilon)}{\varepsilon^{1-\alpha}} + \min\left\{\frac{n}{f(n, \varepsilon) \cdot \varepsilon}, n\right\} + \frac{n}{f(n, \varepsilon)} = O\left(\frac{n^{\alpha'}}{\varepsilon^{1-\alpha'}}\right).$$

Proof. If $f(n, \varepsilon) = n^{1-\alpha'}/\varepsilon^{\alpha'}$, then one can easily verify the equation in the lemma via a direct computation (by bounding each of the three terms on the left-hand side). It suffices to verify the equation for the case $f(n, \varepsilon) = n/2$. In this case, we have $n^{1-\alpha'}/\varepsilon^{\alpha'} \geq n/2$, implying that $n = O(1/\varepsilon)$. It follows that $n^\alpha/\varepsilon^{1-\alpha} = O(n^{\alpha'}/\varepsilon^{1-\alpha'})$. So the first term in the left-hand side is bounded by $O(n^{\alpha'}/\varepsilon^{1-\alpha'})$. The second term is $O(\min\{1/\varepsilon, n\})$, which is bounded by $O(n^{\alpha'}/\varepsilon^{1-\alpha'})$. The third term is clearly $O(1)$. This proves the equation in the lemma. \square

C Implementation details and detailed time analysis for the dynamic quadrant set cover data structure

We present the implementation details of our dynamic quadrant set cover data structure \mathcal{D}_{new} in Section 4.2 as well as a detailed time analysis. Since we are interested in the asymptotic bounds, we may assume that the approximation factor ε is sufficiently small, say $\varepsilon < 1$. Assume the function f we choose satisfies two properties: **(1)** $\sqrt{m}/2 \leq f(m, \varepsilon) \leq m/2$ for any m , and **(2)** $f(\Theta(m), \varepsilon) = \Theta(f(m, \varepsilon))$. Note that property **(1)** implies that $r = \lceil n_0/f(n_0, \varepsilon) \rceil = O(f(n_0, \varepsilon))$ and $r^2 = O(n_0)$.

First, we discuss how to construct \mathcal{D}_{new} . Constructing the data structure \mathcal{A} takes $\tilde{O}(n_0)$ time, as it is basic. The grid can be built in $\tilde{O}(n_0)$ time by sorting the points in S and the vertices of the quadrants in \mathcal{Q} . Once the grid is computed, we build a (static) point location data structure \mathcal{B}_1 which can report in $O(\log r)$ time, for a given point $a \in \mathbb{R}^2$, the grid cell that contains a . Since the grid has r^2 cells, \mathcal{B}_1 can be built in $\tilde{O}(r^2)$ time. With \mathcal{B}_1 in hand, we can determine in $\tilde{O}(1)$ time, for each point $a \in S$ (resp., each quadrant $Q \in \mathcal{Q}$), the cell contains a (resp., the vertex of Q). By doing this for all points in S and all quadrants in \mathcal{Q} , we obtain all $S_{i,j}$ and the non-special quadrants in all $\mathcal{Q}_{i,j}$ in $\tilde{O}(n_0 + r^2)$ time, i.e., $\tilde{O}(n_0)$ time. In order to compute the special quadrants, we need the following support data structure \mathcal{B}_2 .

Lemma 32. *One can store \mathcal{Q} in a basic data structure \mathcal{B}_2 such that given an orthogonal rectangle \square , the leftmost (resp., rightmost) quadrant in \mathcal{Q} that right (resp., left) intersects \square and the topmost (resp., bottommost) quadrant in \mathcal{Q} that bottom (resp., top) intersects \square can be reported in $\tilde{O}(1)$ time (if they exist).*

Proof. It suffices to show how to compute the leftmost quadrant in \mathcal{Q} that right intersects a given orthogonal rectangle \square . Note that only southeast and northeast quadrants can left intersects a rectangle, and without loss of generality, it suffices to see how to find the leftmost southeast quadrant in \mathcal{Q} that right intersects \square . Let $\mathcal{Q}^{\text{SE}} \subseteq \mathcal{Q}$ consist of the southeast quadrants. We store the vertices of the quadrants in \mathcal{Q}^{SE} in a basic 3-sided range-minimum data structure (Lemma 36) with $\tilde{O}(1)$ query time, by setting the weight of each vertex to be its x -coordinate. Let $\square = [x_1, x_2] \times [y_1, y_2]$ be a given orthogonal rectangle. Observe that a quadrant in \mathcal{Q}^{SE} right intersects \square iff its vertex lies in the 3-sided rectangle $(x_1, x_2] \times [y_2, \infty)$. Thus, the leftmost quadrant in \mathcal{Q}^{SE} that right intersects \square corresponds to the lightest vertex (i.e., the vertex with the smallest weight) in $(x_1, x_2] \times [y_2, \infty)$, which can be reported by the 3-sided range-minimum data structure in $\tilde{O}(1)$ time. \square

Using the support data structure \mathcal{B}_2 , we can find the special quadrants in all $\mathcal{Q}_{i,j}$ in $\tilde{O}(r^2)$ time. After this, we can build the data structures $\mathcal{D}_{\text{old}}^{(i,j)}$'s. Constructing each $\mathcal{D}_{\text{old}}^{(i,j)}$ takes $\tilde{O}(|S_{i,j}| + |\mathcal{Q}_{i,j}|)$ time. By Lemma 10, the time for constructing all $\mathcal{D}_{\text{old}}^{(i,j)}$'s is $\tilde{O}(n_0 + r^2)$. Therefore, the entire construction time of \mathcal{D}_{new} is $\tilde{O}(n_0 + r^2)$, i.e., $\tilde{O}(n_0)$.

The support data structures \mathcal{B}_1 and \mathcal{B}_2 will be used later in the implementation of the update procedure of \mathcal{D}_{new} . Thus, \mathcal{B}_2 will be updated after each operation (while \mathcal{B}_1 is static). Besides \mathcal{B}_1 and \mathcal{B}_2 , we need another support data structure \mathcal{B}_3 defined as follows.

Lemma 33. *One can store \mathcal{Q} in a basic data structure \mathcal{B}_3 such that given an orthogonal rectangle \square , a quadrant $Q \in \mathcal{Q}$ that contains \square can be reported in $\tilde{O}(1)$ time (if it exists).*

Proof. It suffices to consider the southeast quadrants. Let $\mathcal{Q}^{\text{SE}} \subseteq \mathcal{Q}$ consist of the southeast quadrants. We store \mathcal{Q}^{SE} in a binary search tree \mathcal{T} where the key of a quadrant is the x -coordinate of its vertex. We augment each node $\mathbf{u} \in \mathcal{T}$ with an additional field which stores the topmost quadrant in the subtree rooted at \mathbf{u} . Given an orthogonal rectangle $\square = [x_1, x_2] \times [y_1, y_2]$, we first look for the topmost quadrant Q whose key is smaller than or equal to x_1 . With the augmented fields, Q can be found in $\tilde{O}(1)$ time using \mathcal{T} . If Q contains \square , then we report Q , otherwise no quadrant in \mathcal{Q}^{SE} contains \square (because a southeast quadrant contains \square iff the x -coordinate of its vertex is smaller than or equal to x_1 and the y -coordinate of its vertex is greater than or equal to y_2). Clearly, \mathcal{T} can be built in $\tilde{O}(n_0)$ time and dynamized with $\tilde{O}(1)$ update time, and thus \mathcal{T} is basic. \square

Next, we consider how to implement the update procedure of \mathcal{D}_{new} . We first discuss the update of the \mathcal{D}_{old} data structures. If the operation is an insertion or deletion on S , then we use \mathcal{B}_1 to find the cell $\square_{i,j}$ that contains the inserted/deleted point and update the data structure $\mathcal{D}_{\text{old}}^{(i,j)}$. The situation is more complicated when the operation happens on \mathcal{Q} . Suppose the operation inserts a quadrant Q to \mathcal{Q} . Without loss of generality, assume Q is a southeast quadrant. Using \mathcal{B}_1 , we can find the cell $\square_{i,j}$ that contains the vertex of Q . We then update the data structure $\mathcal{D}_{\text{old}}^{(i,j)}$ by inserting Q to $\mathcal{Q}_{i,j}$. Besides, the insertion of Q may also change $\mathcal{Q}_{i,k}$ for $k \in \{j+1, \dots, r\}$ and $\mathcal{Q}_{k,j}$ for $k \in \{i+1, \dots, r\}$. Fix an index $k \in \{j+1, \dots, r\}$. The quadrant Q bottom intersects $\square_{i,k}$. We use the support data structure \mathcal{B}_2 to find the topmost quadrant Q' in \mathcal{Q} (before the insertion of Q) that bottom intersects $\square_{i,k}$. Then Q' is a special quadrant in $\mathcal{Q}_{i,k}$. If the y -coordinate of the vertex of Q is greater than the y -coordinate of the vertex of Q' , then we delete Q' from $\mathcal{Q}_{i,k}$ and insert Q to $\mathcal{Q}_{i,k}$, and we update $\mathcal{D}_{\text{old}}^{(i,k)}$ twice for these two operations. The data structures $\mathcal{D}_{\text{old}}^{(k,j)}$'s can be updated similarly. Updating each $\mathcal{D}_{\text{old}}^{(i,k)}$ takes $\tilde{O}(m_{i,k}^\alpha/\varepsilon^{1-\alpha})$ amortized time, where $m_{i,k}$ is the size of the current $(S_{i,k}, \mathcal{Q}_{i,k})$, and updating each $\mathcal{D}_{\text{old}}^{(k,j)}$ takes $\tilde{O}(m_{k,j}^\alpha/\varepsilon^{1-\alpha})$ amortized time. Therefore, the total amortized time cost for updating these \mathcal{D}_{old} data structures is bounded by $\tilde{O}(\sum_{k=1}^r m_{i,k}^\alpha/\varepsilon^{1-\alpha} + \sum_{k=1}^r m_{k,j}^\alpha/\varepsilon^{1-\alpha})$. By Lemma 10 and the fact $r = O(f(n_0, \varepsilon))$, we have $\sum_{k=1}^r m_{i,k} = O(f(n_0, \varepsilon))$ and $\sum_{k=1}^r m_{k,j} = O(f(n_0, \varepsilon))$ at any time of the first period. Since $\alpha \leq 1$, by Hölder's inequality and Lemma 10,

$$\sum_{k=1}^r m_{i,k}^\alpha \leq \left(\frac{\sum_{k=1}^r m_{i,k}}{r} \right)^\alpha \cdot r = O(r^{1-\alpha} \cdot f^\alpha(n_0, \varepsilon))$$

and similarly $\sum_{k=1}^r m_{k,j}^\alpha = O(r^{1-\alpha} \cdot f^\alpha(n_0, \varepsilon))$. It follows that updating the \mathcal{D}_{old} data structures takes $\tilde{O}(r^{1-\alpha} \cdot f^\alpha(n_0, \varepsilon)/\varepsilon^{1-\alpha})$ amortized time. Updating the data structure \mathcal{A} and the support data structures \mathcal{B}_2 and \mathcal{B}_3 can be done in $\tilde{O}(1)$ time since they are basic.

Then we discuss the maintenance of the solution. The time for simulating the output-sensitive algorithm is $\tilde{O}(\mu \cdot \delta)$, i.e., $\tilde{O}(\min\{r^2/\varepsilon, n\})$. If the algorithm gives the solution $\mathcal{Q}_{\text{appx}}$, we compute $|\mathcal{Q}_{\text{appx}}|$ and store $\mathcal{Q}_{\text{appx}}$ in a binary search tree; by doing this, we can answer the size, membership, and reporting queries for $\mathcal{Q}_{\text{appx}}$ in the required query times. This step takes $\tilde{O}(\mu \cdot \delta)$ time, i.e., $\tilde{O}(\min\{r^2/\varepsilon, n\})$ time, since $|\mathcal{Q}_{\text{appx}}| \leq \mu \cdot \delta$ in this case. If the output-sensitive algorithm fails, we compute the sets P and P' . This can be done in $\tilde{O}(r^2)$ time by using the support data structure \mathcal{B}_3 . After this, we compute \mathcal{Q}^* , which again takes $\tilde{O}(r^2)$ time by using \mathcal{B}_3 ; specifically, we consider each $(i, j) \in P$ and use \mathcal{B}_3 to find a quadrant in \mathcal{Q} that contains $\square_{i,j}$. We have $\mathcal{Q}_{\text{appx}} = \mathcal{Q}^* \sqcup (\bigsqcup_{(i,j) \in P'} \mathcal{Q}_{i,j}^*)$. To support the size query for $\mathcal{Q}_{\text{appx}}$ in $O(1)$ time, we need to compute $|\mathcal{Q}_{\text{appx}}| = |\mathcal{Q}^*| + \sum_{(i,j) \in P'} |\mathcal{Q}_{i,j}^*|$. This can be done in $O(r^2)$ time, because we can query $\mathcal{D}_{\text{old}}^{(i,j)}$ to obtain $|\mathcal{Q}_{i,j}^*|$ in $O(1)$ time. To support the reporting query in $O(|\mathcal{Q}_{\text{appx}}|)$ time, we only need to store \mathcal{Q}^* and P' , and store at each $(i, j) \in P'$ a pointer pointing to the data structure $\mathcal{D}_{\text{old}}^{(i,j)}$. In this way, we can report the quadrants in \mathcal{Q}^* and for each $(i, j) \in P'$, report the quadrants in $\mathcal{Q}_{i,j}^*$ in $O(|\mathcal{Q}_{i,j}^*|)$ time by querying $\mathcal{D}_{\text{old}}^{(i,j)}$. Supporting the membership query in $O(\log |\mathcal{Q}_{\text{appx}}|)$ time is more difficult, since a quadrant may belong to many $\mathcal{Q}_{i,j}^*$'s. To handle this issue, the idea is to collect all the special quadrants in the $\mathcal{Q}_{i,j}^*$'s. Specifically, let $\mathcal{P}_{i,j}^* \subseteq \mathcal{Q}_{i,j}^*$ consist of the (at most) four special quadrants. We can compute $\mathcal{P}_{i,j}^*$ for each $(i, j) \in P'$ in $\tilde{O}(1)$ time by first finding the (at most) four special quadrants in $\mathcal{Q}_{i,j}$ using \mathcal{B}_2 and computing the multiplicity of each special quadrant in $\mathcal{Q}_{i,j}^*$ by querying $\mathcal{D}_{\text{old}}^{(i,j)}$. We then store $\mathcal{Q}^* \sqcup (\bigsqcup_{(i,j) \in P'} \mathcal{P}_{i,j}^*)$ in a binary search tree \mathcal{T} . Given a query quadrant $Q \in \mathcal{Q}$, we first use \mathcal{T} to compute the multiplicity p_1 of Q in $\mathcal{Q}^* \sqcup (\bigsqcup_{(i,j) \in P'} \mathcal{P}_{i,j}^*)$ in $O(\log |\mathcal{Q}_{\text{appx}}|)$ time. Then we use \mathcal{B}_1 to find the cell $\square_{i,j}$ that contains the vertex of Q in $O(\log r)$ time and query $\mathcal{D}_{\text{old}}^{(i,j)}$ to find the multiplicity p_2 of Q in $\mathcal{Q}_{i,j}^*$ in $O(\log |\mathcal{Q}_{i,j}^*|)$ time. One can easily verify that $p_1 + p_2$ is the multiplicity of Q in $\mathcal{Q}_{\text{appx}}$. The query takes $\tilde{O}(\log |\mathcal{Q}_{\text{appx}}| + \log r + \log |\mathcal{Q}_{i,j}^*|)$ time. Note that $|\mathcal{Q}_{i,j}^*| \leq |\mathcal{Q}_{\text{appx}}|$ and $|\mathcal{Q}_{\text{appx}}| = \Omega(r)$, where the latter follow from the fact that $|\mathcal{Q}_{\text{appx}}| > \delta$ (as the output-sensitive algorithm fails). Therefore, we can support the membership query in $O(\log |\mathcal{Q}_{\text{appx}}|)$ time. The above work for storing $\mathcal{Q}_{\text{appx}}$ takes $\tilde{O}(r^2)$ time, since $|\mathcal{Q}^*| = O(r^2)$ and $|P'| = O(r^2)$. To summarize, maintaining the solution takes $\tilde{O}(\min\{r^2/\varepsilon, n\} + r^2)$ time.

After processing $f(n_0, \varepsilon)$ operations, we need to reconstruct the entire data structure \mathcal{D}_{new} . The reconstruction is the same as the initial construction, except that n_0 is replaced with n_1 , the size of (S, \mathcal{Q}) at the time of reconstruction. Thus, the reconstruction takes $\tilde{O}(n_1)$ time. We amortize the time cost over all the $f(n_0, \varepsilon)$ operations in the period. Since $n_1 \leq n_0 + f(n_0, \varepsilon)$, the amortized time for reconstruction is $\tilde{O}(n_0/f(n_0, \varepsilon))$, i.e., $\tilde{O}(r)$.

Combining the time for updating the \mathcal{D}_{old} data structures, the time for maintaining the solution, and the time for reconstruction, we see that the amortized update time of \mathcal{D}_{new} is $\tilde{O}(r^{1-\alpha} \cdot f^\alpha(n_0, \varepsilon)/\varepsilon^{1-\alpha} + \min\{r^2/\varepsilon, n\} + r^2)$ during the first period (since \mathcal{D}_{new} is reconstructed periodically, it suffices to analyze the update time in the first period). By property **(1)** of f , we have $n = \Theta(n_0)$ at any time in the period, i.e., the size of (S, \mathcal{Q}) is $\Theta(n_0)$ at any time in the period. By property **(2)** of f , we further have $f(n, \varepsilon) = \Theta(f(n_0, \varepsilon))$ at any time in the period. It follows that the amortized update time of \mathcal{D}_{new} is $\tilde{O}(n^{1-\alpha}/(f^{1-2\alpha}(n, \varepsilon) \cdot \varepsilon^{1-\alpha}) + \min\{n^2/(f^2(n, \varepsilon) \cdot \varepsilon), n\} + n^2/f^2(n, \varepsilon))$. To minimize the time complexity while guaranteeing the two conditions of f , we set $f(n, \varepsilon) = \min\{n^{1-\alpha'/2}/(\sqrt{\varepsilon})^{\alpha'}, n/2\}$ where α' is as defined in Theorem 7, i.e., $\alpha' = 2\alpha/(1+2\alpha)$. Note that by doing this we have $f(n, \varepsilon) \geq \sqrt{n}$ because of our assumption $\varepsilon < 1$. The following lemma shows that our choice of f makes the time bound $\tilde{O}(n^{\alpha'}/\varepsilon^{1-\alpha'})$.

Lemma 34. *When $f(n, \varepsilon) = \min\{n^{1-\alpha'/2}/(\sqrt{\varepsilon})^{\alpha'}, n/2\}$, we have*

$$\frac{n^{1-\alpha}}{f^{1-2\alpha}(n, \varepsilon) \cdot \varepsilon^{1-\alpha}} + \min\left\{\frac{n^2}{f^2(n, \varepsilon) \cdot \varepsilon}, n\right\} + \frac{n^2}{f^2(n, \varepsilon)} = O\left(\frac{n^{\alpha'}}{\varepsilon^{1-\alpha'}}\right).$$

Proof. If $f(n, \varepsilon) = n^{1-\alpha'/2}/(\sqrt{\varepsilon})^{\alpha'}$, then one can easily verify the equation in the lemma via a direct computation (by bounding each of the three terms on the left-hand side). It suffices to verify the equation for the case $f(n, \varepsilon) = n/2$. In this case, we have $n^{1-\alpha'/2}/(\sqrt{\varepsilon})^{\alpha'} \geq n/2$, implying that $n = O(1/\varepsilon)$. It follows that $n^{\alpha'}/\varepsilon^{1-\alpha'} = O(n^{\alpha'}/\varepsilon^{1-\alpha'})$. So the first term in the left-hand side is bounded by $O(n^{\alpha'}/\varepsilon^{1-\alpha'})$. The second term is $O(\min\{1/\varepsilon, n\})$, which is bounded by $O(n^{\alpha'}/\varepsilon^{1-\alpha'})$. The third term is clearly $O(1)$. This proves the equation in the lemma. \square

D Missing details in the output-sensitive quadrant set cover algorithm

D.1 Handling the no-solution case

Let $U = \bigcup_{Q \in \mathcal{Q}} Q$. Observe that no matter whether (S, \mathcal{Q}) has no set cover or not, the algorithm in Section 4.2.2 gives an $O(1)$ -approximate optimal set cover \mathcal{Q}^* for $(S \cap U, \mathcal{Q})$. Thus, in order to handle the no-solution case, we only need to check whether \mathcal{Q}^* covers all points in S , after \mathcal{Q}^* is computed. Define $U^* = \bigcup_{Q \in \mathcal{Q}^*} Q$. Note that \mathcal{Q}^* is a set cover for S iff $\mathbb{R}^2 \setminus U^*$ does not contain any point in S . The area $\mathbb{R}^2 \setminus U^*$ is a rectilinear domain (not necessarily connected). Since U^* is the union of $O(|\mathcal{Q}^*|)$ quadrants, the complexity of U^* is $O(|\mathcal{Q}^*|)$, so is the complexity of $\mathbb{R}^2 \setminus U^*$. Furthermore, it is easy to compute $\mathbb{R}^2 \setminus U^*$ and decompose it into $O(|\mathcal{Q}^*|)$ rectangles in $\tilde{O}(|\mathcal{Q}^*|)$ time, given \mathcal{Q}^* in hand. We then only need to test for each such rectangle R whether R contains any point in S or not, which can be done via an orthogonal range-emptiness query on S ; there are existing basic data structures that support orthogonal range-emptiness queries in $\tilde{O}(1)$ time [24].

D.2 Implementing the algorithm using basic data structures

In this section, we show that the output-sensitive quadrant set cover algorithm can be performed in $\tilde{O}(\text{opt})$ time by storing the instance (S, \mathcal{Q}) in some basic data structure. As argued in Section 4.2.2, it suffices to show how to implement the following operations in $\tilde{O}(1)$ time using basic data structures (please refer to Section 4.2.2 for the notations).

- Computing the point σ .
- Given a point $a \in \mathbb{R}^2$, testing whether $a \in U^{\text{NE}}$ and $a \in U^{\text{NW}}$.
- Given a point $a \in \mathbb{R}^2$, computing the quadrants $\Phi_{\rightarrow}(a, \mathcal{Q}^{\text{SW}})$, $\Phi_{\rightarrow}(a, \mathcal{Q}^{\text{NW}})$, $\Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}})$, and $\Phi_{\uparrow}(a, \mathcal{Q}^{\text{NE}})$.
- Given a number \tilde{y} , computing the point $\phi(\tilde{y})$.

Computing σ . Recall that γ is the boundary of U^{SE} , which is a staircase curve from bottom-left to top-right. The area U^{SW} contains the bottom-left end of γ (if $\gamma \cap U^{\text{SW}} \neq \emptyset$). The point σ is the “endpoint” of $\gamma \cap U^{\text{SW}}$, i.e., the point on γ closest to the top-right end of γ that is contained in U^{SW} . We define the *height* of U^{SW} at $x \in \mathbb{R}$, denoted by $\text{ht}_x(U^{\text{SW}})$, as the largest number $y \in \mathbb{R}$ such that the point with coordinates (x, y) is contained in U^{SW} ; similarly, we can define the height of U^{SE} at $x \in \mathbb{R}$, denoted by $\text{ht}_x(U^{\text{SE}})$. We say U^{SW} is higher than U^{SE} at $x \in \mathbb{R}$ if $\text{ht}_x(U^{\text{SW}}) \geq \text{ht}_x(U^{\text{SE}})$; otherwise, we say U^{SW} is lower than U^{SE} at x . It is easy to see the following three facts about the point σ .

1. The x -coordinate of σ , denoted by x_σ , is equal to $x(Q)$ for some $Q \in \mathcal{Q}^{\text{SE}} \cup \mathcal{Q}^{\text{SW}}$; recall that $x(Q)$ is the x -coordinate of the vertex of Q .
2. U^{SW} is higher than U^{SE} at all $x < x_\sigma$ and U^{SW} is lower than U^{SE} at all $x > x_\sigma$.
3. The coordinates of σ is $(x_\sigma, \min\{\text{ht}_{x_\sigma}(U^{\text{SW}}), \text{ht}_{x_\sigma}(U^{\text{SE}})\})$.

Based on these facts, we compute σ as follows. First, we need a basic data structure built on \mathcal{Q}^{SW} (resp., \mathcal{Q}^{SE}) that can compute the height function in $\tilde{O}(1)$ time.

Lemma 35. *One can store \mathcal{Q}^{SW} (resp., \mathcal{Q}^{SE}) in a basic data structure which can report $\text{ht}_x(U^{\text{SW}})$ (resp., $\text{ht}_x(U^{\text{SE}})$) for a given number $x \in \mathbb{R}$ in $\tilde{O}(1)$ time.*

Proof. It suffices to consider \mathcal{Q}^{SW} . We store \mathcal{Q}^{SW} in a standard binary search tree \mathcal{T} , by using $x(Q)$ as the key of each quadrant $Q \in \mathcal{Q}^{\text{SW}}$. At each node $\mathbf{u} \in \mathcal{T}$, we store a field $Y(\mathbf{u})$ which is the maximum of $y(Q)$ for all quadrants Q in the subtree rooted at \mathbf{u} . Clearly, \mathcal{T} can be constructed in $\tilde{O}(n)$ time where $n = |\mathcal{Q}^{\text{SW}}|$ and can be dynamized with $\tilde{O}(1)$ update time to support insertions and deletions on \mathcal{Q}^{SW} , hence it is a basic data structure. Next, we consider how to compute $\text{ht}_x(U^{\text{SW}})$ for a given $x \in \mathbb{R}$ using \mathcal{T} . One can easily see that $\text{ht}_x(U^{\text{SW}}) = \max\{y(Q) : Q \in \mathcal{Q}^{\text{SW}} \text{ with } x(Q) \geq x\}$. In other words, $\text{ht}_x(U^{\text{SW}})$ is the maximum of $y(Q)$ for all $Q \in \mathcal{Q}^{\text{SW}}$ corresponding to the nodes in \mathcal{T} whose keys are at least x . Therefore, using the field $Y(\mathbf{u})$, $\text{ht}_x(U^{\text{SW}})$ can be computed in $O(\log n)$ time simply via a top-down walk in \mathcal{T} . The walk begins at the root of \mathcal{T} , and we set $\text{ht} = -\infty$ initially. If the key of the current node is smaller than x , then we just go to its right child. Otherwise, we update ht as $\text{ht} \leftarrow \max\{\text{ht}, y(Q), Y(\mathbf{r})\}$ where $Q \in \mathcal{Q}^{\text{SW}}$ is the quadrant corresponding to the current node and \mathbf{r} is the right child of the current node (if the current node has no right child, we update ht as $\text{ht} \leftarrow \max\{\text{ht}, y(Q)\}$), and go to the left child of the current node. When the walk ends at a leaf node, the number ht is just equal to $\text{ht}_x(U^{\text{SW}})$. The walk takes $O(\log n)$ time, hence $\text{ht}_x(U^{\text{SW}})$ can be computed in $\tilde{O}(1)$ time using \mathcal{T} . This completes the proof of the lemma. \square

With the above lemma in hand, we may now assume that the height functions $\text{ht}_x(U^{\text{SW}})$ and $\text{ht}_x(U^{\text{SE}})$ can be computed in $\tilde{O}(1)$ time for any $x \in \mathbb{R}$. In particular, we can test in $\tilde{O}(1)$ time whether U^{SW} is higher or lower than U^{SE} at any $x \in \mathbb{R}$. We then build a binary search tree \mathcal{T} on $\mathcal{Q}^{\text{SW}} \cup \mathcal{Q}^{\text{SE}}$, by using $x(Q)$ as the key of each quadrant $Q \in \mathcal{Q}^{\text{SW}} \cup \mathcal{Q}^{\text{SE}}$. Clearly, \mathcal{T} can be constructed in $\tilde{O}(n)$ time where $n = |\mathcal{Q}^{\text{SW}} \cup \mathcal{Q}^{\text{SE}}|$ and can be dynamized with $\tilde{O}(1)$ update time, hence it is a basic data structure. We observe that, using \mathcal{T} , we can determine in $\tilde{O}(1)$ time for any given number $p \in \mathbb{R}$ which one of the following three is true: **(1)** $p < x_\sigma$, **(2)** $p = x_\sigma$, **(3)** $p > x_\sigma$. To see this, consider a given number $p \in \mathbb{R}$. We first search in \mathcal{T} to see whether $p = x(Q)$ for some $Q \in \mathcal{Q}^{\text{SW}} \cup \mathcal{Q}^{\text{SE}}$. If not, we know $p \neq x_\sigma$, because $x_\sigma = x(Q)$ for some $Q \in \mathcal{Q}^{\text{SW}} \cup \mathcal{Q}^{\text{SE}}$ (as we observed before). In this case, we can decide whether $p < x_\sigma$ or $p > x_\sigma$ by simply testing whether U^{SW} is higher or lower than U^{SE} at p . Specifically, if U^{SW} is higher than U^{SE} at p , then

$p < x_\sigma$, otherwise $p > x_\sigma$, because U^{SW} is higher (resp., lower) than U^{SE} at all $x < x_\sigma$ (resp., $x > x_\sigma$) as we observed before. The remaining case is that $p = x(Q)$ for some $Q \in \mathcal{Q}^{\text{SW}} \cup \mathcal{Q}^{\text{SE}}$. In this case, we also test whether U^{SW} is higher or lower than U^{SE} at p ; by doing this, we can decide whether $p \leq x_\sigma$ or $p \geq x_\sigma$. Suppose $p \leq x_\sigma$. To see whether $p < x_\sigma$ or $p = x_\sigma$, we search in \mathcal{T} to find the smallest key p' that is larger than p . Let $\tilde{p} \in \mathbb{R}$ be any number such that $p < \tilde{p} < p'$. If U^{SW} is higher than U^{SE} at \tilde{p} , then we have $p < \tilde{p} \leq x_\sigma$. If U^{SW} is lower than U^{SE} at \tilde{p} , then we know $p \leq x_\sigma \leq \tilde{p}$. Note that any number in $(p, \tilde{p}]$ is not equal to $x(Q)$ for any $Q \in \mathcal{Q}^{\text{SW}} \cup \mathcal{Q}^{\text{SE}}$, due to the choice of p' and the inequality $p < \tilde{p} < p'$. Therefore, $x_\sigma \notin (p, \tilde{p}]$, which implies $p = x_\sigma$. To summarize, we can determine in $\tilde{O}(1)$ time for any $p \in \mathbb{R}$ which one of the following three is true: **(1)** $p < x_\sigma$, **(2)** $p = x_\sigma$, **(3)** $p > x_\sigma$.

This allows us to compute x_σ using a binary search manner. Specifically, we do a top-down walk from the root of \mathcal{T} . If the key of the current node is equal to x_σ , we are done. Otherwise, if the key of the current node is smaller (resp., larger) than x_σ , we go to its right (resp., left) child, because the keys of the nodes in the left (resp., right) subtree are all smaller (resp., larger) than x_σ . During the walk, we can definitely find a node whose key is x_σ , since $x_\sigma = x(Q)$ for some $Q \in \mathcal{Q}^{\text{SW}} \cup \mathcal{Q}^{\text{SE}}$, i.e., x_σ is the key of some node in \mathcal{T} . In this way, we can compute x_σ in $\tilde{O}(1)$ time. As we observed before, the coordinates of σ is $(x_\sigma, \min\{\text{ht}_{x_\sigma}(U^{\text{SW}}), \text{ht}_{x_\sigma}(U^{\text{SE}})\})$. Thus, once we know x_σ , it suffices to compute $\text{ht}_{x_\sigma}(U^{\text{SW}})$ and $\text{ht}_{x_\sigma}(U^{\text{SE}})$, which takes $\tilde{O}(1)$ time by Lemma 35. We conclude that computing σ can be done in $\tilde{O}(1)$ time (by properly store \mathcal{Q}^{SW} and \mathcal{Q}^{SE} in basic data structures).

Testing whether $a \in U^{\text{NE}}$ and $a \in U^{\text{NW}}$. It suffices to consider how to test whether $a \in U^{\text{NE}}$. We store \mathcal{Q}^{NE} in a binary search tree \mathcal{T} , by using $y(Q)$ as the key of each quadrant $Q \in \mathcal{Q}^{\text{NE}}$. We augment each node $\mathbf{u} \in \mathcal{T}$ with a field which stores the leftmost quadrant in the subtree rooted at \mathbf{u} . Clearly, \mathcal{T} can be constructed in $\tilde{O}(n)$ time where $n = |\mathcal{Q}^{\text{NE}}|$ and can be dynamized with $\tilde{O}(1)$ update time, hence it is basic. Given a point $a \in \mathbb{R}^2$, we first look for the leftmost quadrant Q in \mathcal{T} whose key is smaller than or equal to the y -coordinate of a . With the augmented fields, Q can be found in $\tilde{O}(1)$ time. If $a \in Q$, then we know $a \in U^{\text{NE}}$. Otherwise, we claim that $a \notin U^{\text{NE}}$. Indeed, a quadrant in \mathcal{T} contains a only if its key is smaller than or equal to the y -coordinate of a . Since Q is the leftmost one among such quadrants and $a \notin Q$, we know that a is not contained in any quadrant in \mathcal{T} , i.e., $a \notin U^{\text{NE}}$. We conclude that testing whether $a \in U^{\text{NE}}$ and $a \in U^{\text{NW}}$ for a given point $a \in \mathbb{R}^2$ can be done in $\tilde{O}(1)$ time (by properly store \mathcal{Q}^{NE} and \mathcal{Q}^{NW} in basic data structures).

Computing Φ . Recall that for a point $a \in \mathbb{R}^2$ and a collection \mathcal{P} of quadrants, $\Phi_{\rightarrow}(a, \mathcal{P})$ and $\Phi_{\uparrow}(a, \mathcal{P})$ denote the rightmost and topmost quadrants in \mathcal{P} that contain a , respectively. We want to compute $\Phi_{\rightarrow}(a, \mathcal{Q}^{\text{SW}})$, $\Phi_{\rightarrow}(a, \mathcal{Q}^{\text{NW}})$, $\Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}})$, and $\Phi_{\uparrow}(a, \mathcal{Q}^{\text{NE}})$ in $\tilde{O}(1)$ time for a given point $a \in \mathbb{R}^2$, using basic data structures. Here we only consider how to compute $\Phi_{\rightarrow}(a, \mathcal{Q}^{\text{SW}})$, the other three can be computed in the same way. We store \mathcal{Q}^{SW} in a binary search tree \mathcal{T} , by using $y(Q)$ as the key of each quadrant $Q \in \mathcal{Q}^{\text{SW}}$. At each node $\mathbf{u} \in \mathcal{T}$, we store a field that is the rightmost quadrant in the subtree rooted at \mathbf{u} . Clearly, \mathcal{T} can be constructed in $\tilde{O}(n)$ time where $n = |\mathcal{Q}^{\text{SW}}|$ and can be dynamized with $\tilde{O}(1)$ update time, hence it is a basic data structure. Given a point $a \in \mathbb{R}^2$, we first look for the rightmost quadrant Q in \mathcal{T} whose key is greater than or equal to the y -coordinate y_a of a . With the augmented fields, Q can be found in $\tilde{O}(1)$ time. If Q contains a , then Q is the rightmost quadrant in \mathcal{T} (i.e., in \mathcal{Q}^{SW}) that contains a , because any quadrant Q' that contains a must satisfy $y(Q') \geq y_a$. Otherwise, no quadrant in \mathcal{Q}^{SW} contains a . We conclude that computing $\Phi_{\rightarrow}(a, \mathcal{Q}^{\text{SW}})$, $\Phi_{\rightarrow}(a, \mathcal{Q}^{\text{NW}})$, $\Phi_{\uparrow}(a, \mathcal{Q}^{\text{SE}})$, and $\Phi_{\uparrow}(a, \mathcal{Q}^{\text{NE}})$ for a given point $a \in \mathbb{R}^2$ can be done in $\tilde{O}(1)$ time (by properly store \mathcal{Q}^{SW} , \mathcal{Q}^{NW} , \mathcal{Q}^{SE} , \mathcal{Q}^{NE} in basic data structures).

Computing $\phi(\tilde{y})$. Recall that for a number $\tilde{y} \in \mathbb{R}$, $\phi(\tilde{y})$ is the leftmost point in $S \cap U^{\text{SE}}$ whose y -coordinate is greater than \tilde{y} . We want to store S and \mathcal{Q}^{SE} in some basic data structure such that $\phi(\tilde{y})$ can be computed in $\tilde{O}(1)$ time for any given $\tilde{y} \in \mathbb{R}$. For simplicity of exposition, let us make a general-position assumption: the points in S and the vertices of the quadrants in \mathcal{Q}^{SE} have distinct y -coordinates. The first thing we need is a data structure that supports the so-called *3-sided range-minimum* query. A 3-sided range-minimum

query on a set of *weighted* points in \mathbb{R}^2 ask for the *lightest* point (i.e., the point with the smallest weight) contained in a given 3-sided query rectangle $R = [x_0, \infty) \times [y_1, y_2]$.

Lemma 36. *There exists a basic data structure that supports 3-sided range-minimum queries in $\tilde{O}(1)$ time.*

Proof. The standard range trees can answer *static* 3-sided range-minimum queries in $\tilde{O}(1)$ time, and can be constructed in $\tilde{O}(n)$ time, where n is the size of the dataset. Since range-minimum queries are decomposable, the approach of [4] can be applied to dynamize the static data structure with $\tilde{O}(1)$ update time, by paying an extra logarithmic factor in the query time. This gives us the basic data structure that supports 3-sided range-minimum queries in $\tilde{O}(1)$ time. \square

We store S in the basic 3-sided range-minimum data structure \mathcal{A} of the above lemma, by setting the weight of each point in S to be its x -coordinate. Besides \mathcal{A} , we need a (1D) range tree \mathcal{T} built on $S \cup V(\mathcal{Q}^{\text{SE}})$ for y -coordinates, where $V(\mathcal{Q}^{\text{SE}})$ is the set of the vertices of the quadrants in \mathcal{Q}^{SE} . By the definition of a range tree, the points in $S \cup V(\mathcal{Q}^{\text{SE}})$ are one-to-one corresponding to the leaves of \mathcal{T} , where the left-right order of the leaves corresponds to the small-large order of the y -coordinates of the points. The *canonical subset* of each node $\mathbf{u} \in \mathcal{T}$ refers to the subset of $S \cup V(\mathcal{Q}^{\text{SE}})$ consisting of the points stored in the subtree rooted at \mathbf{u} . For a node $\mathbf{u} \in \mathcal{T}$, we denote by $S(\mathbf{u})$ the set of the points in S that are contained in the canonical subset of \mathbf{u} , and denote by $\mathcal{Q}(\mathbf{u})$ the collection of the quadrants in \mathcal{Q}^{SE} whose vertices are contained in the canonical subset of \mathbf{u} . Also, we write $U(\mathbf{u}) = \bigcup_{Q \in \mathcal{Q}(\mathbf{u})} Q$. At each node $\mathbf{u} \in \mathcal{T}$, we store the following three fields.

- $y^-(\mathbf{u})$: the y -coordinate of the bottommost point in the canonical subset of \mathbf{u} .
- $y^+(\mathbf{u})$: the y -coordinate of the topmost point in the canonical subset of \mathbf{u} .
- $L(\mathbf{u})$: the leftmost quadrant in $\mathcal{Q}(\mathbf{u})$.
- $a(\mathbf{u})$: the leftmost point in $S(\mathbf{u})$ that is contained in $U(\mathbf{u})$.

Let $\mathbf{u} \in \mathcal{T}$ be a node and \mathbf{l}, \mathbf{r} be its left and right children, respectively. It is clear that $y^-(\mathbf{u})$, $y^+(\mathbf{u})$, $L(\mathbf{u})$ can be computed in $O(1)$ time knowing the $y^-(\cdot)$, $y^+(\cdot)$, and $L(\cdot)$ fields of \mathbf{l} and \mathbf{r} . We claim that $a(\mathbf{u})$ can be computed in $\tilde{O}(1)$ time based on the information stored at \mathbf{u} , \mathbf{l} , \mathbf{r} , and the 3-sided range-minimum data structure \mathcal{A} . By definition, $a(\mathbf{u})$ is the leftmost point in $S(\mathbf{u})$ that is contained in $U(\mathbf{u})$. Since $S(\mathbf{u}) = S(\mathbf{l}) \cup S(\mathbf{r})$, it suffices to compute the leftmost point in $S(\mathbf{l})$ contained in $U(\mathbf{u})$ and the leftmost point in $S(\mathbf{r})$ contained in $U(\mathbf{u})$. Note that any point in $S(\mathbf{r})$ is not contained in $U(\mathbf{l})$, since any quadrant in $\mathcal{Q}(\mathbf{l})$ is “below” any point in $S(\mathbf{r})$. It follows that the leftmost point in $S(\mathbf{r})$ that is contained in $U(\mathbf{u})$ is just $a(\mathbf{r})$. To compute the leftmost point in $S(\mathbf{l})$ that is contained in $U(\mathbf{u})$, we only need to compute the leftmost point in $S(\mathbf{l})$ contained in $U(\mathbf{l})$ and the leftmost point in $S(\mathbf{l})$ contained in $U(\mathbf{r})$, because $U(\mathbf{u}) = U(\mathbf{l}) \cup U(\mathbf{r})$. The leftmost point in $S(\mathbf{l})$ contained in $U(\mathbf{l})$ is just $a(\mathbf{l})$. In order to compute the leftmost point in $S(\mathbf{l})$ contained in $U(\mathbf{r})$, we observe that a point in $S(\mathbf{l})$ is contained in $U(\mathbf{r})$ iff it is contained in the quadrant $L(\mathbf{r})$. Thus, the leftmost point in $S(\mathbf{l})$ contained in $U(\mathbf{r})$ is just the leftmost point in $S(\mathbf{l})$ contained in $L(\mathbf{r})$. Note that $S(\mathbf{l})$ is exactly the set of the points in S that lie in the strip $P = \mathbb{R} \times [y^-(\mathbf{l}), y^+(\mathbf{l})]$, i.e., $S(\mathbf{l}) = S \cap P$. Hence, we can query the data structure \mathcal{A} with the 3-sided rectangle $P \cap L(\mathbf{r})$ to obtain the leftmost point in $S(\mathbf{l})$ contained in $L(\mathbf{r})$, i.e., the leftmost point in $S(\mathbf{l})$ contained in $U(\mathbf{r})$, which takes $\tilde{O}(1)$ time. We conclude that, by using the 3-sided range-minimum data structure \mathcal{A} , all the fields of a node \mathbf{u} can be computed in $\tilde{O}(1)$ time based on the information stored at \mathbf{u} and their children. Therefore, the range tree \mathcal{T} with the augmented fields can be dynamized with $\tilde{O}(1)$ update time using the standard technique for dynamizing augmented trees [12]. The construction time of \mathcal{T} is clearly $\tilde{O}(n)$ where $n = |S| + |\mathcal{Q}^{\text{SE}}|$, thus \mathcal{T} is a basic data structure.

Next, we consider how to use \mathcal{T} to compute $\phi(\tilde{y})$ in $\tilde{O}(1)$ time for a given $\tilde{y} \in \mathbb{R}$. We first find the $t = O(\log n)$ canonical nodes $\mathbf{u}_1, \dots, \mathbf{u}_t \in \mathcal{T}$ corresponding to the range $[\tilde{y}, \infty)$. Suppose $\mathbf{u}_1, \dots, \mathbf{u}_t$ are sorted from left to right in \mathcal{T} . By the property of canonical nodes, the canonical subsets of $\mathbf{u}_1, \dots, \mathbf{u}_t$ are disjoint and their union is the subset of $S \cup V(\mathcal{Q}^{\text{SE}})$ consisting of the points whose y -coordinates are in the

range $[\tilde{y}, \infty)$. The point $\phi(\tilde{y})$ we are looking for is just the leftmost point in $\bigcup_{i=1}^t S(\mathbf{u}_i)$ that is contained in U^{SE} . Note that $\phi(\tilde{y})$ is not contained in any southeast quadrant Q with $y(Q) < \tilde{y}$. Thus, $\phi(\tilde{y})$ is the leftmost point in $\bigcup_{i=1}^t S(\mathbf{u}_i)$ that is contained in $\bigcup_{i=1}^t U(\mathbf{u}_i)$. To compute $\phi(\tilde{y})$, it suffices to know the leftmost point in $S(\mathbf{u}_i)$ that is contained in $U(\mathbf{u}_j)$ for all $i, j \in \{1, \dots, t\}$. If $i > j$, then any point in $S(\mathbf{u}_i)$ is not contained in $U(\mathbf{u}_j)$. If $i = j$, then the leftmost point in $S(\mathbf{u}_i)$ contained in $U(\mathbf{u}_j)$ is just $a(\mathbf{u}_i) = a(\mathbf{u}_j)$. If $i < j$, then a point in $S(\mathbf{u}_i)$ is contained in $U(\mathbf{u}_j)$ iff it is contained in the quadrant $L(\mathbf{u}_j)$. Note that the points in $S(\mathbf{u}_i)$ are exactly the set of the points in S that lie in the strip $P_i = \mathbb{R} \times [y^-(\mathbf{u}_i), y^+(\mathbf{u}_i)]$, i.e., $S(\mathbf{u}_i) = S \cap P_i$. Therefore, the leftmost point in $S(\mathbf{u}_i)$ that is contained in $L(\mathbf{u}_j)$ can be computed in $\tilde{O}(1)$ time by querying the data structure \mathcal{A} with the 3-sided rectangle $P_i \cap L(\mathbf{u}_j)$. To summarize, the leftmost point in $S(\mathbf{u}_i)$ that is contained in $U(\mathbf{u}_j)$ can be computed in $\tilde{O}(1)$ time for any $i, j \in \{1, \dots, t\}$. Since $t = O(\log n)$, computing the leftmost point in $S(\mathbf{u}_i)$ that is contained in $U(\mathbf{u}_j)$ for all $i, j \in \{1, \dots, t\}$ takes $\tilde{O}(1)$ time. Once we have these points, the leftmost one among them is just the leftmost point in $\bigcup_{i=1}^t S(\mathbf{u}_i)$ that is contained in $\bigcup_{i=1}^t U(\mathbf{u}_i)$, i.e., the point $\phi(\tilde{y})$. We conclude that computing $\phi(\tilde{y})$ can be done in $\tilde{O}(1)$ time (by properly storing S and \mathcal{Q}^{SE} in some basic data structure).