

Hitting Axis-Parallel Segments with Weighted Points

Rajiv Raman 

Indraprastha Institute of Information Technology Delhi, New Delhi, India

Siddhartha Sarkar 

Indian Institute of Science, Bengaluru, India

Jatin Yadav 

Indian Institute of Technology Delhi, New Delhi, India

Abstract

We study a geometric hitting-set problem in which the input consists of a set P of weighted points and a family $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ of axis-parallel segments in the plane. The goal is to select a minimum-weight subset of P that hits every segment in $\mathcal{S}$. Even restricted geometric hitting-set problems are known to be computationally hard, and for axis-parallel segments the standard decomposition into horizontal and vertical sub-instances yields only a simple factor-2 approximation.

We present an LP-rounding algorithm that breaks the factor-2 barrier. For the weighted problem, we obtain a randomized $(1 + 2/e)$ -approximation by combining systematic rounding on horizontal lines with an exact repair step on residual vertical sub-instances. In the unweighted case, a sharper analysis gives a $(1 + 1/(e - 1))$ -approximation. Finally, we consider the case where one of the sub-instances consists of lines instead of line segments, a problem considered by Fekete et al. (Geometric Hitting Set for Segments of Few Orientations, Theor. Comp. Sys., 62 (2) 2018),. In this case, we improve their result to obtain an approximation factor of $1 + 1/e$ and show that the problem is APX-hard. We also present algorithms for the generalization to d orientations, as well as PTASes for bounded-complexity subclasses of the unweighted Hitting Set problem.

2012 ACM Subject Classification Theory of computation $\rightarrow$ Design and analysis of algorithms; Theory of computation $\rightarrow$ Computational geometry; Theory of computation $\rightarrow$ Approximation algorithms analysis

Keywords and phrases Geometric Hitting Set, Approximation Algorithms, Computational Geometry, LP Rounding, Axis-Parallel Segments, PTAS, APX-hardness

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2026.14

Category APPROX

Related Version *Extended Version*: <https://arxiv.org/abs/2605.14499> [33]

1 Introduction

A *set system* $(P, \mathcal{S})$ consists of a set P of elements and a collection $\mathcal{S}$ of subsets of P . Given a set system, $(P, \mathcal{S})$, a set $P' \subseteq P$ is a *hitting set* if each set in $\mathcal{S}$ contains an element of P' . In the *unweighted* Hitting Set Problem, the goal is to find a hitting set of minimum cardinality, and in the *weighted* Hitting Set Problem, we are additionally given a weight function $w : P \rightarrow \mathbb{R}$ and the goal is to select a minimum weight hitting set P' , i.e., a hitting set such that $\sum_{p \in P'} w(p)$ is minimized. The set systems considered in this paper are finite. We assume that $|P| = n$ and $|\mathcal{S}| = m$.

Without restrictions, the Hitting Set Problem is equivalent to the Set Cover Problem on the *dual set system*, where points in P and sets in $\mathcal{S}$ swap roles. For these general systems, $O(\log m)$ -approximation algorithms exist with nearly matching lower bounds. However, we can often bypass these limits by exploiting the structure of restricted instances. This paper



© Rajiv Raman, Siddhartha Sarkar, and Jatin Yadav;
licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2026).

Editors: Mohit Singh and Tom Gur; Article No. 14; pp. 14:1–14:22



Leibniz International Proceedings in Informatics

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

focuses on *geometric set systems*, where elements and sets are defined by objects in the plane. While Hitting Set and Set Cover remain NP-hard in these settings, they frequently admit better approximation ratios than general systems. Despite this, a unified set of techniques and a complete characterization of these systems remain elusive.

There are two broad frameworks that yield improved approximation algorithms for Hitting Set and Set Cover. The first is via ϵ -nets and LP-rounding ([20, 8, 16, 36, 10]) that yield algorithms with approximation factors depending on the *combinatorial complexity* (such as *VC-dimension*, or *shallow-cell complexity*) of the instance. However, these techniques at best yield $O(1)$ -approximation algorithms with large hidden constants even for simple objects such as disks or squares.

A second framework is *local search*, that yields PTASes when the geometric set system is *non-piercing*, but works only in the unweighted setting [28, 13, 32]. Informally, a set $\mathcal{S}$ of objects in the plane is non-piercing if the boundary of each object in $\mathcal{S}$ is a simple Jordan curve, and for any $A, B \in \mathcal{S}$, $A \setminus B$ and $B \setminus A$ are both connected sets (see [32] for a precise definition). On the other hand, Chan and Grant [10] (see also Har-Peled and Quanrud [19]) showed even with nearly congruent *almost-squares*, the Set Cover and Hitting Set problems become APX-hard. Recently, Banik et al. [6] showed that the local search technique yields a PTAS for the Set Cover problem even when the geometric shapes are not non-piercing, i.e., they *pierce*, but in a bounded manner. However, no such result is known for the Hitting Set Problem.

Our goal is to develop improved approximation algorithms for structured, weighted geometric set systems. Specifically, we address the Hitting Set problem for axis-parallel line segments (formally defined below).

Weighted Geometric Hitting Set for Axis-Parallel Segments

Input: A finite set P of points in the plane with a non-negative weight function $w: P \rightarrow \mathbb{R}_{\geq 0}$, and a family $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ of axis-parallel segments (where $\mathcal{H}$ are horizontal and $\mathcal{V}$ are vertical).

Goal: Find a subset $Q \subseteq P$ that minimizes the total weight $\sum_{p \in Q} w(p)$ such that Q hits every segment in $\mathcal{S}$ (i.e., $Q \cap s \neq \emptyset$ for all $s \in \mathcal{S}$).

Axis-parallel segments generalize several tractable special cases. When all segments in $\mathcal{S}$ are parallel to a single axis, the problem becomes one-dimensional; alternatively, if segments only intersect those of orthogonal orientations, the problem maps to bipartite graphs. In both scenarios, the Hitting Set, Set Cover, and also the Independent Set problems are solvable in polynomial time because their natural LP-relaxations are integral, owing to totally unimodular constraint matrices [34]. Crucially, this total unimodularity breaks down in our setting, where segments possess both orientations and parallel segments overlap.

For Hitting Set, Set Cover, and Independent Set, a trivial 2-approximation follows from independently solving the one-dimensional sub-instances for each axis and combining the results. Remarkably, while this remains the best-known approximation factor for Set Cover and Independent Set, our primary contribution is breaking this factor-2 barrier for the weighted Hitting Set problem.

1.1 Our Contributions

Our main contribution is a suite of improved approximation algorithms for the Hitting Set problem on axis-parallel line segments and points in the plane, summarized in Table 1. For line segments in two orientations, we achieve a $(1 + 2/e) \approx 1.74$ -approximation in the

weighted setting and a $(1 + 1/(e - 1)) \approx 1.58$ -approximation in the unweighted setting. When one orientation consists of infinite lines instead of segments, we improve this to a $(1 + 1/e) \approx 1.36$ -approximation. We further generalize our results to segments parallel to one of d directions. These results are obtained via a randomized rounding procedure on the natural linear program, but can be derandomized using the standard method of conditional expectations (See [33]). We also present a PTAS for specific restricted instances.

■ **Table 1** The table above summarizes our results. Note that the previously best-known 2-approximation is considered a folklore result.

Problem	Approx. Factor	Complexity	Previous bounds
Wt. axis-parallel	$1 + 2/e$ (Th. 1)		2
Unwt. axis-parallel	$1 + 1/(e - 1)$ (Th. 2)		2, APX-hard [17]
d -dir wt.	$\ln d + \ln \ln d + O(1)$ (Th. 3)		d
Hor-seg, vert lines	$1 + 1/e$ (Thm. 5)	APX-hard (Thm. 6)	$5/3$ [17]
k -restricted unwt.	PTAS (Thm. 8)		NP-hard
Lines in 3-directions		APX-hard (Thm. 7)	$4/3$ [15], NP-hard [17]

► **Theorem 1 (Weighted).** *Let $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ be a set of horizontal and vertical line segments in the plane and let P be a set of points with a weight function $w : P \rightarrow \mathbb{R}_{\geq 0}$ such that each segment in $\mathcal{S}$ contains at least one point of P . Then, there is a randomized $(1 + 2/e)$ -approximation algorithm for the weighted Hitting Set problem defined by $\mathcal{S}$ and P .*

For the unweighted setting, we obtain a better approximation factor.

► **Theorem 2 (Unweighted).** *Let $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ be a set of horizontal and vertical line segments in the plane and let P be a set of points such that each segment in $\mathcal{S}$ contains at least one point of P . Then, there is a randomized $(1 + 1/(e - 1))$ -approximation algorithm for the Hitting Set problem defined by $\mathcal{S}$ and P .*

We can generalize the result for axis-parallel segments in the setting where each segment is parallel to one of d directions. Fekete et al. [17] studied the setting where each object to be hit is not a single segment, but is the union of h segments (each coming from one of d orientations). The authors provide an $(h \cdot d)$ -approximation for this setting. Applying the result below, it is possible to obtain an improved $h \cdot (\log d + \log \log d + O(1))$ -approximation.

► **Theorem 3 (d-Directions).** *Let $\mathcal{S}$ be a set of segments, each parallel to one of d directions in the plane for some $d \geq 2$, and let P be a set of points with a weight function $w : P \rightarrow \mathbb{R}_{\geq 0}$ such that each segment in $\mathcal{S}$ contains at least one point of P . Then, there is a randomized algorithm that is a $(\ln d + \ln \ln d + O(1))$ -approximation relative to the natural LP.*

We show that the natural LP-relaxation of the unweighted problem has an integrality gap of at least $5/4$.

► **Theorem 4 (Integrality gap).** *Let $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ be a set of horizontal and vertical line segments in the plane and let P be a set of points such that each segment in $\mathcal{S}$ contains at least one point of P . The natural LP-relaxation of the problem has an integrality gap of at least $5/4$.*

We also consider the setting when the segments in one of the orientations is replaced by lines, a variant considered by Fekete et al. [17], who gave a $(5/3)$ -approximation for the unweighted case. We obtain a $1 + 1/e \approx 1.36$ -approximation.

► **Theorem 5** (Lines and Segments). *Let $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ be a set of horizontal line segments and vertical lines in the plane and let P be a set of points with a weight function $w : P \rightarrow \mathbb{R}_{\geq 0}$ such that each object in $\mathcal{S}$ contains at least one point of P . Then, there is a randomized $(1 + 1/e)$ -approximation algorithm for the weighted Hitting Set problem defined by $\mathcal{S}$ and P .*

This setting with lines in one direction was shown to be NP-hard by Fekete et al. [17] if the point set consists of all points in the plane. However, if the set P is finite, we show that in fact, the problem is APX-hard.

► **Theorem 6** (APX-Hardness). *Let $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ be a set of horizontal line segments and vertical lines in the plane and P be a finite set of points in the plane such that each object in $\mathcal{S}$ contains at least one point of P . The Hitting Set problem with $\mathcal{S}$ and P is APX-hard.*

If instead of line segments, we have lines in two dimensions, the Hitting Set problem with a set P of points is polynomial time solvable via a reduction to Edge Cover in bipartite graphs. Fekete [17] showed that if the set P consists of all points in the plane, then the problem is NP-hard. We sharpen their result to show that if we are additionally given a finite set P of points, then the problem becomes APX-hard.

► **Theorem 7** (3-Line-APX-hardness). *Let $\mathcal{S} = \mathcal{X} \cup \mathcal{Y} \cup \mathcal{Z}$ be three sets of lines such that lines in the plane in each set $\mathcal{X}, \mathcal{Y}$ and $\mathcal{Z}$ are parallel. Let P be a set of points such that each line contains at least one point. The Hitting Set problem defined by $\mathcal{S}$ and P is APX-hard.*

Finally, we consider a restricted setting where each segment intersects segments of the opposite orientation that lie on at most k distinct parallel lines. We call this the k -restricted setting. In this setting, we show that the problem admits a PTAS.

► **Theorem 8** (k -restricted). *Let $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ be a set of horizontal and vertical line segments in the plane such that each segment intersects segments of the opposite orientation lying on at most k parallel lines. Let P be a set of points such that each segment in $\mathcal{S}$ contains at least one point of P . Then, there is a PTAS for the Hitting Set problem with $\mathcal{S}$ and P .*

1.2 Related Work

ϵ -nets. Geometric Hitting Set and Set Cover approximations heavily depend on shape complexity. The ϵ -net framework¹ (see [26] for results on ϵ -nets) connects ϵ -net sizes to LP integrality gaps, where bounded VC-dimension yields $O(\log \text{OPT})$ -approximations [8, 16, 20]. Many geometric set systems have bounded VC-dimension, thus yielding an $O(\log \text{OPT})$ -approximation algorithms. Since the ϵ -net is obtained by uniform random sampling, the approximation algorithms extend to the weighted setting as well. While ϵ -nets of size $O(1/\epsilon)$ are known for some objects such as *pseudodisks*² and halfspaces [25, 31], translating to $O(1)$ -approximation for Hitting Set, these smaller nets do not yield improved approximation algorithms for the weighted problem. If a geometric set system has linear *shallow-cell complexity*, then the *quasi-uniform sampling* technique [36, 12] yield improved approximation algorithms for the weighted problems. However, point-segment systems lack linear shallow-cell complexity.

¹ For a set system $(P, \mathcal{S})$, a set $P' \subseteq P$ is an ϵ -net if it is a hitting set for all sets $S \in \mathcal{S}$ such that $|S| \geq \epsilon|P|$.

² A set $\mathcal{S}$ of objects in the plane, each of whose boundary is a simple Jordan curve is a set of pseudodisks if for any $A, B \in \mathcal{S}$, the boundaries of A and B intersect in either 2 or 0 points.

Piercing and non-piercing. For the weighted setting, there is a QPTAS for the Set Cover problem [27] with non-piercing regions in the plane, but no such result is known for the Hitting Set problem. For the unweighted problem, local search provides PTASes for Hitting Set of pseudodisks, as shown by Mustafa and Ray [28]. The technique also extends to non-piercing regions [7, 32], and also yields PTASes for the Set Cover, Independent Set and Dominating Set problems. Recently, Banik et al. [6] showed that local search also yields PTASes for the Set Cover and Independent Set problems when the objects have a *bounded amount of piercing* [6]. Allowing objects to pierce in an unbounded manner however, renders the Hitting Set and Set Cover problems APX-hard, even for nearly congruent shapes like almost-square rectangles [10] and nearly equilateral triangles [19].

Segments. *Axis-Parallel Covering:* Fekete et al. [17] proved that the Hitting Set problem for axis-parallel segments is APX-hard. For the restricted case where one direction consists of lines, they established NP-hardness and provided a $(5/3)$ -approximation. For Set Cover, Kowalska and Pilipczuk [22] showed APX-hardness for horizontal and vertical segments (even with δ -extensions) and W[1]-hardness for the exact weighted problem. They also developed several FPT algorithms and parameterized approximation schemes parameterized by the solution size k .

Independent Set: For the Independent Set (IS) problem, Marx [24] showed the axis-parallel case is in FPT, whereas arbitrary orientations are W[1]-hard. Parameterized approximation schemes also exist for the axis-parallel setting [14]. Furthermore, Caoduro et al. [9] established a $2 - \epsilon$ lower bound on the ratio between clique cover and independence numbers for segment intersection graphs. Consequently, the natural LP-relaxation for IS on axis-parallel segments has an integrality gap of 2, contrasting with our Hitting Set results.

Arbitrary Orientations: For segments with arbitrary orientations, ϵ -nets have a known size of $O(1/\epsilon \log(1/\epsilon))$ because the VC-dimension is 2. However, Alon [2] and Balogh et al. [5] established an ϵ -net lower bound of $\Omega(1/\epsilon \log^{1/3}(1/\epsilon))$, limiting the best possible ϵ -net-based approximation to $O(\log^{1/3-o(1)} \text{OPT})$. Finally, for the IS problem on arbitrary segments, there is a weighted QPTAS [1] and an unweighted $O(n^\epsilon)$ -approximation [18].

Rectangles. Axis-parallel segments can also be seen as a special case of axis-parallel rectangles. For the Hitting Set problem with axis-parallel rectangles, Aronov, Ezra and Sharir [3] obtained an $O(\log \log \text{OPT})$ -approximation by showing that this set system admits an ϵ -net of size $O(1/\epsilon \log \log 1/\epsilon)$, and this is the best approximation factor known. Pach and Tardos [29] showed that in fact, this bound is tight, and therefore, the natural LP-relaxation for the Hitting Set problem has an integrality gap of $\Omega(\log \log \text{OPT})$.

Organization

Section 2 presents the main algorithm and analysis. Section 3 gives an explicit integrality-gap example. The extension to d -orientations is in Section 4. Section 5 records a separator-based PTAS for bounded-complexity subclasses. Section 6 and Section 7 give the hardness results, and we conclude in Section 8.

2 LP-rounding Algorithm

Let P be a set of n points in the plane with nonnegative weight function $w: P \rightarrow \mathbb{R}_{\geq 0}$. Let $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ be a set of m axis-parallel segments, where $\mathcal{H}$ is the set of horizontal segments and $\mathcal{V}$ is the set of vertical segments. For a segment s , let $P \cap s$ denote the candidate points

14:6 Hitting Axis-Parallel Segments with Weighted Points

lying on s . We always assume $P \cap s \neq \emptyset$ for every $s \in \mathcal{S}$. When we refer to a horizontal or vertical line, we mean a line that contains at least one point of P . The natural LP relaxation for the Hitting Set Problem is

$$\begin{aligned} \min \quad & \sum_{p \in P} w(p) x_p \\ \text{s.t.} \quad & \sum_{p \in P \cap s} x_p \geq 1 \quad \forall s \in \mathcal{S}, \\ & 0 \leq x_p \leq 1 \quad \forall p \in P. \end{aligned}$$

Fix an optimal LP solution x , and let

$$\text{LP} := \sum_{p \in P} w(p) x_p.$$

Since this is a relaxation of the hitting-set problem, $\text{LP} \leq \text{OPT}$.

We use one elementary fact repeatedly: if all candidate points and all segments lie on a single line, then the induced stabbing problem is an interval stabbing instance, whose constraint matrix has the consecutive-ones property and is therefore totally unimodular [34]. Consequently, the one-dimensional problem can be solved exactly in polynomial time.

Algorithm. The algorithm has two phases. In the first phase, we use randomized rounding to obtain a hitting set for the horizontal segments. In doing so, we ensure that each point p is in the hitting set produced with probability x_p , its LP-value. This solution partitions the vertical segments into *blocks*. Each block B is a one-dimensional problem, and the LP solution restricted to B is feasible for the intervals contained in B . Therefore, we can obtain a solution of value at most $\sum_{p \in B} w(p)x_p$. While the algorithm remains the same for the weighted and unweighted problems, the analysis for the unweighted problem can be somewhat sharpened.

We now describe Phase I, where we obtain a hitting set for the horizontal segments. Fix a horizontal line h and let $p_1, \dots, p_k$ be the points on h from left to right. We define the cumulative sums a_i , $i = 0, \dots, k$ as follows:

$$a_0 := 0, \quad a_i := \sum_{j=1}^i x_{p_j} \quad (i = 1, \dots, k),$$

For each i let $I_i := [a_{i-1}, a_i)$. Choose a random shift U_h uniformly from $[0, 1]$. We select point p_i in Phase I if and only if I_i contains an integer of the shifted lattice $U_h + \mathbb{Z}$. We perform this rounding independently for every horizontal line.

We claim that each point p is selected in Phase I with probability x_p , and that the points selected in Phase I constitute a hitting set for the horizontal segments.

► **Lemma 9.** *For each point $p \in P$, the probability that p is selected in Phase I is exactly x_p .*

Proof. If $p = p_i$, then the interval I_i has length $x_{p_i} \leq 1$. As the shift U_h varies uniformly over $[0, 1)$, the set of shifts for which I_i contains a point of $U_h + \mathbb{Z}$ has measure exactly $|I_i| = x_{p_i}$. Hence $\mathbb{P}[p_i \text{ is selected in Phase I}] = x_{p_i}$. ◀

► **Lemma 10.** *Every horizontal segment is hit by Phase I with probability 1.*

Proof. Let $s \in \mathcal{H}$ be a horizontal segment on line h , and suppose it contains the consecutive points $p_i, p_{i+1}, \dots, p_j$. The union of the corresponding intervals is $[a_{i-1}, a_j]$, whose length is

$$a_j - a_{i-1} = \sum_{t=i}^j x_{p_t} \geq 1$$

by the LP constraint for s . Any interval of length at least 1 contains a point of $U_h + \mathbb{Z}$ for every shift $U_h \in [0, 1)$. Therefore at least one point of $s \cap P$ is selected. $\blacktriangleleft$

In Phase II, we solve the residual problem obtained on removing the points selected in Phase I and the segments in $\mathcal{S}$ hit by these points. This residual problem consists of a disjoint collection of vertical segments, each of which is called a *block*. Each block consists of a maximal consecutive run of unselected points on a vertical line and the segments that lie entirely in the block. Each block can be solved optimally either via dynamic programming, or by solving a natural LP-relaxation for the block. We now analyze the cost of the solution returned by the algorithm. First, we observe that the expected cost of the points selected in Phase I is exactly equal to the LP cost.

► **Lemma 11.** $\mathbb{E}[\text{cost}(\text{Phase I})] = \text{LP}$

Proof. $\mathbb{E}[\text{cost}(\text{Phase I})] = \sum_{p \in P} w(p) \mathbb{P}[p \text{ selected in Phase I}] = \sum_{p \in P} w(p) x_p = \text{LP}$, where the second-last equality follows from Lemma 9. $\blacktriangleleft$

We analyze Phase II separately for the weighted case, the unweighted case, and the case where vertical segments are replaced by lines.

2.1 Weighted Hitting Set

Consider a block B in Phase II. Let $x(B)$ denote the sum of the LP values of all the points contained in the block. If $x(B) < 1$, then observe that no segment is completely contained in B . We can therefore, safely ignore such blocks and only deal with blocks with $x(B) \geq 1$. Let the weight of a block B be $w(B) = \sum_{i \in B} w(i) x_i$. For a point $p \in P$, let A_p be the event that p is unselected in Phase I and the block containing p has total LP mass at least 1. Observe that an LP-relaxation for the Hitting Set problem restricted to a block B has an optimal solution at most $w(B)$ if its mass $x(B)$ is at least 1. Therefore,

$$\mathbb{E}[\text{cost}(\text{Phase 2})] \leq \sum_{p \in P} w(p) x_p \mathbb{P}[A_p].$$

We now obtain an upper bound on $\mathbb{P}[A_p]$. Fix a vertical line ℓ with ordered points $1, 2, \dots, m$ from bottom to top, and fix a point $p \in \{1, \dots, m\}$. Define random variables D_p and U_p as follows.

- D_p is the total LP mass of the maximal consecutive run of unselected points ending at p . If p itself is selected in Phase I, then $D_p = 0$.
- U_p is the total LP mass of the maximal consecutive run of unselected points immediately above p .

If p is unselected, then $D_p + U_p$ is exactly the total LP mass of the block containing p . Therefore, $\mathbb{P}[A_p] \leq \mathbb{P}[D_p + U_p \geq 1]$. We now show that the probability that D_p or U_p has mass at least t decreases exponentially.

► **Lemma 12.** For every $t \geq 0$, $\mathbb{P}[D_p \geq t] \leq e^{-t}$ and $\mathbb{P}[U_p \geq t] \leq e^{-t}$. Moreover, D_p and U_p are independent.

Proof. We prove the claim for D_p ; the proof for U_p is identical. Since the claim is trivially true for $t = 0$, we consider $t > 0$. If no suffix of points ending at p has total LP mass at least t , then $\mathbb{P}[D_p \geq t] = 0$. Otherwise, let $j \leq p$ be the largest index such that $\sum_{i=j}^p x_i \geq t$. If $D_p \geq t$, then all points $j, j+1, \dots, p$ must be unselected in Phase I. Since Phase I decisions on a vertical line are independent, $\mathbb{P}[D_p \geq t] \leq \prod_{i=j}^p (1 - x_i) \leq \exp\left(-\sum_{i=j}^p x_i\right) \leq e^{-t}$.

Finally, D_p depends only on the decisions for points at or below p , while U_p depends only on the decisions strictly above p . These sets of decisions are independent, so D_p and U_p are independent. $\blacktriangleleft$

We require the following technical lemma on stochastic dominance (See [35], Theorem 1.A.3, or the full version of our paper [33] for a proof).

► **Lemma 13.** *Let X_1, X_2, Y_1, Y_2 be random variables such that X_1 is independent of X_2 and Y_1 is independent of Y_2 . If $X_1 \preceq_{\text{st}} Y_1$ and $X_2 \preceq_{\text{st}} Y_2$, then $X_1 + X_2 \preceq_{\text{st}} Y_1 + Y_2$.*

Let $E_1, E_2 \sim \text{Exp}(1)$ be independent exponential random variables. The tail bound above implies $D_p \preceq_{\text{st}} E_1$ and $U_p \preceq_{\text{st}} E_2$. By Lemma 13, we have $D_p + U_p \preceq_{\text{st}} E_1 + E_2$. Therefore, $\mathbb{P}[D_p + U_p \geq 1] \leq \mathbb{P}[E_1 + E_2 \geq 1]$. Since $E_1 + E_2$ has the $\Gamma(2, 1)$ distribution, $\mathbb{P}[E_1 + E_2 \geq 1] = e^{-1}(1 + 1) = 2/e$. Hence $\mathbb{P}[A_p] \leq 2/e$ for every point p , and therefore

$$\mathbb{E}[\text{cost}(\text{Phase 2})] \leq \frac{2}{e} \sum_{p \in P} w(p)x_p = \frac{2}{e} \text{LP}.$$

► **Theorem 1 (Weighted).** *Let $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ be a set of horizontal and vertical line segments in the plane and let P be a set of points with a weight function $w : P \rightarrow \mathbb{R}_{\geq 0}$ such that each segment in $\mathcal{S}$ contains at least one point of P . Then, there is a randomized $(1 + 2/e)$ -approximation algorithm for the weighted Hitting Set problem defined by $\mathcal{S}$ and P .*

Proof. Combining the bound on Phase I from Lemma 11, and the bound on A_p above, we obtain

$$\mathbb{E}[\text{cost}(\text{output})] \leq \text{LP} + \frac{2}{e} \text{LP} = \left(1 + \frac{2}{e}\right) \text{LP} \leq \left(1 + \frac{2}{e}\right) \text{OPT}. \quad \blacktriangleleft$$

2.2 Unweighted Hitting Set

Fix a vertical line ℓ containing at least one vertical segment in $\mathcal{V}$. Let the points on this line be $q_1, \dots, q_m$ ordered bottom to top. Let t denote the highest point such that the sum of the LP values of the points above t is at least 1. That is, $t = \max \left\{ j \in \{1, \dots, m\} : \sum_{i=j}^m x_i \geq 1 \right\}$. Since the line contains at least one segment, the LP value on ℓ is ≥ 1 . Therefore, we can assume that t is defined. In the unweighted setting, we can improve the analysis by observing that the value of an optimal solution for a block B is at most $\lfloor \sum_{p \in B} x_p \rfloor$.

For ease of analysis, we introduce a sentinel point 0 immediately below point 1, and declare it always selected. For every $i \in \{0, 1, \dots, m\}$, let $B(i)$ denote the maximal block of consecutive unselected points immediately above i , and define $Y_i := \lfloor x(B(i)) \rfloor$, where $x(B(i)) := \sum_{p \in B(i)} x_p$. Every block on the line is either $B(0)$ or a block $B(i)$ following a selected point $i \geq 1$. Thus the Phase II cost on this line is at most $Y_0 + \sum_{i=1}^m \mathbf{1}[i \text{ selected}] Y_i$, where $\mathbf{1}[i \text{ selected}]$ is the indicator random variable that is 1 if i is selected, and 0 otherwise. Moreover, if $i \geq t$, then $\sum_{j=i+1}^m x_j < 1$, so $Y_i = 0$. Therefore only $Y_0, Y_1, \dots, Y_{t-1}$ matter.

► **Lemma 14.** *For every $i \in \{0, 1, \dots, t-1\}$, $\mathbb{E}[Y_i \mid i \text{ is selected}] \leq \frac{1}{e-1}$, where for $i = 0$ the conditioning is vacuous.*

Proof. Fix $i \in \{0, 1, \dots, t-1\}$. For every integer $k \geq 1$,

$$\mathbb{P}[x(B(i)) \geq k \mid i \text{ is selected}] \leq e^{-k}.$$

Indeed, if the above probability is non-zero, let $r > i$ be the smallest index such that $\sum_{j=i+1}^r x_j \geq k$. Then all points $i+1, \dots, r$ must not be selected in Phase I. Since the rounding decisions on a vertical line are independent and the event that i is selected depends only on point i , conditioning on the selection of i does not affect the points above it. Thus

$$\mathbb{P}[x(B(i)) \geq k \mid i \text{ is selected}] \leq \prod_{j=i+1}^r (1 - x_j) \leq \exp\left(-\sum_{j=i+1}^r x_j\right) \leq e^{-k}.$$

Using the tail-sum formula³,

$$\begin{aligned} \mathbb{E}[Y_i \mid i \text{ is selected}] &= \sum_{k \geq 1} \mathbb{P}[Y_i \geq k \mid i \text{ is selected}] \leq \sum_{k \geq 1} \mathbb{P}[x(B(i)) \geq k \mid i \text{ is selected}] \\ &\leq \sum_{k \geq 1} e^{-k} = \frac{1}{e-1}. \end{aligned} \quad \blacktriangleleft$$

► **Lemma 15.** *For every vertical line, $\mathbb{E}[\text{cost}(\text{Phase II on this line})] \leq \frac{1}{e-1} \sum_{i=1}^m x_i$.*

Proof. Using the block decomposition above,

$$\mathbb{E}[\text{cost}(\text{Phase II on this line})] \leq \mathbb{E}[Y_0] + \sum_{i=1}^{t-1} \mathbb{P}[i \text{ selected}] \mathbb{E}[Y_i \mid i \text{ selected}].$$

By Lemma 14, every conditional expectation is at most $1/(e-1)$, and point i is selected with probability x_i . Hence,

$$\mathbb{E}[\text{cost}(\text{Phase II on this line})] \leq \frac{1}{e-1} + \sum_{i=1}^{t-1} \frac{x_i}{e-1} = \frac{1 + \sum_{i=1}^{t-1} x_i}{e-1}.$$

By the definition of t , we have $\sum_{i=t}^m x_i \geq 1$, so $1 + \sum_{i=1}^{t-1} x_i \leq \sum_{i=1}^m x_i$. Substituting this inequality completes the proof. ◀

► **Theorem 2 (Unweighted).** *Let $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ be a set of horizontal and vertical line segments in the plane and let P be a set of points such that each segment in $\mathcal{S}$ contains at least one point of P . Then, there is a randomized $(1 + 1/(e-1))$ -approximation algorithm for the Hitting Set problem defined by $\mathcal{S}$ and P .*

Proof. Phase I still has expected cost exactly $\text{LP} = \sum_{p \in P} x_p$. Summing the previous line bound over all vertical lines gives $\mathbb{E}[\text{cost}(\text{Phase 2})] \leq 1/(e-1) \text{LP}$. Therefore

$$\mathbb{E}[\text{cost}(\text{output})] \leq \text{LP} + \frac{1}{e-1} \text{LP} = \left(1 + \frac{1}{e-1}\right) \text{LP} \leq \left(1 + \frac{1}{e-1}\right) \text{OPT}. \quad \blacktriangleleft$$

The analysis above is asymptotically tight as we show below.

³ If X is a non-negative integer-valued random variable ($X \in \{0, 1, 2, \dots\}$), the expected value is: $\mathbb{E}[X] = \sum_{n=1}^{\infty} \mathbb{P}(X \geq n)$

2.3 Tightness of the Unweighted Analysis

We record a family of instances showing that the analysis in Section 2.2 is asymptotically tight. Let k tend to infinity, and set $n := k^2$. The point set is the $n \times n$ grid $P = [n] \times [n]$. The input contains every horizontal and vertical segment consisting of exactly k consecutive grid points.

First, $\text{OPT} = n^2/k = k^3$. Indeed, the k disjoint length- k horizontal segments in each row show that every feasible solution has size at least $nk = k^3$. Conversely, the set

$$Q^* := \{(i, j) \in [n] \times [n] : j - i \equiv 0 \pmod{k}\}$$

has exactly k points in every row and every column, with consecutive chosen points spaced exactly k apart. Hence every length- k horizontal or vertical segment contains one point of Q^* , so $|Q^*| = k^3 = n^2/k$.

The same disjoint-segment lower bound also applies to the LP, and the uniform assignment $x_p = 1/k$ for every grid point is feasible. Hence this assignment is an optimal LP solution. We now run the algorithm with this symmetric optimum. On each row, Phase I chooses exactly one residue class modulo k , independently and uniformly from row to row. Thus Phase I selects exactly k points on each row, and

$$\text{cost}(\text{Phase I}) = k^3 = \text{OPT}.$$

It remains to compute the expected repair cost on the columns. Fix a column. A row is already selected in this column with probability $p := 1/k$, independently across rows. Let $q := 1 - p$ and let Z be the Phase II repair cost on this column. If a maximal run of unselected points has length L , then the residual vertical segments inside this run are exactly the length- k intervals contained in the run, and the optimum one-dimensional repair cost is $\lfloor L/k \rfloor$: the lower bound comes from the disjoint length- k intervals in the run, and selecting every k th point attains it.

For $a \geq 1$, let N_a be the number of unselected runs in the column whose length is at least ak . Then $Z = \sum_{a=1}^k N_a$. A run of length at least ak starts in the first row with probability q^{ak} ; it starts at any later row with probability pq^{ak} , requiring the previous row to be selected and the next ak rows to be unselected. Therefore, with $\alpha := q^k = (1 - 1/k)^k$,

$$\mathbb{E}[Z] = \sum_{a=1}^k (q^{ak} + (n - ak)pq^{ak}) = \sum_{a=1}^k (k - a + 1)\alpha^a \geq k \frac{\alpha(1 - \alpha^k)}{1 - \alpha} - \frac{\alpha}{(1 - \alpha)^2},$$

using $\sum_{a=1}^k (a-1)\alpha^a \leq \sum_{a=1}^{\infty} a \cdot \alpha^a = \alpha/(1-\alpha)^2$. Now, there are k^2 columns each contributing $\mathbb{E}[Z]$ to the expected cost of Phase II and $\text{OPT} = k^3$. Hence, the ratio of the expected cost of Phase II and OPT satisfies:

$$\frac{\alpha(1 - \alpha^k)}{1 - \alpha} - \frac{1}{k} \frac{\alpha}{(1 - \alpha)^2} \leq \frac{\mathbb{E}[\text{cost}(\text{Phase II})]}{\text{OPT}} \leq \frac{1}{e - 1}, \quad (1)$$

where the second inequality follows from our analysis of the unweighted case in Section 2.2. Now, as $k \rightarrow \infty$, we have $\alpha = (1 - 1/k)^k \rightarrow e^{-1}$, thus the left hand side of (1) tends to $e^{-1}/(1 - e^{-1}) = 1/(e - 1)$. Thus, the approximation factor $(1 + 1/(e - 1))$ obtained in the unweighted analysis is asymptotically tight for the algorithm.

2.4 Horizontal Segments and Vertical Lines

In this section, we show that the approximation factor can be improved when in one of the directions, we have lines instead of line segments.

► **Theorem 5** (Lines and Segments). *Let $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ be a set of horizontal line segments and vertical lines in the plane and let P be a set of points with a weight function $w : P \rightarrow \mathbb{R}_{\geq 0}$ such that each object in $\mathcal{S}$ contains at least one point of P . Then, there is a randomized $(1 + 1/e)$ -approximation algorithm for the weighted Hitting Set problem defined by $\mathcal{S}$ and P .*


Proof. Each vertical line ℓ that is not hit by the Phase I points contributes $\min_{p \in \ell} w(p) \leq \sum_{p \in \ell} w(p)x_p$ to the Phase II cost, and each vertical line that is hit by the Phase I points contributes 0. Thus,

$$\mathbb{E}[\text{cost}(\text{Phase II})] \leq \sum_{\ell \in \mathcal{V}} \mathbb{P}[\ell \text{ is not hit by Phase I}] \sum_{p \in \ell} w(p)x_p.$$

Now, since $\sum_{p \in \ell} x_p \geq 1$, we have that $\mathbb{P}[\ell \text{ is not hit by Phase I}] \leq \frac{1}{e}$, and hence the total expected cost of Phase II is upper bounded by $\frac{\text{LP}}{e}$. Thus, from Lemma 11, we get that the total expected cost is at most $(1 + 1/e)\text{LP}$. ◀

3 An Explicit Integrality-Gap Construction

In this section, we present an instance on which the integrality gap of the standard LP for the unweighted problem is 1.25.



■ **Figure 1** An integrality-gap instance for unweighted axis-parallel segment hitting. Segments are drawn with slight offsets for readability. The blue segments denote the horizontal segments $\mathcal{H}$, and the red segments denote the vertical segments $\mathcal{V}$.

Let $P = \{0, 1, \dots, 15\}$ be the candidate points shown in Figure 1. Segments in the instance are formed by pairs of points that are on the same horizontal or vertical level, and have no other point on the segment connecting them. Thus, along the top chain we include every segment corresponding to consecutive pairs of points from $\{0, 1, 2, 3, 4, 5\}$, along the bottom chain every segment corresponding to consecutive pairs of points from $\{10, 11, 12, 13, 14, 15\}$, and so on.

► **Theorem 16.** *The instance of Figure 1 has integrality gap exactly $5/4$ for the natural LP relaxation.*

Proof. Our hitting set instance is precisely an instance of vertex cover where the vertices are the points and for each segment, there is an edge between the vertices corresponding to the two points on the segment. Let G denote this graph. Set $x_p := 1/2$ for every point $p \in P$. Every segment in the construction contains at least two marked points, so this assignment is feasible. Hence the LP value is at most $\sum_{p \in P} x_p = 16 \cdot \frac{1}{2} = 8$.

14:12 Hitting Axis-Parallel Segments with Weighted Points

In fact this bound is tight. The eight pairwise vertex-disjoint segments $\{0, 1\}$, $\{2, 3\}$, $\{4, 5\}$, $\{6, 7\}$, $\{8, 9\}$, $\{10, 11\}$, $\{12, 13\}$, $\{14, 15\}$ already force $\sum_{p \in P} x_p \geq 8$, so $\text{LP} = 8$. It remains to compute the optimum integral value. Since this is a vertex-cover instance,

$$\text{OPT} = |V(G)| - \alpha(G) = 16 - \alpha(G),$$

where $\alpha(G)$ is the maximum independent-set size. The set $\{0, 2, 4, 8, 12, 14\}$ is independent, so $\alpha(G) \geq 6$. For the reverse inequality, consider the 6-cycle $C = \{2, 3, 9, 13, 12, 6\}$ and the two 5-cycles $L = \{0, 1, 8, 11, 10\}$ and $R = \{4, 5, 15, 14, 7\}$. Every independent set contains at most 3 vertices from C and at most 2 vertices from each of L and R . If an independent set I has $|I \cap C| \leq 2$, then immediately $|I| \leq 2 + 2 + 2 = 6$.

Suppose instead that $|I \cap C| = 3$. Then $I \cap C$ must be one of the two alternating triples $\{2, 9, 12\}$ or $\{3, 6, 13\}$. If I contains $\{2, 9, 12\}$, then the edges $(1, 2)$, $(8, 9)$, and $(11, 12)$ exclude 1, 8, 11 from the left 5-cycle. Only 0 and 10 remain available in L , and they are adjacent, so $|I \cap L| \leq 1$. Since $|I \cap R| \leq 2$, we get $|I| \leq 3 + 1 + 2 = 6$. The case of the other alternating triple $\{3, 6, 13\}$ is symmetric.

Therefore $\alpha(G) = 6$, and hence $\text{OPT} = 16 - 6 = 10$. The integrality gap is

$$\frac{\text{OPT}}{\text{LP}} = \frac{10}{8} = \frac{5}{4}.$$

4 Generalization to d orientations (weighted case)

Suppose segments are allowed to have d distinct orientations $\theta_1, \theta_2, \dots, \theta_d$, and let $S^{(r)}$ denote the set of segments of orientation θ_r . The LP relaxation is unchanged:

$$\min \sum_{p \in P} w(p)x_p \quad \text{s.t.} \quad \sum_{p \in P \cap s} x_p \geq 1 \quad \forall s \in \bigcup_{r=1}^d S^{(r)}, \quad 0 \leq x_p \leq 1.$$

Let x be an optimal LP solution with value LP .

Algorithm with a parameter k . Fix $k \in \{1, 2, \dots, d\}$ and call $\theta_1, \dots, \theta_k$ the primary orientations. For every line of each primary orientation, independently run the same random-start systematic rounding used above, after rotating coordinates along that line. The output of the primary step is the union of all points selected by these k rounds. Therefore its expected cost is at most

$$\sum_{p \in P} w(p) \mathbb{P}[p \text{ is selected in at least one primary round}] \leq \sum_{p \in P} w(p) \sum_{t=1}^k x_p = k \cdot \text{LP}.$$

All segments in the primary orientations are hit with probability one.

For every remaining orientation $r > k$, we then repair the residual one-dimensional instances on lines of orientation θ_r . These repairs may be performed independently for all remaining orientations and then unioned with the primary solution. This can only overcount the final cost, so it is enough to bound the sum of the repair costs orientation by orientation.

► **Lemma 17.** Fix a remaining orientation θ_r and a point p on a line of orientation θ_r . Let A_p be the event that p survives the primary step and lies in a surviving block of LP mass at least 1 along that line. Then $\mathbb{P}[A_p] \leq (k+1)e^{-k}$.

Proof. Fix a line ℓ of orientation θ_r and order its points along ℓ . In one primary orientation, a point i survives with probability $1 - x_i$; hence after all k primary rounds it survives with probability $(1 - x_i)^k$. Moreover, for any primary orientation θ_t and any two distinct points on ℓ , the two θ_t -lines through these points are distinct; otherwise a θ_t -line would meet ℓ in two points, forcing it to be ℓ , contrary to $\theta_t \neq \theta_r$. Thus the collections of primary-line shifts determining the survival of distinct points on ℓ are disjoint and independent, so the survival events of distinct points on ℓ are independent.

Now repeat the weighted-case block analysis from the proof of Lemma 12 with “survives the primary step” in place of “is unselected in Phase I”. Let D_p be the LP mass of the maximal surviving run ending at p , and let U_p be the LP mass of the maximal surviving run immediately after p . If A_p occurs, then $D_p + U_p \geq 1$.

For any $t \geq 0$, the same suffix argument as in the proof of Lemma 12 gives

$$\mathbb{P}[D_p \geq t] \leq \prod_{i \in J} (1 - x_i)^k \leq \exp\left(-k \sum_{i \in J} x_i\right) \leq e^{-kt},$$

where J is a minimal consecutive suffix ending at p with LP mass at least t ; if no such suffix exists, the probability is 0. Similarly, $\mathbb{P}[U_p \geq t] \leq e^{-kt}$, and D_p, U_p are independent. Thus D_p and U_p are stochastically dominated by independent exponential random variables E_1, E_2 of rate k . Therefore,

$$\mathbb{P}[A_p] \leq \mathbb{P}[D_p + U_p \geq 1] \leq \mathbb{P}[E_1 + E_2 \geq 1] = (k + 1)e^{-k},$$

where the second inequality follows from Lemma 13. ◀

Thus, we get that for each remaining orientation $r > k$,

$$\mathbb{E}[\text{repair cost for orientation } \theta_r] \leq \sum_{p \in P} x_p w(p) \mathbb{P}[A_p] \leq (k + 1)e^{-k} \text{LP}.$$

Summing over the $d - k$ remaining orientations and adding the primary cost yields

$$\mathbb{E}[\text{total cost}] \leq (k + (d - k)(k + 1)e^{-k}) \text{LP}. \quad (2)$$

► **Theorem 3 (d-Directions).** *Let $\mathcal{S}$ be a set of segments, each parallel to one of d directions in the plane for some $d \geq 2$, and let P be a set of points with a weight function $w : P \rightarrow \mathbb{R}_{\geq 0}$ such that each segment in $\mathcal{S}$ contains at least one point of P . Then, there is a randomized algorithm that is a $(\ln d + \ln \ln d + O(1))$ -approximation relative to the natural LP.*

Proof. Take $k = \min(d, \lceil \ln d + \ln \ln d \rceil)$. For large enough d , k equals $\lceil \ln d + \ln \ln d \rceil$, and thus $(d - k)(k + 1)e^{-k} = O(1)$, yielding the claimed approximation ratio when substituted in (2). ◀

We can further generalize this result to the case when each object to be hit is a union of up to h segments, each of which has one of d orientations.

► **Corollary 18.** *There exists a randomized algorithm that, given an instance of the weighted Hitting Set problem where each object is a union of at most h segments of d distinct orientations, computes a solution with an expected approximation factor of $h \cdot (\ln d + \ln \ln d + O(1))$ relative to the natural LP.*

Proof. Let x be a feasible solution to the natural LP. Since each object O is a union of at most h segments, the LP constraint $\sum_{p \in O} x_p \geq 1$ implies there is some segment $s \subseteq O$ such that $\sum_{p \in s} x_p \geq 1/h$. Choose one such segment for each object, and define $y_p := \min\{1, hx_p\}$. Then $0 \leq y_p \leq 1$, $\sum_p w(p)y_p \leq h \cdot \text{LP}$, and for every chosen segment s we have $\sum_{p \in s} y_p \geq 1$: if some $p \in s$ has $hx_p \geq 1$ this is immediate, and otherwise $\sum_{p \in s} y_p = h \sum_{p \in s} x_p \geq 1$. Thus y is feasible for the chosen segment instance, and hitting these chosen segments hits every original object. Running the algorithm for d orientations on y yields the claimed expected approximation factor. $\blacktriangleleft$

5 Local Search PTASes for k -restricted instances

In this section, we show that if we restrict the input so that each segment intersects segments of the opposite orientation lying on at most k -distinct lines, then we obtain a PTAS for the unweighted Hitting Set problem. The results of Raman and Ray [32] only work if the objects are non-piercing. While Banik et al. [6] works for the case where each object has a bounded amount of piercing, their results only work for the Set Cover and Independent Set problems. The restriction on the number of such intersections is essential, as otherwise, as Section 6 shows, the problem is APX-hard.

► **Definition 19.** A family $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ of axis-parallel segments is k -restricted if every segment in $\mathcal{H}$ intersects segments of $\mathcal{V}$ lying on at most k parallel lines, and every segment in $\mathcal{V}$ intersects segments of $\mathcal{H}$ lying on at most k parallel lines.

Observe that in a k -restricted instance, a segment can intersect an unbounded number of segments of the opposite orientation. The only restriction is that these segments lie on at most k distinct parallel lines.

The local-search algorithm is standard. Fix an integer parameter t . Starting from any feasible hitting set, repeatedly replace at most t objects of the current solution by a strictly smaller feasible set, as long as such an improving swap exists. For fixed t , this process terminates in polynomial time. The key issue is to prove that a t -locally optimal solution is globally near-optimal when $t = t(\epsilon)$ is chosen large enough [32, 4, 13, 28].

The analysis relies on showing the existence of a suitable *exchange graph* $G = (L \cup O, E)$ on a locally optimal solution L and an optimal solution O . This exchange graph has two properties: the *local property* that for any $L' \subseteq L$, $(L' \setminus L) \cup N_G(L')$ is also a hitting set for $\mathcal{S}$, where $N_G(L')$ are the neighbors of L' in the exchange graph G , and the *separator property* that G comes from a hereditary family $\mathcal{G}$ of graphs with sub-linear size separators. That is, there is an absolute constant $\epsilon > 0$ such that any graph $H \in \mathcal{G}$ has a separator of size $O(|V(H)|^{1-\epsilon})$.

Given two disjoint hitting sets R and B for a set $\mathcal{S}$ of axis-parallel segments, we show the existence of a suitable exchange graph. A PTAS then follows by standard arguments.

► **Lemma 20.** Let $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ be a set of k -restricted horizontal and vertical segments and let R and B be two disjoint hitting sets for the segments $\mathcal{S}$. Then, there is a local-exchange graph $G = (R \cup B, E)$ that satisfies the local exchange and the separator properties.

⁴ A balanced separator, or just a separator for a graph $G = (V, E)$ is a subset $S \subseteq V$ of vertices whose removal partitions the graph such that no resulting connected component contains more than $\alpha|V|$ vertices, where $\alpha < 1$ is a fixed constant, typically $1/2$ or $2/3$. A graph admits a sub-linear separator if there exists an absolute constant $\epsilon > 0$ such that $|S| = O(|V|^{1-\epsilon})$.

Proof. The graph $G = (R \cup B, E)$ is constructed as follows. Starting with the empty graph on $R \cup B$, for each segment $s \in \mathcal{S}$, let R_s and B_s denote respectively, the points of R in s and the points of B in s . Suppose no point in R_s is adjacent to a point in B_s , we choose an arbitrary pair of points, one in R_s and the other in B_s that are consecutive along s and add the edge $\{r, b\}$ to G as a straight-line segment along s . We claim that G is the desired local-exchange graph. We say that the points r and s are *responsible* for hitting s .

Let $R' \subseteq R$. We claim that $T = (R \setminus R') \cup N_G(R')$ is a hitting set. Indeed, for a segment s , if the vertex r responsible for hitting s is not selected in R' , then s is still hit by a point in the set T . If not, since the vertex $b \in B$ responsible for hitting s is adjacent to r , it implies that s is hit by b . Therefore, G satisfies the local-exchange property.

Consider the natural embedding of G in the plane, where the edges are drawn as straight-line segments between the points. By construction, each edge in G is only intersected by edges of the opposite orientation. Since any edge lies along a segment, and the instance is k -restricted, it implies that each edge of G is crossed by at most k other edges. That is G is a k -planar graph⁵. k -planar graphs form a hereditary family as removing vertices leaves the induced graph k -planar.

Since k -planar graphs have $O(\sqrt{kn})$ edges [30], where $n = |R| + |B|$ and any two segments intersect in at most one point we obtain a plane graph H with $O(k^{3/2}n)$ vertices, by splitting each pair of crossing edges by adding a new *splitting vertex* at their intersection point. We assign a weight of 1 to each vertex of G and a weight of 0 to each splitting vertex. By the Planar Separator Theorem [23], there is a balanced separator S of H of size $O(\sqrt{|V(H)|})$ such that both sides of the partition has weight at most $2/3$ the total weight of the vertices in H . To obtain a separator in G , for each vertex in S , if a splitting vertex s lies in S , we add the four end-points corresponding to the pair of crossing edges defined by s . The resulting set S' is a balanced separator in G , and has size $O(|S|)$. Therefore, we obtain a sub-linear sized separator in G . ◀

► **Theorem 8** (k -restricted). *Let $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ be a set of horizontal and vertical line segments in the plane such that each segment intersects segments of the opposite orientation lying on at most k parallel lines. Let P be a set of points such that each segment in $\mathcal{S}$ contains at least one point of P . Then, there is a PTAS for the Hitting Set problem with $\mathcal{S}$ and P .*

Proof sketch. Let L be a t -locally optimal Hitting Set and let O be an optimal Hitting Set. We can assume that $L \cap O = \emptyset$, as otherwise, we can remove the points in $L \cap O$ and the segments hit by them, and consider the residual instance. This can only improve the approximation factor. By Lemma 20, there is an exchange graph satisfying the local-exchange property. Now, a PTAS follows directly via the techniques in [28]. ◀

6 APX-hardness

The proof of the following theorem is based on the proof of APX-hardness of Minimum Hitting Set for orthogonal segments by Fekete et al. [17]. Fekete et al. show NP-hardness for the *piercing problem*, i.e., when we are allowed to choose any point in $\mathbb{R}^2$ to hit the segments. On the other hand, we show that if P is restricted to a discrete set, the problem becomes APX-hard.

⁵ A graph is k -planar if it has an embedding in the plane such that each edge is crossed by at most k other edges


► **Theorem 6** (APX-Hardness). *Let $\mathcal{S} = \mathcal{H} \cup \mathcal{V}$ be a set of horizontal line segments and vertical lines in the plane and P be a finite set of points in the plane such that each object in $\mathcal{S}$ contains at least one point of P . The Hitting Set problem with $\mathcal{S}$ and P is APX-hard.*

Proof. We give an approximation-preserving reduction from MAX-2SAT(3), the maximum satisfiability problem for CNF formulae with two literals per clause, in which each variable appears in at most three clauses. This problem is APX-hard [21].

Let Φ be an instance of MAX-2SAT(3) with variables $x_1, \dots, x_n$ and clauses $C_1, \dots, C_m$. We may assume that every variable appears in at least one clause; isolated variables can be deleted. Since each clause contains two literals, the total number of literal occurrences is $2m$; therefore, $n \leq 2m$.

Let OPT denote the maximum number of clauses of Φ that can be satisfied. Also, since a uniformly random truth assignment satisfies each 2-literal-clause with probability at least $1/2$, we have $\text{OPT} \geq \frac{m}{2}$.

We construct an instance of Minimum Hitting Set as follows. For each variable x_i , let $d_i \in \{1, 2, 3\}$ be the number of occurrences of x_i in Φ . We construct a variable gadget G_i consisting of $2d_i + 4$ geometric objects: $2d_i + 2$ horizontal segments and two vertical lines. For each variable x_i , there is vertical line L_i at the left end and there is a vertical line R_i at the right end. These objects are arranged so that their intersection graph is an even cycle of length $2d_i + 4$, and no point intersects three objects of the variable gadget. Equivalently, the only points that hit two objects of G_i are the consecutive intersections along this cycle. Refer to Figure 2.



Moreover, in a hitting set of size exactly b_i , every selected point must lie at an intersection of two consecutive objects of G_i . Therefore, it corresponds to a perfect matching of the cycle. Thus there are exactly two canonical optimum hitting sets for G_i : the set of all red points and the set of all green points. We interpret the green solution as setting $x_i = \text{true}$ and the red solution as setting $x_i = \text{false}$.

For each clause C_j , introduce a vertical line ℓ_j . If the literal x_i occurs in C_j , we make ℓ_j pass through the green point of G_i reserved for this occurrence. If the literal $\bar{x}_i$ occurs in C_j , we make ℓ_j pass through the corresponding red point of G_i . The lines and input points are placed such that these are the only points at which a clause line passes through. Thus, each clause line passes through exactly two input points. Let $B := \sum_{i=1}^n b_i$. Since $\sum_i d_i = 2m$, we get $B = 6m$.

We first relate satisfying assignments to hitting sets. Given an optimal truth assignment satisfying OPT clauses, choose in each variable gadget the canonical red or green hitting set corresponding to the truth value of the variable. This uses exactly B points. Every satisfied clause line is hit at the corresponding literal point. Each unsatisfied clause line can be hit by one additional point placed anywhere on that line. Therefore, if k denotes the optimum solution size of the constructed hitting-set instance, then $k \leq B + (m - \text{OPT})$. In particular, $k \leq B + m \leq 7m$.

Conversely, let H be any hitting set for the constructed instance. We normalize H without increasing its cardinality. Consider a variable gadget G_i . Since G_i has $2b_i$ objects and no point hits more than two objects of G_i , the points of H used to hit objects of G_i have cardinality at least b_i .

If the selected points of H for the variable gadget G_i are consistent, that is, all are red or all are green, then we replace the points of H used inside G_i by the corresponding canonical optimum set of b_i red or green points. This does not increase the number of points.

If the selected literal points are inconsistent, then H uses literal points of both colors. Since the two canonical solutions are the only size- b_i hitting sets for G_i , the portion of H used for G_i has size at least $b_i + 1$. We replace it by one canonical optimum set of size b_i , choosing the majority color among the selected literal points. Since $d_i \leq 3$, at most one clause line that was previously hit through a minority-colored literal point becomes unhit. We use the one saved point to hit that clause line separately. Again, the total number of points does not increase.

Applying this operation independently to all variable gadgets, we obtain a hitting set of the form $H_v \cup H_c$, with $|H_v| = B$ and $|H_v| + |H_c| \leq |H|$.

The set H_v is canonical in every variable gadget, and hence defines a truth assignment. The set H_c consists of additional points placed on clause lines. Every clause not satisfied by the assignment induced by H_v must be hit by a point of H_c . Therefore, if alg is the number of clauses satisfied by the induced assignment, then

$$\text{alg} \geq m - |H_c|.$$

Now suppose there is a polynomial-time $(1+\varepsilon)$ -approximation algorithm for the hitting-set problem on horizontal segments and vertical lines. Applied to the constructed instance, it returns a hitting set H with $|H| \leq (1+\varepsilon)k$, and after the normalization above, $|H_c| \leq |H| - B \leq (1+\varepsilon)k - B$. Hence,

$$\text{alg} \geq m - |H_c| \geq m - (1+\varepsilon)k + B = (m + B - k) - \varepsilon k.$$

Since $k \leq B + (m - \text{OPT})$, we have $m + B - k \geq \text{OPT}$. Thus, $\text{alg} \geq \text{OPT} - \varepsilon k$. Using $k \leq 7m$ and $\text{OPT} \geq m/2$, we obtain

$$\varepsilon k \leq 7\varepsilon m \leq 14\text{OPT}.$$

Therefore,

$$\text{alg} \geq (1 - 14\varepsilon) \text{OPT}.$$

Consequently, a PTAS for minimum hitting set with horizontal segments and vertical lines would imply a PTAS for MAX-2SAT(3), contradicting its APX-hardness. Hence the hitting-set problem for horizontal segments and vertical lines is APX-hard. ◀

7 Discrete 3-Slope Line Cover is APX-hard

DISCRETE 3-SLOPE-LINE-COVER (3SLC) refers to the problem of computing a minimum-cardinality hitting set for a set $\mathcal{S}$ of lines having three distinct slopes with a subset of a given set P of points. In this section, we prove that this problem is APX-hard [17].

► **Definition 21.** (*L-reduction*) [11] A pair of functions f and g , both computable in polynomial time, is an *L-reduction* from an optimization problem $\mathcal{P}$ to an optimization problem $\mathcal{P}'$ if there are positive constants α and β such that for each instance x of $\mathcal{P}$, the following hold:

- (L1) The function f maps instances of $\mathcal{P}$ to instances of $\mathcal{P}'$ such that $\text{OPT}(f(x)) \leq \alpha \cdot \text{OPT}(x)$.
- (L2) The function g maps feasible solutions y of $f(x)$ to feasible solutions $g(y)$ of x such that $|c_x(g(y)) - \text{OPT}(x)| \leq \beta \cdot |c_{f(x)}(y) - \text{OPT}(f(x))|$, where c_x and $c_{f(x)}$ are the cost functions of the instances x and $f(x)$ respectively.

► **Theorem 7** (3-Line-APX-hardness). Let $\mathcal{S} = \mathcal{X} \cup \mathcal{Y} \cup \mathcal{Z}$ be three sets of lines such that lines in the plane in each set $\mathcal{X}, \mathcal{Y}$ and $\mathcal{Z}$ are parallel. Let P be a set of points such that each line contains at least one point. The Hitting Set problem defined by $\mathcal{S}$ and P is APX-hard.


Proof. We give an *L-reduction* from SPECIAL-3SC, which is APX-hard [11]. Recall that an instance of SPECIAL-3SC is obtained from a cubic graph $G = (V, E)$. For every vertex $v_t \in V$, there are four elements w_t, x_t, y_t, z_t , and for every edge $e_i \in E$, there is an element a_i . If the three edges incident to v_t are e_i, e_j, e_k , with $i < j < k$, then the instance contains the five sets $\{w_t, a_i\}, \{w_t, x_t\}, \{x_t, y_t, a_j\}, \{y_t, z_t\}, \{z_t, a_k\}$. Moreover, every element a_i belongs to exactly two sets.

We first construct an equivalent finite discrete 3SLC instance in the dual plane. Let L_1, L_2, L_3 be three vertical lines, from left to right. We place the elements w_t, y_t on L_1 , the elements x_t, z_t on L_2 , and the elements a_i on L_3 . The placement is chosen so that, for each v_t with incident edges e_i, e_j, e_k , $i < j < k$, there are five non-vertical candidate lines whose intersections with the constructed point set are exactly

$$\{w_t, a_i\}, \quad \{w_t, x_t\}, \quad \{x_t, y_t, a_j\}, \quad \{y_t, z_t\}, \quad \{z_t, a_k\}.$$

We briefly justify that such a placement can be found in polynomial time. Put the edge points $a_1, \dots, a_n$ on L_3 in this order. For each vertex v_t , reserve two small open intervals $I_t^1 \subset L_1$ and $I_t^2 \subset L_2$, with these intervals pairwise disjoint over different vertices, such that every line joining a point of I_t^1 to an edge point a_r meets L_2 inside I_t^2 . Such intervals are obtained, for example, by taking disjoint thin vertical cross-sections of triangles with common base on L_3 .

Now choose $y_t \in I_t^1$ and let x_t be the intersection of $\overline{y_t a_j}$ with L_2 . Choose $w_t \in I_t^1$ and $z_t \in I_t^2$. At each step, the forbidden choices are only those causing one of the five intended candidate lines to pass through an unintended point. There are finitely many such forbidden positions. Hence we can choose rational points avoiding all of them. See Figure 3.



■ **Figure 3** Two vertex gadgets for $v_s, v_t \in V$ in the dual construction, where $(v_s, v_t) \in E$. For each vertex v_t incident to edges e_i, e_j, e_k , $i < j < k$, the points w_t, y_t lie on L_1 and x_t, z_t lie on L_2 . The three solid lines encode the sets $\{w_t, a_i\}$, $\{x_t, y_t, a_j\}$, and $\{z_t, a_k\}$; the dashed magenta lines encode the non-edge sets $\{w_t, x_t\}$ and $\{y_t, z_t\}$.

It follows that the constructed dual finite line-cover instance has exactly the same set system as the SPECIAL-3SC instance: selecting a candidate line is equivalent to selecting the corresponding set, and a subfamily covers all points if and only if the corresponding sets cover the universe. Therefore

$$\text{OPT}_{\text{dual}} = \text{OPT}_{\text{SC}},$$

where OPT_{dual} denotes the optimal solution value for the dual line-cover instance and OPT_{SC} denotes the optimal solution value for the SPECIAL-3SC instance.

It remains to pass from this dual finite line-cover instance to DISCRETE 3-SLOPE-LINE-COVER. We use the standard point-line duality

$$p = (p_x, p_y) \longleftrightarrow p^* : b = p_x a - p_y,$$

and

$$\ell : y = ax - b \longleftrightarrow \ell^* = (a, b),$$

which preserves incidences. The three vertical lines L_1, L_2, L_3 in the dual plane become three slopes in the primal plane. The dual points become the input lines S , and the dual candidate lines become the finite candidate point set P . Hence we obtain an instance of DISCRETE

3-SLOPE LINE COVER. Since incidence and cardinality are preserved, feasible covers in the dual line-cover instance are in one-to-one correspondence with feasible hitting sets in the resulting DISCRETE 3-SLOPE-LINE-COVER instance.

Thus, for the reduction f ,

$$\text{OPT}_{3\text{SLC}}(f(I)) = \text{OPT}_{\text{SC}}(I),$$

so condition (L1) of an L -reduction holds with $\alpha = 1$. Given any feasible solution to $f(I)$, mapping each selected point back by duality gives the corresponding selected set of the SPECIAL-3SC instance, with the same cardinality. Therefore

$$|c_{\text{SC}}(g(Y)) - \text{OPT}_{\text{SC}}(I)| = |c_{3\text{SLC}}(Y) - \text{OPT}_{3\text{SLC}}(f(I))|,$$

so condition (L2) holds with $\beta = 1$. The construction is polynomial-time computable. Hence DISCRETE 3-SLOPE-LINE-COVER is APX-hard. ◀

8 Conclusion

We gave a short LP-rounding framework for hitting axis-parallel segments with weighted points. The algorithm improves the immediate factor-2 bound by rounding one orientation exactly and repairing the residual subproblem on the other orientation using exact one-dimensional optimization. This yields a deterministic $(1+2/e)$ -approximation for the weighted problem and a deterministic $(1 + 1/(e-1))$ -approximation in the unweighted case. For bounded-complexity subclasses, local search yields a PTAS for the Hitting Set problem, and we also extend our results for segments in d orientations.

Two natural directions remain open. First, it would be interesting to determine whether the analysis can be strengthened further, or whether these bounds are close to the true integrality behavior of the natural LP. For the Set Cover and Independent Set problems, the best known approximation is the straightforward factor-2 bound. The integrality gap of the natural LP-formulation for Independent Set was shown to be 2 [9]. Hence, breaking the factor-2 for the independent set problem cannot go via rounding the natural LP and new ideas are required. Finally, the approach suggests exploring whether similar one-orientation-first rounding schemes can be extended to richer families such as thin rectangles.

References

- 1 Anna Adamaszek, Sarel Har-Peled, and Andreas Wiese. Approximation schemes for independent set and sparse subsets of polygons. *Journal of the ACM*, 66(4):29:1–29:40, 2019. doi:10.1145/3326122.
- 2 Noga Alon. A non-linear lower bound for planar epsilon-nets. *Discret. Comput. Geom.*, 47(2):235–244, 2012. doi:10.1007/S00454-010-9323-7.
- 3 Boris Aronov, Esther Ezra, and Micha Sharir. Small-size ϵ -nets for axis-parallel rectangles and boxes. *SIAM J. Comput.*, 39(7):3248–3282, 2010. doi:10.1137/090762968.
- 4 Rom Aschner, Matthew J. Katz, Gila Morgenstern, and Yelena Yuditsky. Approximation schemes for covering and packing. In Subir Kumar Ghosh and Takeshi Tokuyama, editors, *WALCOM: Algorithms and Computation, 7th International Workshop, WALCOM 2013, Kharagpur, India, February 14-16, 2013. Proceedings*, volume 7748 of *Lecture Notes in Computer Science*, pages 89–100. Springer, 2013. doi:10.1007/978-3-642-36065-7_10.
- 5 Jozsef Balogh and Jozsef Solymosi. On the number of points in general position in the plane. *Discrete Analysis*, pages 1–20, 2018. doi:10.19086/da.4438.

- 6 Aritra Banik, Rajiv Raman, and Saurabh Ray. Geometric optimization parameterized by piercing complexity. In *EATCS International Colloquium on Automata, Languages, and Programming (ICALP) 2026*, 2026. to appear.
- 7 Aniket Basu Roy, Sathish Govindarajan, Rajiv Raman, and Saurabh Ray. Packing and covering with non-piercing regions. *Discrete & Computational Geometry*, 2018.
- 8 Hervé Brönnimann and Michael T. Goodrich. Almost optimal set covers in finite VC-dimension. *Discrete & Computational Geometry*, 14(4):463–479, 1995. doi:10.1007/BF02570718.
- 9 Marco Caoduro, Jana Cslovjecssek, Michal Pilipczuk, and Karol Węgrzycki. On the independence number of intersection graphs of axis-parallel segments. *J. Comput. Geom.*, 14(1):144–156, 2023. doi:10.20382/JOCG.V14I1A5.
- 10 Timothy M. Chan and Elyot Grant. Exact algorithms and APX-hardness results for geometric packing and covering problems. *Computational Geometry*, 47(2):112–124, 2014. doi:10.1016/j.comgeo.2012.04.001.
- 11 Timothy M. Chan and Elyot Grant. Exact algorithms and APX-hardness results for geometric packing and covering problems. *Computational Geometry Theory and Applications*, 47(2):112–124, 2014. doi:10.1016/J.COMGEO.2012.04.001.
- 12 Timothy M. Chan, Elyot Grant, Jochen Könemann, and Malcolm Sharpe. Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms*, Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 1576–1585, 2012. doi:10.1137/1.9781611973099.125.
- 13 Timothy M. Chan and Sarel Har-Peled. Approximation algorithms for maximum independent set of pseudo-disks. *Discrete and Computational Geometry*, 48(2):373–392, 2012. doi:10.1007/S00454-012-9417-5.
- 14 Jana Cslovjecssek, Michal Pilipczuk, and Karol Węgrzycki. Parameterized approximation for maximum weight independent set of rectangles and segments. In Timothy M. Chan, Johannes Fischer, John Iacono, and Grzegorz Herman, editors, *32nd Annual European Symposium on Algorithms, ESA 2024, Royal Holloway, London, United Kingdom, September 2-4, 2024*, LIPIcs, pages 43:1–43:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.ESA.2024.43.
- 15 Rong-chii Duh and Martin Fürer. Approximation of k -set cover by semi-local optimization. In Frank Thomson Leighton and Peter W. Shor, editors, *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 256–264. ACM, 1997. doi:10.1145/258533.258599.
- 16 Guy Even, Dror Rawitz, and Shimon Shahar. Hitting sets when the vc-dimension is small. *Inf. Process. Lett.*, 95(2):358–362, 2005. doi:10.1016/j.ipl.2005.03.010.
- 17 Sándor P. Fekete, Kan Huang, Joseph S. Mitchell, Ojas Parekh, and Cynthia A. Phillips. Geometric hitting set for segments of few orientations. *Theor. Comp. Sys.*, 62(2):268–303, February 2018. doi:10.1007/s00224-016-9744-7.
- 18 Jacob Fox and János Pach. Applications of a new separator theorem for string graphs. *Comb. Probab. Comput.*, 23(1):66–74, 2014. doi:10.1017/S0963548313000412.
- 19 Sarel Har-Peled and Kent Quanrud. Approximation algorithms for polynomial-expansion and low-density graphs. In *Algorithms – ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14–16, 2015, Proceedings*, pages 717–728, 2015. doi:10.1007/978-3-662-48350-3_60.
- 20 David Haussler and Emo Welzl. epsilon-nets and simplex range queries. *Discrete & Computational Geometry*, 2:127–151, 1987. doi:10.1007/BF02187876.
- 21 J Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182(1):105–142, 1999.
- 22 Katarzyna Kowalska and Michał Pilipczuk. Parameterized and approximation algorithms for coverings points with segments in the plane. In *41st International Symposium on Theoretical Aspects of Computer Science*, 2024.

- 23 Richard J. Lipton and Robert E. Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979.
- 24 Dániel Marx. Parameterized complexity of independence and domination on geometric graphs. In *2nd 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.
- 25 Jirí Matousek, Raimund Seidel, and Emo Welzl. How to net a lot with little: Small epsilon-nets for disks and halfspaces. In Raimund Seidel, editor, *Proceedings of the Sixth Annual Symposium on Computational Geometry, Berkeley, CA, USA, June 6-8, 1990*, pages 16–22. ACM, 1990. doi:10.1145/98524.98530.
- 26 Nabil H. Mustafa. *Sampling in combinatorial and geometric set systems*, volume 265. American Mathematical Society, 2022.
- 27 Nabil H. Mustafa, Rajiv Raman, and Saurabh Ray. Quasi-polynomial time approximation scheme for weighted geometric set cover on pseudodisks and halfspaces. *SIAM Journal on Computing*, 44(6):1650–1669, 2015. doi:10.1137/14099317X.
- 28 Nabil H. Mustafa and Saurabh Ray. Improved results on geometric hitting set problems. *Discrete & Computational Geometry*, 44(4):883–895, 2010. doi:10.1007/S00454-010-9285-9.
- 29 János Pach and Gábor Tardos. Tight lower bounds for the size of epsilon-nets. In *Proceedings of the twenty-seventh annual symposium on Computational geometry*, pages 458–463, 2011. doi:10.1145/1998196.1998271.
- 30 János Pach and Géza Tóth. Graphs drawn with few crossings per edge. *Comb.*, 17(3):427–439, 1997. doi:10.1007/BF01215922.
- 31 Evangelia Pyrga and Saurabh Ray. New existence proofs for ϵ -nets. In *Proceedings of the Twenty-fourth Annual Symposium on Computational Geometry*, SoCG '08, pages 199–207, New York, NY, USA, 2008. ACM.
- 32 Rajiv Raman and Saurabh Ray. Constructing planar support for non-piercing regions. *Discrete and Computational Geometry*, 64(3):1098–1122, 2020. doi:10.1007/S00454-020-00216-W.
- 33 Rajiv Raman, Siddhartha Sarkar, and Jatin Yadav. Hitting axis-parallel segments with weighted points, 2026. doi:10.48550/arXiv.2605.14499.
- 34 Alexander Schrijver. *Theory of linear and integer programming*. John Wiley & Sons, 1998.
- 35 Moshe Shaked and J. George Shanthikumar. *Stochastic Orders*. Springer Series in Statistics. Springer New York, New York, NY, 2007. doi:10.1007/978-0-387-34675-5.
- 36 Kasturi Varadarajan. Weighted geometric set cover via quasi-uniform sampling. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, pages 641–648, 2010. doi:10.1145/1806689.1806777.