Parameterized and Approximation Algorithms for Coverings Points with Segments in the Plane

Katarzyna Kowalska ✉
Faculty of Mathematics, Informatics, and Mechanics, University of Warsaw, Poland

Michał Pilipczuk ✉
Institute of Informatics, University of Warsaw, Poland

Abstract

We study parameterized and approximation algorithms for a variant of Set Cover, where the universe of elements to be covered consists of points in the plane and the sets with which the points should be covered are segments. We call this problem Segment Set Cover. We also consider a relaxation of the problem called δ*-extension*, where we need to cover the points by segments that are extended by a tiny fraction, but we compare the solution's quality to the optimum without extension.

For the unparameterized variant, we prove that Segment Set Cover does not admit a PTAS unless P=NP, even if we restrict segments to be axis-parallel and allow $\frac{1}{2}$-extension. On the other hand, we show that parameterization helps for the tractability of Segment Set Cover: we give an FPT algorithm for unweighted Segment Set Cover parameterized by the solution size k, a parameterized approximation scheme for Weighted Segment Set Cover with k being the parameter, and an FPT algorithm for Weighted Segment Set Cover with δ-extension parameterized by k and δ. In the last two results, relaxing the problem is probably necessary: we prove that Weighted Segment Set Cover without any relaxation is W[1]-hard and, assuming ETH, there does not exist an algorithm running in time $f(k) \cdot n^{o(k/\log k)}$. This holds even if one restricts attention to axis-parallel segments.

2012 ACM Subject Classification Theory of computation → Parameterized complexity and exact algorithms; Theory of computation → Approximation algorithms analysis

Keywords and phrases Geometric Set Cover, fixed-parameter tractability, weighted parameterized problems, parameterized approximation scheme

Digital Object Identifier 10.4230/LIPIcs.STACS.2024.47

Related Version *Full Version*: `https://arxiv.org/abs/2402.16466` [9]

Funding *Michał Pilipczuk*: This work is a part of project BOBR that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 948057).

Acknowledgements The authors would like to thank Krzysztof Maziarz for help with proofreading the manuscript.

1 Introduction

In the classic Set Cover problem, we are given a set of elements (universe) $\mathcal{U}$ and a family of sets $\mathcal{F}$ that are subsets of $\mathcal{U}$ and sum up to the whole $\mathcal{U}$. The task is to find a subfamily $\mathcal{S} \subseteq \mathcal{F}$ such that $\bigcup \mathcal{S} = \mathcal{U}$ and the size of $\mathcal{S}$ is minimum possible.

In the most general form, Set Cover is NP-complete, inapproximable within factor $(1-\delta)\ln|\mathcal{U}|$ for any $\delta > 0$ assuming P $\neq$ NP [5], and W[2]-complete with the natural parameterization by the size of the solution [4, Theorem 13.21]. However, restricting the problem to various specialized settings can lead to more tractable special cases. Particularly well-studied setting is that of Geometric Set Cover, where $\mathcal{U}$ consists of points in some Euclidean space V (most often the plane $\mathbb{R}^2$), while $\mathcal{F}$ consists of various geometric objects in V. In this paper we take a closer look at the Segment Set Cover problem, where we assume that $\mathcal{U}$ is a finite set of points in the plane and $\mathcal{F}$ consists of segments in the plane (not necessarily axis-parallel). Each of these problems has also a natural weighted variant, where each set $A \in \mathcal{F}$ comes with a nonnegative real weight $\mathbf{w}(A)$ and the task is to find a solution with the minimum possible total weight.

Approximation. Over the years there has been a lot of work related to approximation algorithms for Geometric Set Cover. Notably, Geometric Set Cover with unweighted unit disks or weighted unit squares admits a PTAS [6,17]. When we consider the same problem with weighted disks or squares (not necessarily unit), the problem admits a QPTAS [16], see also [19]. On the other hand, Chan and Grant proved that unweighted Geometric Set Cover with axis-parallel fat rectangles is APX-hard [3]. They also showed similar hardness for Geometric Set Cover with many other standard geometric objects. See the introductory section of [3] for a wider discussion of approximation algorithms for Geometric Set Cover with various kinds of geometric objects.

Parameterization. We consider Geometric Set Cover parameterized by the size of solution: Given an instance $(\mathcal{U}, \mathcal{F})$ and a parameter k, the task is to decide whether there is a solution of cardinality at most k. In the weighted setting, we look for a minimum-weight solution among those of cardinality at most k, and k remains a parameter.

(Unweighted) Geometric Set Cover where $\mathcal{F}$ consists of lines in the plane is called Point Line Cover, and it is a textbook example of a problem that admits a quadratic kernel and a $2^{\mathcal{O}(k \log k)} \cdot n^{\mathcal{O}(1)}$-time fixed-parameter algorithm (cf. [4, Exercise 2.4]). See also the work of Kratsch et al. [10] for a matching lower bound on the kernel size and a discussion of the relevant literature. The simple branching and kernelization ideas behind the parameterized algorithms for Point Line Cover were generalized by Langerman and Morin [11] to an abstract variant of Geometric Set Cover where the sets of $\mathcal{F}$ can be assigned a suitable notion of dimension. This framework in particular applies to the problem of covering points with hyperspaces in $\mathbb{R}^d$.

As proved by Marx, unweighted Geometric Set Cover with unit squares in the plane is already W[1]-hard [12, Theorem 5]. Later, Marx and Pilipczuk showed that there is an algorithm running in time $n^{\mathcal{O}(\sqrt{k})}$ that solves weighted Geometric Set Cover with squares or with disks, and that this running time is tight under the Exponential-Time Hypothesis (ETH) [15]. However, they also showed that any small deviations from the setting of squares or disks – for instance considering thin rectangles or rectangles with sidelengths in the interval $[1, 1 + \delta]$ for any $\delta > 0$ – lead to problems for which there are lower bounds refuting running times of the form $f(k) \cdot n^{o(k)}$ or $f(k) \cdot n^{o(k/\log k)}$, for any computable f. See [15] for a broader exposition of these results and for more literature pointers.

We are not aware of any previous work that concretely considered the Segment Set Cover problem. In particular, it seems that the framework of Langerman and Morin [11] does *not* apply to this problem, since no suitable notion of dimension can be assigned to segments in the plane (more concretely, the fundamental [11, Lemma 1] fails, which renders further arguments not applicable). In [13] Marx considered the related Dominating Set problem in intersection graphs of axis-parallel segments, and proved it to be W[1]-hard. The parameterized complexity of the Independent Set problem for segments in the plane was studied in the same work of Marx, and independently by Kára and Kratochvíl [8].

δ-extension. We also consider the δ-extension relaxation of the Segment Set Cover problem. Formally, for a center-symmetric object $L \subseteq \mathbb{R}^2$ with center of symmetry $S = (x_s, y_s)$, the *δ-extension* of L is the set:

$$L^{+\delta} = \{(1+\epsilon) \cdot (x - x_s, y - y_s) + (x_s, y_s) : (x, y) \in L, 0 \leqslant \epsilon < \delta\}.$$

That is, $L^{+\delta}$ is the image of L under homothety centred at S with scale $(1+\delta)$ but with the extreme points excluded. In particular, δ-extension turns a closed segment into a segment without endpoints and a rectangle into the interior of a rectangle; this is a technical detail that will turn out to be useful in presentation.

In Geometric Set Cover with δ-extension, we assume that in the given instance $(\mathcal{U}, \mathcal{F})$, $\mathcal{F}$ consists of center-symmetric objects, and we are additionally given the accuracy parameter $\delta > 0$. The task is to find $\mathcal{S} \subseteq \mathcal{F}$ such that $\mathcal{S}^{+\delta} := \{L^{+\delta} : L \in \mathcal{S}\}$ covers all points in $\mathcal{U}$, but the quality of the solution – be it the cardinality or the weight of $\mathcal{S}$ – is compared to the optimum without assuming extension. Thus, requirements on the the output solution are relaxed: the points of $\mathcal{U}$ have to be covered only after expanding every object of $\mathcal{S}$ a tiny bit. The parameterized variants of Geometric Set Cover with δ-extension are defined naturally: the task is to either find a solution of size at most k that covers all of $\mathcal{U}$ after δ-extension, or conclude that there is no solution of size k that covers $\mathcal{U}$ without extension.

The study of the δ-extension relaxation is motivated by the *δ-shrinking* relaxation considered in the context of the Geometric Independent Set problem: given a family $\mathcal{F}$ of objects in the plane, find the maximum size subfamily of pairwise disjoint objects. In the δ-shrinking model, the output solution is required to be disjoint only after shrinking every object by a $1 - \delta$ multiplicative factor. Geometric Independent Set remains W[1]-hard for as simple objects as unit disks or unit squares [13] and admits a QPTAS for polygons [2], but the existence of a PTAS for the problem is widely open. However, as first observed by Adamaszek et al. [1], and then confirmed by subsequent works of Wiese [20] and of Pilipczuk et al. [18], adopting the δ-shrinking relaxation leads to a robust set of FPT algorithms and (efficient or parameterized) approximation schemes. The motivation of this work is to explore if the analogous δ-extension relaxation of Geometric Set Cover also leads to more positive results.

In fact, we are not the first to consider the δ-extension relaxation of Geometric Set Cover. In [7], Har-Peled and Lee considered the Weighted Geometric Set Cover problem with δ-extension[1] for fat polygons, and proved that the problem admits a PTAS with running time $|\mathcal{F}|^{\mathcal{O}(\epsilon^{-2}\delta^{-2})} \cdot |\mathcal{U}|$. Given this result, our goal is to understand the complexity in the setting of ultimately non-fat polygons: segments.

Our contribution. First, we show that Segment Set Cover does not have a polynomial-time approximation scheme (PTAS) assuming P≠NP, even if segments are axis-parallel and we relax the problem with $\frac{1}{2}$-extension. Thus, there is no hope for the analog of the result of Har-Peled and Lee [7] in the setting of segments.

▶ **Theorem 1.** *There exists a constant $\gamma > 0$ such that, unless P=NP, there is no polynomial-time algorithm that given an instance $(\mathcal{U}, \mathcal{F})$ of (unweighted) Segment Set Cover in which all segments are axis-parallel, returns a set $\mathcal{S} \subseteq \mathcal{F}$ such that $\mathcal{S}^{+\frac{1}{2}}$ covers $\mathcal{U}$ and $|\mathcal{S}| \leqslant (1+\gamma) \cdot \mathsf{opt}$, where opt denotes the minimum size of a subset of $\mathcal{F}$ that covers $\mathcal{U}$.*

[1] We note that Har-Peled and Lee considered a different definition of δ-extension, where every object L is extended by all points at distance at most $\delta \cdot \mathsf{rad}(L)$, where $\mathsf{rad}(L)$ is the radius of the largest circle inscribed in L. This definition works well for fat polygons, but not so for segments, hence we adopt the homothetical definition of δ-extension discussed above.

Theorem 1 justifies also considering parameterization by the solution size k. For this parameterization, we provide three parameterized algorithms:

- an FPT algorithm for (unweighted) SEGMENT SET COVER with k being the parameter;
- a *parameterized approximation scheme* (PAS) for WEIGHTED SEGMENT SET COVER: a $(1+\epsilon)$-approximation algorithm with running time of the form $f(k,\epsilon)\cdot(|\mathcal{U}||\mathcal{F}|)^{\mathcal{O}(1)}$; and
- an FPT algorithm for WEIGHTED SEGMENT SET COVER with δ-extension, where both k and $\delta > 0$ are the parameters.

Formal statements of these results follow below.

▶ **Theorem 2.** *There is an algorithm that given a family $\mathcal{F}$ of segments in the plane, a set $\mathcal{U}$ of points in the plane, and a parameter k, runs in time $k^{\mathcal{O}(k)}\cdot(|\mathcal{U}||\mathcal{F}|)^{\mathcal{O}(1)}$, and either outputs a set $\mathcal{S}\subseteq\mathcal{F}$ such that $|\mathcal{S}|\leqslant k$ and $\mathcal{S}$ covers all points in $\mathcal{U}$, or determines that such a set $\mathcal{S}$ does not exist.*

▶ **Theorem 3.** *There is an algorithm that given a family $\mathcal{F}$ of weighted segments in the plane, a set $\mathcal{U}$ of points in the plane, and parameters k and $\epsilon > 0$, runs in time $(k/\epsilon)^{\mathcal{O}(k)}\cdot(|\mathcal{U}||\mathcal{F}|)^{\mathcal{O}(1)}$ and outputs a set $\mathcal{S}$ such that:*

- *$\mathcal{S}\subseteq\mathcal{F}$, $|\mathcal{S}|\leqslant k$, and $\mathcal{S}$ covers all points in $\mathcal{U}$, and*
- *the weight of $\mathcal{S}$ is not greater than $1+\epsilon$ times the minimum weight of a subset of $\mathcal{F}$ of size at most k that covers $\mathcal{U}$,*

or determines that there is no set $\mathcal{S}\subseteq\mathcal{F}$ with $|\mathcal{S}|\leqslant k$ such that $\mathcal{S}$ covers all points in $\mathcal{U}$.

▶ **Theorem 4.** *There is an algorithm that given a family $\mathcal{F}$ of weighted segments in the plane, a set $\mathcal{U}$ of points in the plane, and parameters k and $\delta > 0$, runs in time $f(k,\delta)\cdot(|\mathcal{U}||\mathcal{F}|)^{\mathcal{O}(1)}$ for some computable function f and outputs a set $\mathcal{S}$ such that:*

- *$\mathcal{S}\subseteq\mathcal{F}$, $|\mathcal{S}|\leqslant k$, $\mathcal{S}^{+\delta}$ covers all points in $\mathcal{U}$, and*
- *the weight of $\mathcal{S}$ is not greater than the minimum weight of a subset of $\mathcal{F}$ that covers $\mathcal{U}$ without δ-extension,*

or determines that there is no set $\mathcal{S}\subseteq\mathcal{F}$ with $|\mathcal{S}|\leqslant k$ such that $\mathcal{S}$ covers all points in $\mathcal{U}$.

It is natural to ask whether relying on relaxations – $(1+\epsilon)$-approximation or δ-extension – is really necessary for WEIGHTED SEGMENT SET COVER, as Theorem 2 shows that it is not in the unweighted setting. Somewhat surprisingly, we show that this is the case by proving the following result. Recall that here we consider WEIGHTED SEGMENT SET COVER as a parameterized problem where we seek a solution of the minimum total weight among those of cardinality at most k.

▶ **Theorem 5.** *The WEIGHTED SEGMENT SET COVER problem is W[1]-hard when parameterized by k and assuming ETH, there is no algorithm for this problem with running time $f(k)\cdot(|\mathcal{U}|+|\mathcal{F}|)^{o(k/\log k)}$ for any computable function f. Moreover, this holds even if all segments in $\mathcal{F}$ are axis-parallel.*

Thus, the uncovered parameterized complexity of SEGMENT SET COVER is an interesting one: the problem is FPT when parameterized by the solution size k in the unweighted setting, but this tractability ceases to hold when moving to the weighted setting. However, fixed-parameter tractability in the weighted setting can be restored if one considers any of the following relaxations: $(1+\epsilon)$-approximation or δ-extension.

Organization. In Section 2 we prove Theorems 2, 3 and 4, while in Section 3 we prove Theorem 5. Due to space constraints, the proof of Theorem 1 is presented only in the full version of this article [9].

2 Algorithms

In this section we give our positive results – Theorems 2, 3, and 4. We start with a shared definition. For a set of collinear points C in the plane, *extreme points* of C are the endpoints of the smallest segment that covers all points from set C. In particular, if C consists of one point or is empty, then there are 1 or 0 extreme points, respectively.

2.1 Unweighted segments and a parameterized approximation scheme

We first a give an FPT algorithm for Weighted Segment Set Cover where we additionally consider the number of different weights to be the parameter.

▶ **Theorem 6.** *There is an algorithm that given a family $\mathcal{F}$ of weighted segments in the plane, a set $\mathcal{U}$ of points in the plane, and a parameter k, runs in time $(qk)^{\mathcal{O}(k)} \cdot (|\mathcal{U}||\mathcal{F}|)^{\mathcal{O}(1)}$, where q is the number of different weights used by the weight function, and either outputs a solution $\mathcal{S} \subseteq \mathcal{F}$ such that $|\mathcal{S}| \leqslant k$ and $\mathcal{S}$ covers all points in $\mathcal{U}$, or determines that such a set $\mathcal{S}$ does not exist.*

Clearly, Theorem 2 follows from applying Theorem 6 for $q = 1$. However, later we use Theorem 6 for larger values of q to obtain our parameterized approximation scheme: Theorem 3.

We remark that the proof of Theorem 6 relies on branching and kernelization arguments that are standard in parameterized algorithms. Even though the statement does not formally follow from the work of Langerman and Morin [11], the basic technique is very similar.

Towards the proof of Theorem 6, we may assume that the given instance $(\mathcal{U}, \mathcal{F}, \mathbf{w})$, where $\mathbf{w}\colon \mathcal{F} \to \mathbb{R}_{\geqslant 0}$ denotes the weight function on $\mathcal{F}$, is *reasonable* in the following sense: there do not exist distinct $A, B \in \mathcal{F}$ with the same weight such that $A \cap \mathcal{U} \subseteq B \cap \mathcal{U}$. Indeed, then A could be safely removed from $\mathcal{F}$, since in any solution, taking B instead of A does not increase the weight and may only result in covering more points in $\mathcal{U}$. In the next lemma we show that in reasonable instances we can find a small subset of $\mathcal{F}$ that is guaranteed to intersect every small solution.

▶ **Lemma 7.** *Suppose $(\mathcal{U}, \mathcal{F}, \mathbf{w})$ is a reasonable instance of Weighted Segment Set Cover where the weight function $\mathbf{w}$ uses at most q different values. Suppose further that there exists a line L in the plane with at least $k+1$ points of $\mathcal{U}$ on it. Then there exists a subset $\mathcal{R} \subseteq \mathcal{F}$ of size at most qk such that every subset $\mathcal{S} \subseteq \mathcal{F}$ with $|\mathcal{S}| \leqslant k$ that covers $\mathcal{U}$ satisfies $|\mathcal{R} \cap \mathcal{S}| \geqslant 1$. Moreover, such a subset $\mathcal{R}$ can be found in polynomial time.*

Proof. Let us enumerate the points of $\mathcal{U}$ that lie on L as $x_1, x_2, \ldots, x_t$ in the order in which they appear on L. By reasonability of $(\mathcal{U}, \mathcal{F})$, for every $i \in \{1, \ldots, k\}$ there exist at most q different segments in $\mathcal{F}$ that are collinear with L and cover x_i, but do not cover x_{i-1} (or just cover x_1, in case $i = 1$). Indeed, if $A \in \mathcal{F}$ is collinear with L, covers x_i and does not cover x_{i-1}, then $A \cap \mathcal{U} = \{x_i, \ldots, x_j\}$ for some $j \geqslant i$; so if there was another $B \in \mathcal{F}$ with the same property and the same weight as A, then the reasonability of $(\mathcal{U}, \mathcal{F})$ would imply that $A = B$. Let $\mathcal{R}_i$ be the set of segments with the property discussed above; then $|\mathcal{R}_i| \leqslant q$. Our proposed set is defined as:

$$\mathcal{R} \coloneqq \bigcup_{i=1}^{k} \mathcal{R}_i.$$

Clearly, $\mathcal{R}$ can be found in polynomial time and $|\mathcal{R}| \leqslant qk$. It remains to prove that $\mathcal{R}$ has the desired property. Consider any set $\mathcal{S} \subseteq \mathcal{F}$ of size at most k that covers $\mathcal{U}$.

Let $\mathcal{S}_L$ be the set of segments from $\mathcal{S}$ that are collinear with L. Every segment that is not collinear with L can cover at most one of the points that lie on this line. Hence, if $\mathcal{S}_L$ was empty, then $\mathcal{S}$ would cover at most k points on line L, but L had at least $k+1$ different points from $\mathcal{U}$ on it.

Therefore, we know that $\mathcal{S}_L$ is not empty and hence $|\mathcal{S} - \mathcal{S}_L| \leqslant k-1$. Segments from $\mathcal{S} - \mathcal{S}_L$ can cover at most $k-1$ points among $\{x_1, x_2, \ldots, x_k\}$, therefore at least one of these points must be covered by segments from $\mathcal{S}_L$. Let $i \in \{1, \ldots, k\}$ be the smallest index such that x_i is covered by a segment in $\mathcal{S}_L$. Then, by minimality, this segment cannot cover x_{i-1} (if existent), so it must belong to $\mathcal{R}_i$. We conclude that $\mathcal{R} \cap \mathcal{S}$ is nonempty, as desired. ◀

With Lemma 7 in hand, we prove Theorem 6 using a straightforward branching strategy.

Proof of Theorem 6. Let $(\mathcal{U}, \mathcal{F}, \mathbf{w})$ be the given instance and k be the given parameter: the target size of a solution. We present a recursive algorithm that proceeds as follows:

(1) As long as there are distinct sets $A, B \in \mathcal{F}$ with $A \cap \mathcal{U} \subseteq B \cap \mathcal{U}$ and $\mathbf{w}(A) = \mathbf{w}(B)$, remove A from $\mathcal{F}$. Once this step is applied exhaustively, the instance $(\mathcal{U}, \mathcal{F}, \mathbf{w})$ is reasonable.

(2) If there is a line with at least $k+1$ points from $\mathcal{U}$, we branch over the choice of adding to the solution one of the at most qk possible segments from the set $\mathcal{R}$ provided by Lemma 7. That is, for every $s \in \mathcal{R}$, we recurse on the instance $(\mathcal{U} - s, \mathcal{F} - \{s\}, \mathbf{w})$, and parameter $k-1$. If any such recursive call returned a solution $\mathcal{S}'$, then return the lightest among solutions $\mathcal{S}' \cup \{s\}$ obtained in this way. Otherwise, return that there is no solution.

(3) If every line has at most k points on it and $|\mathcal{U}| > k^2$, then return that there is no solution.

(4) If $|\mathcal{U}| \leqslant k^2$, solve the problem by brute force: check all subsets of $\mathcal{F}$ of size at most k.

That the algorithm is correct is clear: the correctness of step (2) follows from Lemma 7, and to see the correctness of step (3) note that if no line contains more than k points, than no segment of $\mathcal{F}$ can cover more than k points in $\mathcal{U}$, hence having more than k^2 points in $\mathcal{U}$ implies that there is no solution of size at most k.

For the time complexity, observe that in the leaves of the recursion we have $|\mathcal{U}| \leqslant k^2$, so $|\mathcal{F}| \leqslant qk^4$, because every segment can be uniquely identified by its weight and the two extreme points it covers (this follows from reasonability). Therefore, there are $\binom{qk^4}{\leqslant k} \leqslant (qk)^{\mathcal{O}(k)}$ possible solutions to check, each can be checked in polynomial time. Thus, step (4) takes time $(qk)^{\mathcal{O}(k)}$ whenever applied in the leaf of the recursion.

During the recursion, the parameter k is decreased with every recursive call, so the recursion tree has depth at most k and at each node we branch over at most qk possibilities. Thus, there are at most $\mathcal{O}((qk)^k)$ nodes in the recursion tree, and local computation in each of them can be done in time $(|\mathcal{U}||\mathcal{F}|)^{\mathcal{O}(1)} \cdot (qk)^{\mathcal{O}(k)}$ (the second factor is due to possibly applying step (4) in the leaves). Thus, the time complexity of the algorithm is $(qk)^{\mathcal{O}(k)} \cdot (|\mathcal{U}||\mathcal{F}|)^{\mathcal{O}(1)}$. ◀

Finally, we use Theorem 6 to prove Theorem 3, recalled below for convenience. The idea is to multiplicatively round the weights so that we obtain an instance with only few different weight values, on which the algorithm of Theorem 6 can be employed.

► **Theorem 3.** *There is an algorithm that given a family $\mathcal{F}$ of weighted segments in the plane, a set $\mathcal{U}$ of points in the plane, and parameters k and $\epsilon > 0$, runs in time $(k/\epsilon)^{\mathcal{O}(k)} \cdot (|\mathcal{U}||\mathcal{F}|)^{\mathcal{O}(1)}$ and outputs a set $\mathcal{S}$ such that:*

- *$\mathcal{S} \subseteq \mathcal{F}$, $|\mathcal{S}| \leqslant k$, and $\mathcal{S}$ covers all points in $\mathcal{U}$, and*
- *the weight of $\mathcal{S}$ is not greater than $1+\epsilon$ times the minimum weight of a subset of $\mathcal{F}$ of size at most k that covers $\mathcal{U}$,*

or determines that there is no set $\mathcal{S} \subseteq \mathcal{F}$ with $|\mathcal{S}| \leqslant k$ such that $\mathcal{S}$ covers all points in $\mathcal{U}$.

Proof. Let $\mathcal{S}^\star$ be an optimum solution: a minimum-weight set at most k segments in $\mathcal{F}$ that covers $\mathcal{U}$. The algorithm does not know $\mathcal{S}^\star$, but by branching into at most $|\mathcal{F}|$ choices we may assume that it knows the weight of the heaviest segment in $\mathcal{S}^\star$; call this weight W. Thus, we have $W \leqslant \mathbf{w}(\mathcal{S}^\star) \leqslant kW$. We may dispose of all segments in $\mathcal{F}$ whose weight is larger than W, as they will for sure not participate in the solution.

We define a new weight function $\mathbf{w}' \colon \mathcal{F} \to \mathbb{R}_{\geqslant 0}$ as follows. Consider any segment $A \in \mathcal{F}$. If $\mathbf{w}(A) \leqslant \frac{\epsilon}{2k} \cdot W$, then set $\mathbf{w}'(A) \coloneqq \frac{\epsilon}{2k} \cdot W$. Otherwise, set $\mathbf{w}'(A) \coloneqq \frac{W}{(1+\epsilon/2)^i}$, where i is the unique integer such that

$$\frac{W}{(1+\epsilon/2)^{i+1}} < \mathbf{w}(A) \leqslant \frac{W}{(1+\epsilon/2)^i}.$$

Note that the assumption $\mathbf{w}(A) > \frac{\epsilon}{2k} \cdot W$ implies that we always have $i \leqslant \log_{1+\epsilon/2}(2k/\epsilon) = \mathcal{O}(1/\epsilon \log(k/\epsilon))$. As we also have $i \geqslant 0$ due to removing segments of weight larger than W, we conclude that the weight function $\mathbf{w}'$ uses at most $\mathcal{O}(1/\epsilon \log(k/\epsilon))$ different weight values.

Next, observe that for every segment $A \in \mathcal{F}$, we have

$$\mathbf{w}'(A) \leqslant (1+\epsilon/2) \cdot \mathbf{w}(A) + \frac{\epsilon}{2k} \cdot W.$$

Summing this inequality through all segments of $\mathcal{S}^\star$ yields

$$\mathbf{w}'(\mathcal{S}^\star) \leqslant (1+\epsilon/2) \cdot \mathbf{w}(\mathcal{S}^\star) + k \cdot \frac{\epsilon}{2k} \cdot W \leqslant (1+\epsilon/2) \cdot \mathbf{w}(\mathcal{S}^\star) + \epsilon/2 \cdot \mathbf{w}(\mathcal{S}^\star) = (1+\epsilon) \cdot \mathbf{w}(\mathcal{S}^\star).$$

As $\mathcal{S}^\star$ is an optimum solution, we conclude that the optimum solution in the instance $(\mathcal{U}, \mathcal{F}, \mathbf{w}')$ for parameter k is at most $(1+\epsilon)$ times heavier than the optimum solution in the instance $(\mathcal{U}, \mathcal{F}, \mathbf{w})$ for parameter k. Hence, it suffices to apply the algorithm of Theorem 6 to the instance $(\mathcal{U}, \mathcal{F}, \mathbf{w}')$ and parameter k and return the obtained solution. The running time is $(1/\epsilon \cdot k \log(k/\epsilon))^{\mathcal{O}(k)} \cdot (|\mathcal{U}||\mathcal{F}|)^{\mathcal{O}(1)} = (k/\epsilon)^{\mathcal{O}(k)} \cdot (|\mathcal{U}||\mathcal{F}|)^{\mathcal{O}(1)}$, as promised. ◀

2.2 Weighted segments with δ-extension

In this section we prove Theorem 4, restated below for convenience.

▶ **Theorem 4.** *There is an algorithm that given a family $\mathcal{F}$ of weighted segments in the plane, a set $\mathcal{U}$ of points in the plane, and parameters k and $\delta > 0$, runs in time $f(k,\delta) \cdot (|\mathcal{U}||\mathcal{F}|)^{\mathcal{O}(1)}$ for some computable function f and outputs a set $\mathcal{S}$ such that:*
- *$\mathcal{S} \subseteq \mathcal{F}$, $|\mathcal{S}| \leqslant k$, $\mathcal{S}^{+\delta}$ covers all points in $\mathcal{U}$, and*
- *the weight of $\mathcal{S}$ is not greater than the minimum weight of a subset of $\mathcal{F}$ that covers $\mathcal{U}$ without δ-extension,*

or determines that there is no set $\mathcal{S} \subseteq \mathcal{F}$ with $|\mathcal{S}| \leqslant k$ such that $\mathcal{S}$ covers all points in $\mathcal{U}$.

Roughly speaking, our approach to prove Theorem 4 is to find a small kernel for the problem; but we need to be careful with the definition of kernelization, because we work in the δ-extension model. The key technical tool will be the notion of a *dense subset*.

Dense subsets. Intuitively speaking, for a set of collinear points C, a subset $A \subseteq C$ is dense if any small cover of A becomes a cover of C after a tiny extension. This is formalized in the following definition.

▶ **Definition 8.** *For a set of collinear points C, a subset $A \subseteq C$ is* (k,δ)-dense *in C if for any set of segments $\mathcal{R}$ that covers A and such that $|\mathcal{R}| \leqslant k$, it holds that $\mathcal{R}^{+\delta}$ covers C.*

The key combinatorial observation in our approach is expressed in the following Lemma 9: in every collinear set C one can always find a (k,δ)-dense subset of size bounded by a function of k and δ. Later, this lemma will allow us to find a kernel for our original problem.

▶ **Lemma 9.** *For every set C of collinear points in the plane, $\delta > 0$ and $k \geqslant 1$, there exists a (k,δ)-dense set $A \subseteq C$ of size at most $(2+\frac{4}{\delta})^k$. Moreover, such a (k,δ)-dense set can be computed in time $\mathcal{O}(|C| \cdot (2+\frac{4}{\delta})^k)$.*

Proof. We give a proof of the existence of such a dense subset A, and at the end we will argue that the proof naturally gives rise to an algorithm with the promised complexity. We fix δ and proceed by induction on k. Formally, we shall prove the following stronger statement: For any set of collinear points C, there exists a subset $A \subseteq C$ such that:

- A is (k,δ)-dense in C,
- $|A| \leqslant (2+\frac{4}{\delta})^k$, and
- the extreme points of C are in A.

Consider first the base case when $k = 1$. Then it is sufficient to just take A that consists of the (at most 2) extreme points of C. Indeed, if the extreme points of C are covered with one segment, then this segment must cover the whole set C (even without extension). Thus, the set A has size at most $2 < (2+\frac{4}{\delta})^1$, as required.

We now proceed to the inductive step. Assuming inductive hypothesis for any set of collinear points C and for parameter k, we will prove it for $k+1$.

Let s be the minimal segment that includes all points from C. That is, s is the segment whose endpoints are the extreme points of C.

Split s into $M := \lceil 1 + \frac{4}{\delta} \rceil$ subsegments of equal length. We name these segments $v_1, v_2, \dots, v_M$ in order, and we consider them closed. Note that $|v_i| = \frac{|s|}{M}$ for each $1 \leqslant i \leqslant M$, where $|\cdot|$ denotes the length of a segment.

Let C_i be the subset of C consisting of points belonging to v_i. Further, let t_i be the segment with endpoints being the extreme points of C_i. Note that t_i might be a degenerate single-point segment if C_i consists of one point, or even t_i might be empty if C_i is empty. Figure 1 presents an example of the construction.

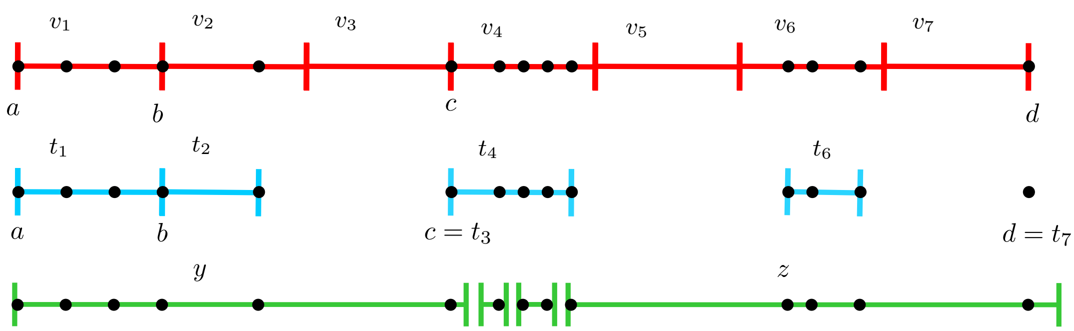

Figure 1 Example of the construction in the proof of Lemma 9 for $M = 7$ and some set of points C (marked with black circles). The top panel shows segments v_i. The middle panel shows segments t_i. Note that t_5 is an empty segment, because there are no points in C that belong to v_5, while each of the segments t_3 and t_7 is degenerated to a single point: c and d, respectively. Segments t_1 and t_2 share one point b. The bottom panel shows an example of the second case in the correctness proof: a solution $\mathcal{R}$ of size 4 whose all segments intersect t_4. Then one of y and z will cover the whole of C_4 after extension.

We use the inductive hypothesis to choose a (k,δ)-dense subset A_i of C_i, for each $i \in \{1,\ldots,M\}$. Note that if $|C_i| \leqslant 1$, then $A_i = C_i$, so A_i is (k,δ)-dense set for C_i. Also, by assumption, A_i contains the extreme points of C_i.

Next, we define $A \coloneqq \bigcup_{i=1}^{M} A_i$. Thus A includes the extreme points of C, because they are included in the sets A_1 and A_M.

By induction, the size of each A_i is at most $(2+\frac{4}{\delta})^k$. Therefore,

$$|A| \leqslant M\left(2+\frac{4}{\delta}\right)^k = \left\lceil 1+\frac{4}{\delta}\right\rceil \cdot \left(2+\frac{4}{\delta}\right)^k \leqslant \left(2+\frac{4}{\delta}\right)^{k+1}.$$

We are left with verifying that A is $(k+1,\delta)$-dense in C. For this, consider any cover of A with $k+1$ segments and call it $\mathcal{R}$.

Consider any segment t_i. If there exists a segment $x \in \mathcal{R}$ that is disjoint with t_i, then $\mathcal{R} - \{x\}$ constitutes a cover of A_i with at most k segments. Since A_i is (k,δ)-dense in C_i, $(\mathcal{R} - \{x\})^{+\delta}$ covers C_i. So $\mathcal{R}^{+\delta}$ covers C_i as well.

On the other hand, if for any fixed t_i a segment $x \in \mathcal{R}$ as above does not exist, then all the $k+1$ segments of $\mathcal{R}$ intersect t_i. An example of such a situation is depicted in the bottom panel of Figure 1. Let us consider any such t_i. By induction, the endpoints of s are in A_1 and A_M respectively, so $\mathcal{R}$ must cover them. So for each endpoint of s, there exists a segment in $\mathcal{R}$ that contains this endpoint and intersects t_i. Let us call these two segments y and z. It follows that $|y|+|z|+|t_i| \geqslant |s|$. Since $|t_i| \leqslant |v_i| = \frac{|s|}{M} \leqslant \frac{|s|}{1+\frac{4}{\delta}} = \frac{\delta|s|}{\delta+4}$, we have

$$\max(|y|,|z|) \geqslant |s|\left(1-\frac{\delta}{\delta+4}\right)/2 = \frac{2|s|}{\delta+4}.$$

After δ-extension, the longer of the segments y and z will expand at both ends by at least:

$$\delta/2 \cdot \max(|y|,|z|) \geqslant \frac{\delta|s|}{\delta+4} = \frac{|s|}{1+\frac{4}{\delta}} \geqslant \frac{|s|}{M} = |v_i| \geqslant |t_i|.$$

Therefore, the longer of segments y and z will cover the whole segment t_i after δ-extension. We conclude that $\mathcal{R}^{+\delta}$ covers C_i as well.

Since $C = \bigcup_{i=1}^{M} C_i$, we conclude that $\mathcal{R}^{+\delta}$ covers C. So indeed, A is $(k+1,\delta)$-dense in C. This concludes the proof of the existence of such a dense set A. To compute A in time $\mathcal{O}\left(|C|\cdot(2+\frac{4}{\delta})^k\right)$ observe that the inductive proof explained above can be easily turned into a recursive procedure that for a given C and k, outputs a (k,δ)-dense subset of C. The recursion tree of this procedure has size $\mathcal{O}\left((2+\frac{4}{\delta})^k\right)$ in total, while every recursive calls uses $\mathcal{O}(|C|)$ time for internal computation, so the total running time is $\mathcal{O}\left(|C|\cdot(2+\frac{4}{\delta})^k\right)$. ◀

Long lines. We need a few additional observations in the spirit of the algorithm of Theorem 6. For a finite set of points $\mathcal{U}$ in the plane, call a line L *k-long* with respect to $\mathcal{U}$ if L contains more than k points from $\mathcal{U}$. We have the following observations.

▶ **Lemma 10.** *Let $\mathcal{U}$ be a finite set of points in the plane such that there are more than k lines that are k-long with respect to $\mathcal{U}$. Then $\mathcal{U}$ cannot be covered with k segments.*

Proof. We proceed by contradiction. Assume there are at least $k+1$ different k-long lines and there is a set of segments $\mathcal{R}$ of size at most k covering all points in $\mathcal{U}$.

Consider any k-long line L. Note that every segment $\mathcal{R}$ which is not collinear with L, covers at most one point that lies on L. Since L is long, there are at least $k+1$ points from $\mathcal{U}$ that lie on L. This implies that there must be a segment in $\mathcal{R}$ that is collinear with L.

Since we have at least $k+1$ different long lines, there are at least $k+1$ segments in $\mathcal{R}$ collinear with different lines. This contradicts the assumption that $|\mathcal{R}| \leqslant k$. ◀

▶ **Lemma 11.** *Let $\mathcal{U}$ be a finite set of points in the plane such that there are more than k^2 points from $\mathcal{U}$ that do not lie on any line that is k-long with respect to $\mathcal{U}$. Then $\mathcal{U}$ cannot be covered with k segments.*

Proof. We proceed by contradiction. Assume that we have more than k^2 points in $\mathcal{U}$ that do not lie on any k-long line. Call this set A. Suppose there is a set of segments $\mathcal{R}$ of size at most k that covers all points in $\mathcal{U}$.

Since any line in the plane can cover only at most k points in A, the same is also true for every segment in $\mathcal{R}$. Therefore, the segments from $\mathcal{R}$ can cover at most k^2 points in A in total. As $|A| > k^2$, $\mathcal{R}$ cannot cover the whole A, which is a subset of $\mathcal{U}$; a contradiction. ◀

We are now ready to give a proof of Theorem 4.

Proof of Theorem 4. Let $(\mathcal{U}, \mathcal{F}, \mathbf{w})$ be the input instance of WEIGHTED SEGMENT SET COVER, where $\mathbf{w}\colon \mathcal{F} \to \mathbb{R}_{\geqslant 0}$ is the weight function. Further, let k and $\delta > 0$ be the input parameters. Our goal is to either conclude that $(\mathcal{U}, \mathcal{F}, \mathbf{w})$ has no solution of cardinality at most k, or to find an instance $(\mathcal{U}', \mathcal{F}', \mathbf{w})$ of size bounded by $f(k, \delta)$ for some computable function f and satisfying $\mathcal{U}' \subseteq \mathcal{U}$ and $\mathcal{F}' \subseteq \mathcal{F}$, such that the following two properties hold:

- *(Property 1)* For every set $\mathcal{S} \subseteq \mathcal{F}$ such that $|\mathcal{S}| \leqslant k$ and $\mathcal{S}$ covers $\mathcal{U}$, there is a set $\mathcal{S}_1 \subseteq \mathcal{F}'$ such that $|\mathcal{S}_1| \leqslant k$, the weight of $\mathcal{S}_1$ is not greater than the weight of $\mathcal{S}$, and $\mathcal{S}_1$ covers $\mathcal{U}'$.
- *(Property 2)* For every set $\mathcal{S} \subseteq \mathcal{F}'$ such that $|\mathcal{S}| \leqslant k$ and $\mathcal{S}$ covers all points in $\mathcal{U}'$, $\mathcal{S}^{+\delta}$ covers all points in the original set $\mathcal{U}$.

Suppose we constructed such an instance $(\mathcal{U}', \mathcal{F}', \mathbf{w})$. Then using *Property 1* we know that an optimum solution of size at most k to $(\mathcal{U}', \mathcal{F}', \mathbf{w})$ has no greater weight than an optimum solution of size at most k to $(\mathcal{U}, \mathcal{F}, \mathbf{w})$. On the other hand, using *Property 2* we know that any solution to $(\mathcal{U}', \mathcal{F}', \mathbf{w})$ after δ-extension covers $\mathcal{U}$. So it will remain to find an optimum solution to the instance $(\mathcal{U}', \mathcal{F}', \mathbf{w})$. This can be done by brute-force in time $|\mathcal{F}'|^{k+\mathcal{O}(1)} \cdot |\mathcal{U}'|^{\mathcal{O}(1)}$, which is bounded by a computable function of k and δ.

It remains to construct the instance $(\mathcal{U}', \mathcal{F}', \mathbf{w})$. Let ℓ be the number of different lines that are k-long with respect to $\mathcal{U}$. By Lemmas 10 and 11, if we had more than k different k-long lines or more than k^2 points from $\mathcal{U}$ that do not lie on any k-long line, then we can safely conclude that $(\mathcal{U}, \mathcal{F}, \mathbf{w})$ has no solution of cardinality at most k, and terminate the algorithm. So assume otherwise, in particular $\ell \leqslant k$.

Next, we cover $\mathcal{U}$ with at most $k + 1$ sets:

- D consists of all points in $\mathcal{U}$ that do not lie on any k-long line. Then we have $|D| \leqslant k^2$.
- For $1 \leqslant i \leqslant \ell$, C_i consists of all points in $\mathcal{U}$ that lie on the i-th long line. Then $|C_i| > k$.

Note that sets C_i do not need to be disjoint.

For every set C_i, we apply Lemma 9 to obtain a subset $A_i \subseteq C_i$ that is (k, δ)-dense and satisfies $|A_i| \leqslant (2 + \frac{4}{\delta})^k$. We define $\mathcal{U}' := D \cup \bigcup_{i=1}^{\ell} A_i$. Thus, $\mathcal{U}'$ has size at most $k^2 + k(2 + \frac{4}{\delta})^k$. Further, we define $\mathcal{F}'$ as follows: for every pair of points in $\mathcal{U}'$, if there are segments in $\mathcal{F}$ that cover this pair of points, we choose one such segment with the lowest weight and include it in $\mathcal{F}'$. Thus $\mathcal{F}'$ has size at most $|\mathcal{U}'|^2$, which means that both $\mathcal{F}'$ and $\mathcal{U}'$ have sizes bounded by $\mathcal{O}\left((k^2 + k(2 + \frac{4}{\delta})^k)^2\right)$. We are left with verifying Properties 1 and 2.

For Property 2, consider any set $\mathcal{S} \subseteq \mathcal{F}'$ such that $|\mathcal{S}| \leqslant k$ and $\mathcal{S}$ covers all points in $\mathcal{U}'$. Then in particular, for every $i \in \{1, \ldots, \ell\}$, $\mathcal{S}$ in covers all points in A_i. As A_i is (k, δ)-dense in C_i, we conclude that $\mathcal{S}^{+\delta}$ covers C_i. Hence $\mathcal{S}^{+\delta}$ covers $D \cup \bigcup_{i=1}^{\ell} C_i = \mathcal{U}$, as required.

For Property 1, consider any solution $\mathcal{S}$ to $(\mathcal{U}, \mathcal{F}, \mathbf{w})$ of size at most k. For every segment $s \in \mathcal{S}$, let B_s be the set of points in $\mathcal{U}'$ that are covered by s. B_s is of course a set of collinear points, hence B_s can be covered by any segment that covers the extreme points of B_s. Therefore, we can replace s with a segment $s' \in \mathcal{F}$ that has the lowest weight among the segments that cover the extreme points of B_s. Such a segment belongs to $\mathcal{F}'$ by construction, and s' has weight no greater than the weight of s, because s also covers the extreme points of B_s. Therefore, if $\mathcal{S}_1 \subseteq \mathcal{F}'$ is the set obtained by performing such replacement for every $s \in \mathcal{S}$, then $\mathcal{S}_1$ has both size and weight not greater than $\mathcal{S}$, and $\mathcal{S}_1$ covers $\mathcal{U}'$. ◀

3 W[1]-hardness of Weighted Segment Set Cover

In this section we prove Theorem 5, recalled below for convenience.

▶ **Theorem 5.** *The Weighted Segment Set Cover problem is W[1]-hard when parameterized by k and assuming ETH, there is no algorithm for this problem with running time $f(k) \cdot (|\mathcal{U}| + |\mathcal{F}|)^{o(k/\log k)}$ for any computable function f. Moreover, this holds even if all segments in $\mathcal{F}$ are axis-parallel.*

To prove Theorem 5, we give a reduction from a W[1]-hard problem: Partitioned Subgraph Isomorphism, defined as follows. An instance of Partitioned Subgraph Isomorphism consists of a *pattern graph* H, a *host graph* G, and a function $\lambda\colon V(G) \to V(H)$ that colors every vertex of G with a vertex of H. The task is to decide whether there exists a subgraph embedding $\phi\colon V(H) \to V(G)$ that respects the coloring λ. That is, the following conditions have to be satisfied.

- $\lambda(\phi(a)) = a$ for each $a \in V(H)$, and
- $\phi(a)$ and $\phi(b)$ are adjacent in G for every edge $ab \in E(H)$.

The following complexity lower bound for Partitioned Subgraph Isomorphism was proved by Marx in [14].

▶ **Theorem 12** ([14]). *Consider the Partitioned Subgraph Isomorphism problem where the pattern graph H is assumed to be 3-regular. Then this problem is W[1]-hard when parameterized by k, the number of vertices of H, and assuming the ETH there is no algorithm solving this problem in time $f(k) \cdot |V(G)|^{o(k/\log k)}$, where f is any computable function.*

In the remainder of this section we prove Theorem 5 by providing a parameterized reduction from Partitioned Subgraph Isomorphism to Weighted Segment Set Cover. The technical statement of the reduction is encapsulated in the following lemma.

▶ **Lemma 13.** *Given an instance (H, G, λ) of Partitioned Subgraph Isomorphism where H is 3-regular and has k vertices, one can in polynomial time construct an instance $(\mathcal{U}, \mathcal{F}, \mathbf{w})$ of Weighted Segment Set Cover and a positive real W such that:*

(1) *all segments in $\mathcal{F}$ are axis-parallel;*
(2) *if the instance (H, G, λ) has a solution, then there exists a solution to $(\mathcal{U}, \mathcal{F}, \mathbf{w})$ of cardinality $\frac{11}{2}k$ and weight at most W; and*
(3) *if there exists a solution to $(\mathcal{U}, \mathcal{F}, \mathbf{w})$ of weight at most W, then the instance (H, G, λ) has a solution.*

Note that in (3) we in fact do not require any bound on the cardinality of the solution, just on its weight.

It is easy to see that Lemma 13 implies Theorem 12. First, Lemma 13 gives a parameterized reduction from the W[1]-hard Partitioned Subgraph Isomorphism problem with 3-regular pattern graphs to Weighted Segment Set Cover parameterized by the cardinality of the sought solution, which shows that the latter problem is also W[1]-hard. Second, combining the reduction with an algorithm for Weighted Segment Set Cover with running time as postulated in Theorem 5 would give an algorithm for Partitioned Subgraph Isomorphism with running time $f(k) \cdot |V(G)|^{o(k/\log k)}$ for a computable function f, which would contradict ETH by Theorem 12. So we are left with giving a proof of Lemma 13, which spans the remainder of this section.

The key element of the proof will be a construction of a *choice gadget* that works within a single line; this construction is presented in the lemma below. Here, a *chain* is a sequence $(A_1, A_2, \ldots, A_\ell)$ of subsets of $\mathbb{R}$ such that for each $i \in \{1, \ldots, \ell-1\}$, all numbers in A_i are strictly smaller than all numbers in A_{i+1}.

▶ **Lemma 14.** *Suppose we are given an integer $N > 100$ and p chains $\{(A_{j,1}, \ldots, A_{j,\ell}) : j \in \{1, \ldots, p\}\}$ of length ℓ each such that the sets $\{A_{j,t} : j \in \{1, \ldots, p\}, t \in \{1, \ldots, \ell\}\}$ are all pairwise disjoint and contained in $\{1, \ldots, N\}$. Then one can in polynomial time construct a set of points $\mathcal{U} \subseteq \mathbb{R}$, $\mathcal{U} \supseteq \{1, \ldots, N\}$, as well as a set of segments $\mathcal{F}$ contained in $\mathbb{R}$ such that the following holds:*

- *For every $j \in \{1, \ldots, p\}$ and every set B that contains exactly one point from each element of the chain $(A_{j,1}, \ldots, A_{j,\ell})$, there exists $\mathcal{R}_B \subseteq \mathcal{F}$ such that $|\mathcal{R}_B| = \ell + 1$, $\mathcal{R}_B$ covers all points of $\mathcal{U}$ except for B, and the total length of the segments in $\mathcal{R}_B$ is equal to $N + 1 - 2\ell/N^2$.*
- *For every subset of segments $\mathcal{R} \subseteq \mathcal{F}$, if $\mathcal{R}$ covers all points in $\mathcal{U} - \{1, \ldots, N\}$, then the total length of segments in $\mathcal{R}$ is at least $N + 1 - 2/N$.*
- *For every subset of segments $\mathcal{R} \subseteq \mathcal{F}$, if the total length of segments of $\mathcal{R}$ is not larger than $N + \frac{3}{2}$ and $\mathcal{R}$ covers all points in $\mathcal{U} - \{1, \ldots, N\}$, then the total length of segments of $\mathcal{R}$ is equal to $N + 1 - 2\ell/N^2$ and there exists $j \in \{1, \ldots, p\}$ such that for every $t \in \{1, \ldots, \ell\}$, $\mathcal{R}$ does not cover the whole set $A_{j,t}$.*

Proof. Denote $I \coloneqq \{1, \ldots, N\}$ and $\epsilon \coloneqq 1/N^2$ for convenience. For every $i \in I$, let

$$i^- \coloneqq i - \epsilon \qquad \text{and} \qquad i^+ \coloneqq i + \epsilon.$$

Define $I^- \coloneqq \{i^- : i \in I\}$, $I^+ \coloneqq \{i^+ : i \in I\}$, and

$$\mathcal{U} \coloneqq \{0\} \cup I^- \cup I \cup I^+.$$

Next, for every $j \in \{1, \ldots, p\}$, define the following set of segments:

$$\mathcal{R}_j \coloneqq \{[0, a^-] : a \in A_{j,1}\} \cup \bigcup_{t=1}^{\ell-1} \{[a^+, b^-] : (a, b) \in A_{j,t} \times A_{j,t+1}\} \cup \{[a^+, N+1] : a \in A_{j,\ell}\}.$$

We set

$$\mathcal{F} \coloneqq \bigcup_{j=1}^{p} \mathcal{R}_j.$$

See Figure 2 for a visualization of the construction. We are left with verifying the three postulated properties of $\mathcal{U}$ and $\mathcal{F}$.

Figure 2 Construction of Lemma 14 for $N = 8$. Elements of $I \cup \{0\}$ are depicted with circles and elements of $I^+ \cup I^-$ are depicted with squares. Blue segments represent the set $\mathcal{R}_B$ for $B = \{3, 7\}$.

For the first property, let b_t be the unique element of $B \cap A_{j,t}$, for $t \in \{1,\ldots,\ell\}$, and let

$$\mathcal{R}_B \coloneqq \{[0, b_1^-], [b_1^+, b_2^-], \ldots, [b_{\ell-1}^+, b_\ell^-], [b_\ell^+, N+1]\}.$$

It is straightforward to see that $\mathcal{R}_B$ covers all the points of $\mathcal{U}$ except for B, and that the total sum of lengths of segments in $\mathcal{R}_B$ is $N+1-2\ell\epsilon = N+1-2\ell/N^2$.

For the second postulated property, observe that each segment of $\mathcal{F}$ that covers any point $i^+ \in I^+$, in fact covers the whole interval $[i^+, (i+1)^-]$ (where $(N+1)^- = N+1$). Similarly, each segment of $\mathcal{F}$ that covers any point $i^- \in I^-$, in fact covers the whole interval $[(i-1)^+, i^-]$ (where $0^+ = 0$). Hence, if $\mathcal{R} \subseteq \mathcal{F}$ covers all points of $\mathcal{U} - I$, in particular $\mathcal{R}$ covers all points in $I^+ \cup I^-$, hence also all intervals of the form $[i^+, (i+1)^-]$ for $i \in \{0,1,\ldots,N\}$. The sum of the lengths of those intevals is equal to $N+1-2\epsilon N = N+1-2/N$. Hence, the sum of length of intervals in $\mathcal{R}$ must be at least $N+1-2/N$.

For the third postulated property, observe that if two segments of $\mathcal{F}$ intersect, then their intersection is a segment of length at least $1-2\epsilon$. Since $\mathcal{R}$ covers all points of $\mathcal{U} - I$, by the second property the sum of lengths of the segments in $\mathcal{R}$ is at least $N+1-2/N$. Now if any of those segments intersected, then the total sum of lengths of the segments in $\mathcal{R}$ would be at least $N+1-2/N+(1-2\epsilon)$, which is larger than $N+\frac{3}{2}$. We conclude that the segments of $\mathcal{R}$ are pairwise disjoint.

Since $0 \in \mathcal{U} - I$, there is a segment $s_1 \in \mathcal{R}$ that covers 0. By construction, there exists $j \in \{1,\ldots,p\}$ such that $s_1 = [0, b_{j,1}^-]$ for some $b_{j,1} \in A_{j,1}$. As the segments of $\mathcal{R}$ are pairwise disjoint and cover all points in I^+, the next (in the natural order on $\mathbb{R}$) segment in $\mathcal{R}$ must start at $b_{j,1}^+$, and in particular $b_{j,1}$ is not covered by $\mathcal{R}$. Since all sets in all chains on input are pairwise disjoint, the segment in $\mathcal{R}$ starting at $b_{j,1}^+$ must be of the form $s_2 = [b_{j,1}^+, b_{j,2}^-]$ for some $b_{j,2} \in A_{j,2}$. Continuing this reasoning, we find that in fact $\mathcal{R} = \mathcal{R}_B$ for some set $B = \{b_{j,1}, b_{j,2}, \ldots, b_{j,\ell}\}$ such that $b_{j,t} \in A_{j,t}$ for each $t \in \{1,\ldots,\ell\}$. In particular, the total length of segments in $\mathcal{R}$ is equal to $N+1-2\epsilon\ell$ and $\mathcal{R}$ does not cover any point in B; the latter implies that for each $t \in \{1,\ldots,\ell\}$, $\mathcal{R}$ does not cover $A_{j,t}$ entirely. ◀

With Lemma 14 established, we proceed to the proof of Lemma 13.

Let (H, G, λ) be the given instance of Partitioned Subgraph Isomorphism where H is a 3-regular graph. Let $k \coloneqq |V(H)|$ and $\ell \coloneqq |E(H)|$; note that $\ell = \frac{3}{2}k$. We may assume that $V(H) = \{1,\ldots,k\}$, and that whenever uv is an edge in G, we have that $\lambda(u)\lambda(v)$ is an edge of H (other edges in G play no role in the problem and can be discarded). We construct an instance $(\mathcal{U}, \mathcal{F}, \mathbf{w})$ of Weighted Segment Set Cover as follows; see Figure 3 for a visualization.

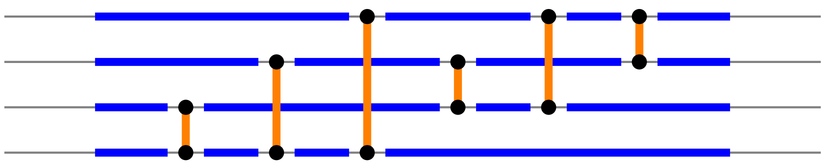

Figure 3 Example solution in the instance $(\mathcal{U}, \mathcal{F})$ constructed in the proof of Lemma 13 for $H = K_4$. Blue segments belong to the sets $\mathcal{S}_i$ for $i \in \{1,2,3,4\}$ and orange segments belong to $\mathcal{D}$.

For each edge $ab \in E(H)$, let E_{ab} be the subset of those edges uv of G for which $\lambda(u) = a$ and $\lambda(v) = b$. Thus, $\{E_{ab} \colon ab \in E(H)\}$ is a partition of $E(G)$. Let $N \coloneqq |E(G)|$ and $\xi \colon E(G) \to \{1,\ldots,N\}$ be any bijection such that for each $ab \in E(H)$, $\xi(E_{ab})$ is a contiguous interval of integers. By copying some vertices of G if necessary, we may assume that $N > 100k$.

Consider any $a \in \{1, \ldots, k\}$ and let b_1, b_2, b_3 be the three neighbors of a in H, ordered so that $(\xi(E_{ab_1}), \xi(E_{ab_2}), \xi(E_{ab_3}))$ is a chain. For each $u \in \lambda^{-1}(a)$, let E_u be the set of edges of G incident to u, and let us construct the chain

$$C_u := (\xi(E_u \cap E_{ab_1}), \xi(E_u \cap E_{ab_2}), \xi(E_u \cap E_{ab_3})).$$

Note that all sets featured in all the chains C_u, for $u \in \lambda^{-1}(a)$, are pairwise disjoint. We now apply Lemma 14 for the integer N and the chains $\{C_u : u \in \lambda^{-1}(a)\}$. This way, we construct a suitable point set $\mathcal{U}_a \subseteq \mathbb{R}$ and a set of segments $\mathcal{F}_a$ contained in $\mathbb{R}$. We put all those points and segments on the line $\{(x, a) : x \in \mathbb{R}\}$; that is, every point $x \in \mathcal{U}_a$ is replaced with the point (x, a), and similarly for the segments of $\mathcal{F}_a$. By somehow abusing the notation, we let $\mathcal{U}_a$ and $\mathcal{F}_a$ be the point set and the segment set after the replacement.

Next, for every edge uv of G, we define s_{uv} to be the segment with endpoints $(\xi(uv), a)$ and $(\xi(uv), b)$, where $a = \lambda(u)$ and $b = \lambda(v)$.

We set

$$\mathcal{U} := \bigcup_{a=1}^{k} \mathcal{U}_a \qquad \text{and} \qquad \mathcal{F} := \{s_{uv} : uv \in E(G)\} \cup \bigcup_{a=1}^{k} \mathcal{F}_a.$$

Note that all segments in sets $\mathcal{F}_a$ are horizontal and each segment s_{uv} is vertical, thus $\mathcal{F}$ consists of axis-parallel segments. Each segment $s \in \bigcup_{a=1}^{k} \mathcal{F}_a$ is assigned weight $\mathbf{w}(s)$ equal to the length of s, and each segment s_{uv} for $uv \in E(G)$ is assigned weight $\mathbf{w}(s_{uv}) = \delta$, where $\delta := 1/N^4$. Finally, we set

$$W := k \cdot (N + 1 - 6/N^2) + \delta\ell.$$

This concludes the construction of the instance $(\mathcal{U}, \mathcal{F}, \mathbf{w})$. We are left with verifying the correctness of the reduction, which is done in the following two claims.

▷ Claim 15. Suppose the input instance (H, G, λ) of PARTITIONED SUBGRAPH ISOMORPHISM has a solution. Then the output instance $(\mathcal{U}, \mathcal{F}, \mathbf{w})$ of WEIGHTED SEGMENT SET COVER has a solution of cardinality $4k + \ell = \frac{11}{2}k$ and weight at most W.

Proof. Let ϕ be the supposed solution to (H, G, λ). By the first property of Lemma 14, for every $a \in \{1, \ldots, k\}$ there is a set $\mathcal{R}_{\phi,a}$ of size 4 and total weight $N + 1 - 6/N^2$ that covers all points from $\mathcal{U}_a$ except for the points

$$(\xi(\phi(a)\phi(b_1)), a), (\xi(\phi(a)\phi(b_2)), a), (\xi(\phi(a)\phi(b_3)), a),$$

where b_1, b_2, b_3 are the neighbors of a in H. Define

$$\mathcal{S} := \{s_{\phi(a)\phi(b)} : ab \in E(H)\} \cup \bigcup_{a=1}^{k} \mathcal{R}_{\phi,a}.$$

Thus, for each $a \in \{1, \ldots, k\}$, the aforementioned points of $\mathcal{U}_a$ not covered by $\mathcal{R}_{\phi,a}$ are actually covered by the segments $s_{\phi(a)\phi(b_1)}, s_{\phi(a)\phi(b_2)}, s_{\phi(a)\phi(b_3)}$. We conclude that $\mathcal{S}$ covers all the points in $\mathcal{U}$ and has cardinality $4k + \ell = \frac{11}{2}k$ and total weight W, as promised. ◁

▷ Claim 16. Suppose the output instance $(\mathcal{U}, \mathcal{F}, \mathbf{w})$ of WEIGHTED SEGMENT SET COVER has a solution of weight at most W. Then the input instance (H, G, λ) of PARTITIONED SUBGRAPH ISOMORPHISM has a solution.

Proof. Let $\mathcal{S}$ be the supposed solution to $(\mathcal{U}, \mathcal{F}, \mathbf{w})$. Denote

$$\mathcal{D} \coloneqq \mathcal{S} \cap \{s_{uv} \colon uv \in E(G)\}$$

and

$$\mathcal{S}_a \coloneqq \mathcal{S} \cap \mathcal{F}_a \qquad \text{for } a \in \{1, \dots, k\}.$$

Fix $a \in \{1, \dots, k\}$ for a moment. Observe that the segments from $\mathcal{D}$ can only cover points with positive integer coordinates within the set $\mathcal{U}_a$, hence the whole point set $\mathcal{U}_a - (\{1, \dots, N\} \times \{a\})$ has to be covered by $\mathcal{S}_a$. By the second property of Lemma 14 we infer that the total weight of $\mathcal{S}_a$ must be at least $N + 1 - 2/N$.

Observe now that

$$W - k \cdot (N + 1 - 2/N) = \delta\ell + 2k/N - 6k/N^2 < \frac{1}{2}.$$

It follows that the total weight of each set $\mathcal{S}_a$ must be smaller than $N + \frac{3}{2}$, for otherwise the sum of weights of sets $\mathcal{S}_a$ would be larger than W. By the third property of Lemma 14, we infer that for every $a \in \{1, \dots, k\}$, the total weight of $\mathcal{S}_a$ is equal to $N + 1 - 6/N^2$ and there exists $\phi(a) \in \lambda^{-1}(a)$ such that $\mathcal{S}_a$ does not entirely cover any of the sets $\xi(E_{\phi(a)} \cap E_{ab_1}), \xi(E_{\phi(a)} \cap E_{ab_2}), \xi(E_{\phi(a)} \cap E_{ab_3})$, where b_1, b_2, b_3 are the three neighbors of a in H. In particular, there are edges $e_{a,b_1} \in E_{ab_1}, e_{a,b_2} \in E_{ab_2}, e_{a,b_3} \in E_{ab_3}$, all sharing the endpoint $\phi(a)$, such that $\mathcal{S}_a$ does not cover the points $(\xi(e_{a,b_1}), a), (\xi(e_{a,b_2}), a), (\xi(e_{a,b_3}), a)$. Call these points X_a and let $X \coloneqq \bigcup_{a=1}^{k} X_a$. Note that

$$|X| = 3k = 2\ell$$

and that X must be entirely covered by $\mathcal{D}$.

Since the weight of $\mathcal{S}_a$ is equal to $N + 1 - 6/N^2$ for each $a \in \{1, \dots, k\}$, the weight of $\mathcal{D}$ is upper bounded by

$$W - k \cdot (N + 1 - 6/N^2) = \delta\ell.$$

As every member of $\mathcal{D}$ has weight δ, we conclude that $|\mathcal{D}| \leqslant \ell$. Now, one can readily verify that every segment $s_{uv} \in \mathcal{D}$ can cover at most two points in X, as X cannot contain more than two points with the same horizontal coordinate (recall that this coordinate is the index of an edge of G). Moreover, s_{uv} can cover two points in X only if $u = \phi(a)$ and $v = \phi(b)$, where $a = \lambda^{-1}(u)$ and $b = \lambda^{-1}(v)$. As $|X| = 2\ell$ and $|\mathcal{D}| \leqslant \ell$, this must be the case for every segment in $\mathcal{D}$. In particular, $\phi(a)\phi(b)$ must be an edge in G for every edge $ab \in E(H)$, so ϕ is a solution to the instance (H, G, λ) of Partitioned Subgraph Isomorphism. $\lhd$

Claims 15 and 16 finish the proof of Lemma 13. So the proof of Theorem 5 is also done.

References

1 Anna Adamaszek, Parinya Chalermsook, and Andreas Wiese. How to tame rectangles: Solving Independent Set and Coloring of rectangles via shrinking. In *18th International Conference on Approximation Algorithms for Combinatorial Optimization Problems, APPROX 2015*, volume 40 of *LIPIcs*, pages 43–60. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPIcs.APPROX-RANDOM.2015.43.

2 Anna Adamaszek, Sariel Har-Peled, and Andreas Wiese. Approximation schemes for independent set and sparse subsets of polygons. *J. ACM*, 66(4):29:1–29:40, 2019. doi:10.1145/3326122.

3 Timothy M. Chan and Elyot Grant. Exact algorithms and APX-hardness results for geometric packing and covering problems. *Comput. Geom.*, 47(2):112–124, 2014. doi:10.1016/j.comgeo.2012.04.001.

4 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized Algorithms.* Springer, 2015. doi:10.1007/978-3-319-21275-3.

5 Irit Dinur and David Steurer. Analytical approach to parallel repetition. In *46th ACM Symposium on Theory of Computing, STOC 2014*, pages 624–633. ACM, 2014. doi:10.1145/2591796.2591884.

6 Thomas Erlebach and Erik Jan van Leeuwen. PTAS for Weighted Set Cover on unit squares. In *13th International Workshop on Approximation and Combinatorial Optimization, APPROX 2010*, volume 6302 of *Lecture Notes in Computer Science*, pages 166–177. Springer, 2010. doi:10.1007/978-3-642-15369-3_13.

7 Sariel Har-Peled and Mira Lee. Weighted geometric set cover problems revisited. *J. Comput. Geom.*, 3(1):65–85, 2012. doi:10.20382/jocg.v3i1a4.

8 Jan Kára and Jan Kratochvíl. Fixed parameter tractability of independent set in segment intersection graphs. In *Second International Workshop on Parameterized and Exact Computation, IWPEC 2006*, volume 4169 of *Lecture Notes in Computer Science*, pages 166–174. Springer, 2006. doi:10.1007/11847250_15.

9 Katarzyna Kowalska and Michał Pilipczuk. Parameterized and approximation algorithms for coverings points with segments in the plane. *CoRR*, abs/2402.16466, 2024. arXiv:2402.16466.

10 Stefan Kratsch, Geevarghese Philip, and Saurabh Ray. Point Line Cover: The easy kernel is essentially tight. *ACM Trans. Algorithms*, 12(3):40:1–40:16, 2016. doi:10.1145/2832912.

11 Stefan Langerman and Pat Morin. Covering things with things. *Discret. Comput. Geom.*, 33(4):717–729, 2005. doi:10.1007/s00454-004-1108-4.

12 Dániel Marx. Efficient approximation schemes for geometric problems? In *13th European Symposium on Algorithms, ESA 2005*, pages 448–459. Springer, 2005.

13 Dániel Marx. Parameterized complexity of independence and domination on geometric graphs. In *Second International Workshop on Parameterized and Exact Computation, IWPEC 2006*, volume 4169 of *Lecture Notes in Computer Science*, pages 154–165. Springer, 2006. doi:10.1007/11847250_14.

14 Dániel Marx. Can you beat treewidth? *Theory Comput.*, 6(1):85–112, 2010. doi:10.4086/toc.2010.v006a005.

15 Dániel Marx and Michał Pilipczuk. Optimal parameterized algorithms for planar facility location problems using voronoi diagrams. *ACM Trans. Algorithms*, 18(2):13:1–13:64, 2022. doi:10.1145/3483425.

16 Nabil H. Mustafa, Rajiv Raman, and Saurabh Ray. Settling the apx-hardness status for geometric set cover. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014*, pages 541–550. IEEE Computer Society, 2014. doi:10.1109/FOCS.2014.64.

17 Nabil H. Mustafa and Saurabh Ray. Improved results on geometric hitting set problems. *Discret. Comput. Geom.*, 44(4):883–895, 2010. doi:10.1007/s00454-010-9285-9.

18 Michał Pilipczuk, Erik Jan van Leeuwen, and Andreas Wiese. Approximation and parameterized algorithms for geometric independent set with shrinking. In *42nd International Symposium on Mathematical Foundations of Computer Science, MFCS 2017*, volume 83 of *LIPIcs*, pages 42:1–42:13. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.MFCS.2017.42.

19 Michał Pilipczuk, Erik Jan van Leeuwen, and Andreas Wiese. Quasi-polynomial time approximation schemes for packing and covering problems in planar graphs. *Algorithmica*, 82(6):1703–1739, 2020. doi:10.1007/s00453-019-00670-w.

20 Andreas Wiese. Independent Set of convex polygons: From n^{ϵ} to $1+\epsilon$ via shrinking. *Algorithmica*, 80(3):918–934, 2018. doi:10.1007/s00453-017-0347-8.