

Geometric Hitting Set for Segments of Few Orientations

Sándor P. Fekete · Kan Huang ·
Joseph S. B. Mitchell · Ojas Parekh ·
Cynthia A. Phillips

Received: date / Accepted: date

Abstract We study several natural instances of the geometric hitting set problem for input consisting of sets of line segments (and rays, lines) having a small number of distinct slopes. These problems model path monitoring (e.g., on road networks) using the fewest sensors (the “hitting points”). We give approximation algorithms for cases including (i) lines of 3 slopes in the plane, (ii) vertical lines and horizontal segments, (iii) pairs of horizontal/vertical segments. We give hardness and hardness of approximation results for these problems. We prove that the hitting set problem for vertical lines and horizontal rays is polynomially solvable.

Keywords set cover · hitting set · approximation algorithms

1 Introduction

The set cover problem is fundamental in combinatorial optimization. It is NP-hard and has an $O(\log n)$ -approximation algorithm, which is best possible

S. Fekete
TU Braunschweig, Braunschweig, Germany
E-mail: s.fekete@tu-bs.de

K. Huang
Stony Brook University, Stony Brook, NY, USA
E-mail: khuang@ams.stonybrook.edu

J. Mitchell
Stony Brook University, Stony Brook, NY, USA
E-mail: jsbm@ams.stonybrook.edu

O. Parekh
Sandia National Labs, Albuquerque, NM, USA
E-mail: odparek@sandia.gov

C. Phillips
Sandia National Labs, Albuquerque, NM, USA
E-mail: caphill@sandia.gov

(unless $P = NP$, [15]). Equivalently, set cover can be cast as a hitting set problem: given a collection, \mathcal{C} , of subsets of set U , find a smallest cardinality set $H \subseteq U$ such that every set in \mathcal{C} contains at least one element of H . Numerous special instances of set cover/hitting set have been studied. Our focus in this paper is on geometric instances that arise in covering (hitting) sets of (possibly overlapping) line segments using the fewest points (“hit points”). A closely related problem is the “Guarding a Set of Segments” (GSS) problem [5, 7, 8, 28], in which the segments may cross arbitrarily, but do not overlap. Since this problem is strongly NP-complete [7] in general, our focus is on special cases, primarily those in which the segments come from a small number of orientations (e.g., horizontal, vertical). We provide several new results on hardness and approximation algorithms.

We are motivated by the path monitoring problem: given a set of trajectories, each a path of line segments in the plane, place the fewest sensors (points) to observe (hit) all trajectories. To gain theoretical insight into this challenging problem, we examine cleaner, but progressively harder, versions of hitting trajectory/line-like objects with points. If the trajectories are on a Manhattan road network, the paths are (possibly overlapping) horizontal/vertical segments. Alternatively, one wishes to place the fewest vendors or service stations in a road network to service a set of customer trajectories.

Our results. We give complexity and approximation results for several geometric hitting set problems on inputs S of line “segments” of special classes, mostly of fixed orientations. The segments are allowed to overlap arbitrarily. We consider various cases of “segments” that may be bounded (line segments), semi-infinite (rays), or unbounded in both directions (lines). Our results are:

(1) Hitting lines of 3 slopes in the plane is NP-hard (greedy is optimal for 2 slopes). For set cover with set size at most 3, standard analysis of the greedy algorithm gives an approximation factor of $H(3) = 1 + (1/2) + (1/3) = (11/6)$, and there is a $4/3$ -approximation based on semi-local optimization [17]. We prove that the greedy algorithm in this special geometric case is a $(7/5)$ -approximation.

(2) Hitting vertical lines and horizontal rays is polytime solvable.

(3) Hitting vertical lines and horizontal (even unit-length) segments is NP-hard. Our proof shows hitting horizontal and vertical unit-length segments is also NP-hard. We prove APX-hardness for hitting horizontal and vertical segments.

(4) Hitting vertical lines and horizontal segments has a $(5/3)$ -approximation algorithm. (This problem has a straightforward 2-approximation.)

(5) Hitting pairs of horizontal/vertical segments has a 4-approximation. Hitting pairs having one vertical and one horizontal segment has a $(10/3)$ -approximation. These results are based on LP-rounding. More generally, hitting sets of k segments from r orientations has a $(k \cdot r)$ -approximation algorithm.

(6) We give a linear-time combinatorial 3-approximation algorithm for hitting triangle-free sets of (non-overlapping) segments. Recently Joshi and

Narayanaswamy [28] gave a 3-approximation for this version of GSS using linear programming.

Related Work. There is a wealth of related work on geometric set cover and hitting set problems; we do not attempt here to give an exhaustive survey. The *point line cover* (PLC) problem (see [26,30]) asks for a smallest set of lines to cover a given set of points; it is equivalent, via point-line duality, to the hitting problem for a set of lines. The PLC (and thus the hitting problem for lines) was shown to be NP-hard [33]; in fact, it is APX-hard [9] and Max-SNP Hard [31]. The problem has an $O(\log OPT)$ -approximation (e.g., greedy – see [29]); in fact, the greedy algorithm for PLC has worst-case performance ratio $\Omega(\log n)$ [18].

Hassin and Megiddo [25] considered hitting geometric objects with the fewest lines having a small number of distinct slopes. They observed that, even for covering with axis-parallel lines, the greedy algorithm has an approximation ratio that grows logarithmically. They gave approximations for the problem of hitting horizontal/vertical segments with the fewest axis-parallel lines (and, more generally, with lines of a few slopes). Gaur and Bhattacharya [22] consider covering points with axis-parallel lines in d -dimensions. They give a $(d - 1)$ -approximation based on rounding the corresponding linear program (LP). Many other stabbing problems (find a small set of lines that stab a given set of objects) have been studied; see, e.g., [16,20,23,24,29,32].

A recent paper [28] gives a 3-approximation for hitting sets of “triangle-free” segments. Brimkov et al. [5,7,8] have studied the hitting set problem on line segments, including various special cases; they refer to the problem as “Guarding a Set of Segments”, or GSS. GSS is a special case of the “art gallery problem:” place a small number of “guards” (e.g., points) so that every point within a geometric domain is “seen” by at least one guard [35,37]. Brimkov et al. [6] provide experimental results for three GSS heuristics, including two variants of “greedy,” showing that in practice the algorithms perform well and are often optimal or very close to optimal. They prove, however, that, in theory, the methods do not provide worst-case constant-factor approximation bounds. For the special case that the segments are “almost tree (1)” (a connected graph is an *almost tree* (k) if each biconnected component has at most k edges not in a spanning tree of the component), a $(2 - \varepsilon)$ -approximation is known [5].

An important distinction between GSS and our problems is that we allow *overlapping* (or partially overlapping) segments (rays, and lines), while, in GSS, each line segment is maximal in the input set of line segments (the union of two distinct input segments is not a segment). A special case of our problem is *interval stabbing* on a line: Given a set of segments (intervals), arbitrarily overlapping on a line, find a smallest hitting set of points that hit all segments. A simple sweep along the line solves this problem optimally: when a segment ends, place a point and remove all segments covered by that point.

If no point lies within three or more objects, then the hitting set problem is an edge cover problem in the intersection graph of the objects. In particular, if no three segments pass through a common point, the problem can be

solved optimally in polynomial time. (This implies that in an arrangement of “random” segments, the GSS problem is almost surely polynomially solvable; see [5].)

Hitting axis-aligned rectangles is related to hitting horizontal and vertical segments. Aronov, Ezra, and Sharir [2] provide an $O(\log \log OPT)$ -approximation for hitting set for axis-aligned rectangles (and axis-aligned boxes in 3D), by proving a bound of $O(\varepsilon^{-1} \log \log(\varepsilon^{-1}))$ on the ε -net size of the corresponding range space. The connection between hitting sets and ε -nets [10, 13, 14, 21] implies a c -approximation for hitting set if one can compute an ε -net of size c/ε ; recent major advances [1, 36] on lower bounds on ε -nets imply that associated range spaces (rectangles and points, lines and points, points and rectangles) have ε -nets of size superlinear in $1/\varepsilon$. Remarkably, improved $(1+\varepsilon)$ -approximation algorithms (i.e., PTASs) for certain geometric hitting set and set cover problems are possible with simple local search. For example, Mustafa and Ray [34] give a local search PTAS for computing a smallest subset, of a given set of disks, that covers a given set of points. Hochbaum and Maas [27] used grid shifting to obtain a much earlier PTAS for the minimum unit disk cover problem when disks can be placed anywhere in the plane, not restricted to a discrete input set.

2 Hitting Segments

Suppose S is a set of line segments in the plane. If all segments are horizontal, then we can compute an optimal hitting set by independently solving the interval stabbing problem along each of the horizontal lines determined by the input.

If the segments are of two different orientations (slopes), then the problem becomes significantly harder. Without loss of generality, assume the segments are horizontal and vertical. We show the problem is hard even if the axis-parallel segments are all the same length. This result (Corollary 1) is a consequence of an even stronger result, Theorem 5, which we establish in Section 5.

We readily obtain a 2-approximation algorithm by solving optimally each of the two orientations, and using the union of the hitting points for both. (This generalizes to a k -approximation for hitting sets of segments of k orientations.)

3 Hitting Lines

When S is a set of n lines in the plane, greedy gives an $O(\log OPT)$ approximation factor; any approximation factor better than logarithmic would be quite interesting. (See [19, 30].) If the lines have only 2 slopes, then greedy is optimal.

3.1 Hardness of Hitting Lines of 3 Slopes in 2D

We prove that the hitting set problem is NP-hard when lines have more than two orientations. Consider the dual formulation: (3-SLOPE-LINE-COVER, 3SLC) Find a minimum-cardinality set of non-vertical lines to cover a set P of points (duals to the set S of lines), which are known to lie on three vertical lines.

If a line ℓ covers i points, we say that ℓ is an i -line. Let P_ℓ denote the set of points of P that are covered by line ℓ . If $P_{\ell_1} \cap P_{\ell_2} = \emptyset$, then we say that lines ℓ_1 and ℓ_2 are *independent*. A set L of lines is *independent* if the lines are pairwise independent; i.e., $\ell_i \in L$ and $\ell_j \in L$ are independent for every $\ell_i \neq \ell_j$.

Theorem 1 *The problem 3SLC is NP-complete.*

Proof Our reduction is from 3-SAT. Let n and m denote the numbers of variables and clauses respectively. The input points P are distributed on three vertical lines, denoted l_1 , l_2 and l_3 from left to right. Figure 1 shows a variable gadget for variable u in a 4-clause instance. Points in the first two columns can be hit by either of two sets of independent lines, red or blue. These represent the “True” or “False” setting of the variable u respectively. Each variable gadget has m points on line l_1 , m points on line l_2 and $2m$ points on line l_3 . We can add the variable gadgets one by one onto the three columns so that all 3-lines are within variable gadgets.

We use a clause gadget to link variables. The i th clause has a single point $c[i]$ on line l_1 . Point $c[i]$ is connected to the negations of the variables in its clause on line l_3 . Figure 2 shows the gadget for clause $c[i] = \bar{u} \vee v \vee \bar{w}$. We add the intersections of these three line segments with line l_2 to the point set. We pick the location of $c[i]$ such that the addition of these four new points create no new 3-lines other than the green lines associated with the clause gadget.

To complete the construction, we place $nm + m$ points on line l_1 such that these new points are not on any 3-line. Thus, by construction, each 3-line is in a variable gadget or a clause gadget. There are $2nm + 2m$ points on line l_1 , $nm + 3m$ points on line l_2 and $2nm$ points on line l_3 for a total of $5nm + 5m$ points on all three lines.

We now argue that the 3-SAT formula is satisfiable if and only if the corresponding 3SLC instance can be covered by $2nm + 2m$ non-vertical lines.

If the 3-SAT formula is satisfiable, there is an independent set of 3-lines in the variable and clause gadgets of size $nm + m$. Specifically, there are nm independent 3-lines from the satisfying truth assignment for the n variables. Since the formula is satisfiable, for each clause, the 3-line corresponding to the negation of one correctly set variable is independent of the variable truth assignment points, and can be part of an independent set. This gives the other m members of the independent set. There are then $nm + m$ points on l_1 , $2m$ points on l_2 and $nm - m$ points on l_3 not part of this maximum independent set. Each of the $nm + m$ remaining points on line l_1 can be paired with one of the remaining $nm + m$ points on either line l_2 or l_3 . Thus, all points are

covered with $(nm + m)$ 3-lines and $(nm + m)$ 2-lines for a total of $2nm + 2m$ lines.

If the instance can be covered with $2nm + 2m$ lines, then, by construction, $nm + m$ of these lines must be 2-lines involving the points on line l_1 that are not in any clause or variable gadgets. These 2-lines can cover $2nm + 2m$ points, leaving $3nm + 3m$ total points to be covered by the remaining $nm + m$ lines. Thus the remaining lines must all be independent 3-lines. nm of these correspond to a truth assignment for the variables. The remaining m must come from clause gadgets. Thus each clause gadget has a 3-line compatible with the truth assignment and the 3-SAT instance is satisfiable. \square

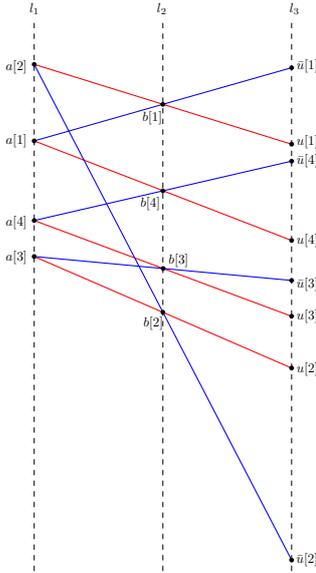


Fig. 1 A variable gadget.

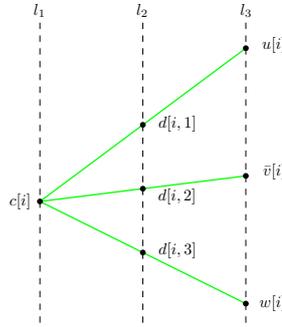


Fig. 2 A clause gadget.

3.2 Analysis of the Greedy Hitting Set Algorithm for Lines of 3 Slopes in 2D

If no point lies in more than k sets, the greedy algorithm's approximation factor is $H(k) = \sum_{i=1}^k (1/i)$ [12]. This property holds for lines of 3 slopes with $k = 3$, giving a greedy approximation factor $H(3) = 11/6$. We give a new analysis, exploiting the special geometric structure of the hitting set problem for lines of 3 slopes, to obtain an approximation factor of $7/5$.

Let x , y and z be the number of lines in each of the three orientations in the plane. Without loss of generality, we assume that $x \geq y \geq z > 0$.

We call a point where at least two lines meet a vertex. A vertex is a 3-intersection if three lines, one from each orientation, meet at that point.

Otherwise it is a 2-intersection. If there is no 3-intersections, then all vertices are 2-intersections, and the greedy algorithm is optimal.

We analyze this case in more detail because it will be useful later in the proof. The greedy algorithm places points at any 2-intersection, removes the two lines, and continues until there are no more 2-intersections. Then it places a hit point on each remaining unhit line (all of which are parallel to each other). Let $OPT_2(x, y, z)$ denote the optimal solution to a 3-slope instance with no 3-intersections. Then we have

$$OPT_2(x, y, z) = \begin{cases} x & \text{if } x \geq y + z; \\ x + \lceil \frac{y+z-x}{2} \rceil & \text{if } x < y + z. \end{cases}$$

When there are 3-intersections, the greedy algorithm iteratively places hit points at 3-intersections (in any order), removing the 3 hits lines, until there are no more 3-intersections. The remaining instance, with only 2-intersections, now can now be solved optimally as described above.

Theorem 2 *The greedy algorithm yields a $\frac{7}{5}$ -approximation for hitting lines of 3 slopes in 2D.*

Proof Consider the graph G whose vertices are the 3-intersections of the input set of lines, with an edge of G between two vertices if and only if the corresponding 3-intersections lie on a common line. Let K denote the cardinality of a maximum independent set, I_{max} , in G . Let N_3 be the number of 3-intersections that the greedy algorithm selects. These N_3 vertices correspond to a maximal independent set in G , so $K \geq N_3$. We have, $K \leq 3N_3$, since (by independence) there is at most one vertex of I_{max} along each of the $3N_3$ lines that the greedy algorithm hits with 3-intersections.

The optimal solution has OPT hit points, with

$$OPT = K + OPT_2(x - K, y - K, z - K) \quad (1)$$

The greedy solution yields a set of N_{greedy} hit points, with

$$N_{greedy} = N_3 + OPT_2(x - N_3, y - N_3, z - N_3). \quad (2)$$

First suppose $x = K$, which means that $x = y = z = K$. Then, $OPT = K$ and $N_{greedy} = N_3 + OPT_2(K - N_3, K - N_3, K - N_3) = K + \lceil \frac{K-N_3}{2} \rceil$. We have three cases, depending on the value of K , mod 3:

1. $K = 3l$, where l is an integer; thus,

$$\frac{N_{greedy}}{OPT} \leq 1 + \left\lceil \frac{3l-l}{2} \right\rceil / 3l = \frac{4}{3}. \quad (3)$$

2. $K = 3l + 1$; thus,

$$\frac{N_{greedy}}{OPT} \leq 1 + \left\lceil \frac{3l+1-(l+1)}{2} \right\rceil / (3l+1) < \frac{4}{3}. \quad (4)$$

3. $K = 3l + 2$; thus,

$$\frac{N_{greedy}}{OPT} \leq 1 + \left\lceil \frac{3l + 2 - (l + 1)}{2} \right\rceil / (3l + 2) \leq \frac{7}{5}. \quad (5)$$

Now, suppose $x \geq K + 1$. Then, we have the following three cases:

1. If $x \geq y + z - N_3$, then $x - N_3 \geq y - N_3 + z - N_3$, and $x - K \geq y - K + z - K$. Thus,

$$\begin{aligned} OPT &= K + x - K = x, \\ N_{greedy} &= N_3 + x - N_3 = x. \end{aligned}$$

2. If $y + z - K \leq x < y + z - N_3$, then

$$\begin{aligned} OPT &= x, \\ N_{greedy} &= x + \left\lceil \frac{y + z - x - N_3}{2} \right\rceil, \\ \frac{N_{greedy}}{OPT} &= 1 + \left\lceil \frac{y + z - x - N_3}{2} \right\rceil / x \leq 1 + \left\lceil \frac{K - N_3}{2} \right\rceil / (K + 1) \leq \frac{4}{3}. \end{aligned} \quad (6)$$

The detailed analysis is similar to the case in which $x = K$.

3. If $x < y + z - K$, then

$$\begin{aligned} OPT &= x + \left\lceil \frac{y + z - x - K}{2} \right\rceil \geq K + 1 + 1 = K + 2, \\ N_{greedy} - OPT &\leq \frac{y + z - x - N_3}{2} + 1 - \frac{y + z - x - K}{2} \leq 1 + \frac{K}{3}, \\ \frac{N_{greedy}}{OPT} &= 1 + \frac{N_{greedy} - OPT}{OPT} \leq 1 + \frac{1 + K/3}{K + 2}. \end{aligned} \quad (7)$$

Thus,

- when $K = 1$, $OPT = N_{greedy}$;
- when $K = 2, 3, 4$, using (6) and the first equality in (7), we have $N_{greedy} - OPT = 1$, so we have

$$\frac{N_{greedy}}{OPT} \leq 1 + \frac{1}{4} = 1.25; \quad (8)$$

- when $K \geq 5$,

$$\frac{N_{greedy}}{OPT} \leq 1 + \frac{8}{21} \approx 1.381.$$

This completes the proof. \square

3.3 Axis-Parallel Lines in 3D

While in 2D the hitting set problem for axis-parallel lines is easily solved, in 3D we prove that the corresponding hitting set problem is NP-hard, using a reduction from 3-SAT.

Theorem 3 *Hitting set for axis-parallel lines in 3D is NP-complete.*

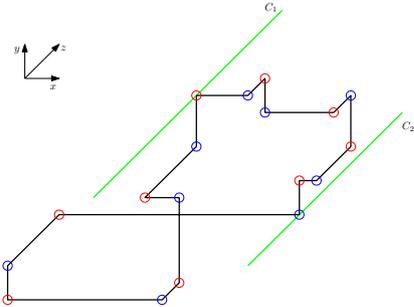


Fig. 3 An example of a variable loop.

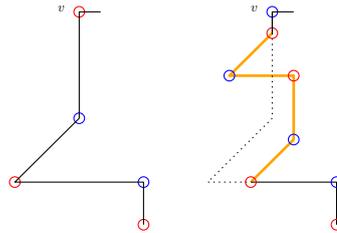


Fig. 4 The insertion of the orange detour changes the color of v from red to blue.

Proof We give a reduction from 3-SAT. We say that a line is a d -line if it is parallel to the d -axis; we say that a plane is an ab -plane if it is parallel to the plane spanned by the a -axis and the b -axis. A clause is represented by a z -line. A variable is represented by a loop of axis-parallel lines with the following properties:

1. No two x -lines lie in the same xz -plane. (This ensures that when a clause z -line meets a vertex of the loop where an x -line meets a y -line, it does not also meet another such vertex.)
2. There are an even number of edges in each loop.
3. Lines from two different loops do not intersect.
4. A loop intersects a clause z -line if and only if the variable represented by that loop is in the clause represented by the z -line. The intersection point represents a literal in the corresponding clause.
5. There are two optimal hitting sets for a loop – the set of odd vertices, and the set of even vertices. All positive (resp., negative) literals should be in the same hitting set.

Figure 3 shows a portion of an instance in which clause C_1 includes literal x_1 and clause C_2 includes literal \bar{x}_1 .

In three dimensions, it is not hard to take detours to avoid unwanted intersections. The number of vertices in a loop can be adjusted by inserting a detour, as shown in Figure 4. Finally, we argue that all clause z -lines can be hit for free by the optimal hitting sets of variable loops if and only if there is a satisfying truth assignment for the corresponding 3-SAT instance.

□

4 Hitting Rays and Lines

Hitting rays is “harder” than hitting lines, since any instance of hitting lines has a corresponding equivalent instance as a hitting rays problem (place the apices of the rays far enough away that they are effectively lines). Unlike lines, there can be many different collinear rays. Divide collinear rays into two groups according to the direction they point along the containing line, ℓ . Because of nesting, we need keep only one of the rays pointing in each of the two directions along ℓ . For example, among left-pointing rays, we keep only the one contained in all other left-pointing rays. This is the one with the right-most apex.

We show that the special case with horizontal rays and vertical lines (abbreviated HRVL) is exactly solvable in polynomial time:

Theorem 4 *The hitting set problem for vertical lines and horizontal rays can be solved in $O(nT)$ time, where n is the number of entities and T is the time for computing a maximum matching in a bipartite graph with n nodes.*

We begin with a high-level overview of the algorithm. A point can cover at most 3 objects: a vertical line, a left-facing ray, and a right-facing ray. This requires the two rays to intersect in a segment, and the vertical line to intersect this segment. We call points at such intersections *3-hitters*. We can compute the maximum possible number of 3-hitters, with no two sharing a line or a segment, via maximum matching in a bipartite graph, where edges represent intersections between vertical lines and horizontal segments. We prove there exists an optimal solution with this maximum number of 3-hitters. The algorithm performs a sweep inward from the left and right, finding a suitable set of 3-hitters, ensuring the remaining lines have the best possible chance to share a point with the remaining rays. Once everything that is 3-hit is removed, the remaining objects intersect in at most pairs. So we can finish the hitting by solving an edge cover problem between rays and lines. We prove this is optimal.

We now give additional algorithmic and proof details. We call a horizontal ray to the left (resp., right) an *l-ray* (resp. *r-ray*). In this section, all lines are vertical. If two collinear rays are disjoint, we shift one ray slightly up or down, so no two disjoint rays are collinear. These rays cannot be covered by a single point, so this does not fundamentally alter the optimal solution. We also assume that no collinear rays have the same orientation, since in this case one ray is a subset of the others, which are redundant.

If a ray is not collinear with any other ray, we add a ray to pair with it. For example, if an r-ray intersects no l-ray, we add an intersecting l-ray whose right endpoint is to the right of all vertical input lines. This additional ray will not change the optimal solution. If an l-ray and r-ray intersect, their intersection is a segment. Since each ray intersects exactly one other ray, we represent each such pair of rays by their segment.

Let H and V denote the number of segments and lines respectively. Any solution requires V points to cover the lines. Those points can help “hit” segments in two possible ways:

- (1) The point on a line may be placed on a segment. We call the corresponding line a *3-hitter* and say that the segment is *3-hit* by the line.
- (2) Points on lines may hit each ray outside its intersecting segment. This requires two points on two distinct lines. We call the left (resp., right) line an *l-hitter* (resp., *r-hitter*). We say the segment is *double-hit* by those two lines.

Let v_1 and v_2 be the number of segments hit by the V points in the first (1) and second (2) types, respectively. Then, the number of points in the solution is $H + V - v_1 - v_2$. Our goal is to put a point on each line to hit as many segments as possible (i.e., to maximize $v_1 + v_2$).

Given an instance of HRVL, we can calculate the maximum number of 3-hitters. We construct the bipartite graph G in which one set of nodes is the lines and the other set of nodes is the segments; there is an edge between two nodes if and only if the line and the segment they represent intersect. We refer to G as the *lines-segments graph*. Maximum matching in a bipartite graph is solvable in polynomial time. A matching in the graph represents a set of independent intersections in the corresponding HRVL. That is, a set of M edges in a matching corresponds to a way to hit M segments and M lines with M points. These are hittings of type (1). The following lemma shows that hitting points of type (1) are preferred over hitting points of type (2).

Lemma 1 *For any instance of HRVL, there is a maximum matching between lines and segments that can be augmented to be an optimal solution.*

Proof The proof is by contradiction. Let v_1^* be the largest v_1 for any minimum hitting set. We assume that v_1^* is less than m , the cardinality of the maximum matching between lines and segments. Thus, there is an augmenting path in the bipartite graph G ; an example of such a path is shown in green in Figure 5. Because the current solution is optimal, any augmenting path cannot improve it. This allows us to infer some properties of the first segment and the last line on the augmenting path. We consider the augmenting path P with the shortest length (fewest elements) and the shortest horizontal distance between the last two lines along the path. Then by case analysis on path P , we argue there exists another augmenting path that increases v_1^* or violates a minimality condition of P .

In more detail, an augmenting path in graph G corresponds to a sequence of alternating segments and lines in the HRVL instance: $\{e_1, l_1, e_2, l_2, \dots, e_n, l_n\}$. In the current solution line l_{i-1} is matched with segment e_i . In Figure 5 the green path is an example of an augmenting path, where $n = 3$ and blue dots correspond to current 3-hitter points. Because of the optimality of the current solution, any augmenting path cannot improve it. The following two properties hold; otherwise, after augmenting, the sum of v_1 and v_2 will stay the same, but v_1^* would be increased by 1:

- e_1 is double-hit by other lines.
- l_n is helping to double-hit another segment.

Without loss of generality, we assume the intersection of e_1 and l_1 is to the left of l_n . Also assume that n , the number of lines in the augmenting path, is as small as possible. Among augmenting paths with smallest n , we pick one with the shortest horizontal distance between lines l_{n-1} and l_n . We consider the following cases:

1. If line l_n is the l-hitter of some segment e_t , then the l-hitter of segment e_1 can take its job. One can do the augmenting and assign the l-hitter of e_1 to l-hit e_t . Therefore the solution is still optimal and v_1^* increases.
2. Suppose line l_n is an r-hitter for a segment e_t , and let l_r be the r-hitter for segment e_1 . If line l_r is to the right of line l_n , then line l_r can take the job of line l_n . That is, line l_r can be the r-hitter for segment e_t . Again, it is now possible to increase v_1^* while maintaining an optimal total number of points.
3. If line l_n is an r-hitter and the r-hitter of e_1 (called l_r) is to the left of l_n , we know that line l_r will intersect a segment, e_q (with $1 \leq q \leq n$) of the augmenting path. This is because we assume line l_n is to the right of the intersection of e_1 and l_1 . In Figure 5, line l_4 is a possible l_r . Thus, using l_r gives a shorter augmenting path, $\{e_1, l_1, \dots, e_q, l_r\}$. This new path either has strictly fewer lines or has a strictly shorter horizontal distance between the last lines. This contradicts the choice of the first augmenting path. \square

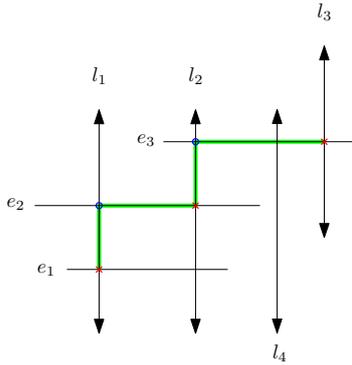


Fig. 5 A green augmenting path: the matching size increases by replacing blue circles with red crosses.

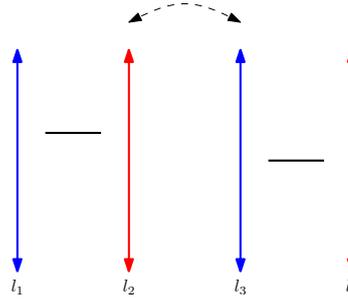


Fig. 6 Swapping the (l-hitter or r-hitter) roles of l_2 and l_3 makes both of them more useful.

The following lemma gives additional useful structure for at least one optimal solution:

Lemma 2 *Given an optimal solution \mathcal{S} , there is an optimal solution \mathcal{S}' that has the same set of 3-hitters as \mathcal{S} , with its l-hitters all left of its r-hitters.*

Proof In Figure 6 two segments are double-hit by two pairs of lines; the blue lines are l-hitters and the red lines are r-hitters. When we pair l_1 to l_3 and

pair l_2 to l_4 , the two segments are still double-hit, because this swap moves the l-hitter further left and the r-hitter further right. A sequence of such swaps moves all l-hitters to the left of all r-hitters. \square

Algorithm 1 below gives an optimal solution for HRVL. The algorithm maximizes the number of 3-intersections and “balances” the remaining lines between the left and right sides as much as possible. We test the “criticality” of a line l by computing a maximum cardinality matching in the lines-segments graph, with and without the line l ; if the matching cardinality drops when line l is not part of the graph, then line l is *critical*. In the algorithm, we check the criticality of lines: given the previous choices, if a critical line is not used as a 3-hitter, there is no way to extend the previous choices to a maximum matching. We add a 3-hitter to our solution if and only if the line involved is critical. Non-critical lines are counted as l-hitters when sweeping from the left and as r-hitters when sweeping from the right. The algorithm removes each newly-discovered non-critical line from consideration as a 3-hitter, and swaps the sweep direction (from left to right, or from right to left). This balances the number of presumed l-hitters and r-hitters during the course of the the algorithm. When the sweep encounters a segment s , this triggers testing of the first line (in the sweep direction) that intersects s , if any. Subsequent sweep steps continue to process segment s until it is matched as part of a 3-hitter, or it is removed, to be double hit at the end of the algorithm.

We now argue that the left-right-balanced approach gives lines that obey Lemma 2. Let S be the solution given by Algorithm 1, and let S' be an optimal solution with the maximum cardinality set of 3-hitters. We know that S and S' have the same number of 3-hitters. Let D and D' denote the sets of lines left behind (not 3-hitters) in S and S' , respectively. We order lines in D and D' from left to right. Let k be $\lfloor \frac{|D|}{2} \rfloor$. Thus, there are at most k pairs of double-hitters in S and S' . Let lh_i (resp., lh'_i) be the i th line of D (resp., D').

Given a solution P and a line l , let $E(l, P)$ denote the number of segments on the left side of l not hit by 3-hitters in P . A line having more segments on its right side is more likely to be an l-hitter. We will show that line lh_i is at least as capable of being an l-hitter as is line lh'_i ; specifically, we show that

$$E(lh_i, S) \leq E(lh'_i, S'), \quad i = 1, 2, \dots, k. \quad (9)$$

We split the proof of (9) into one claim and two lemmas.

Because of the criticality test and the choice of intersecting segment e_2 , we have the following claim:

Claim In S if a 3-hitter is on the left side of an l-hitter, the segment hit by the 3-hitter will not intersect that l-hitter.

Proof The proof is by contradiction. Let l be a 3-hitter, matched to segment e . Let l_h be the closest l-hitter to the right of line l and suppose l_h intersects segment e . There can be $j \geq 0$ lines between l and l_h which must all be 3-hitters. Let this set of 3-hitters with their matched segments be

```

1 Input: set  $L$  of vertical lines, set  $S$  of horizontal segments (ray intersections);
2  $H \leftarrow [0, 0]$  //  $H$  counts 2-hitters at left and right sides;
3  $I_3 \leftarrow \{\}$  //  $I_3$  stores 3-intersections of the solution;
4  $SD \leftarrow 0$ ; //  $SD$  stands for sweep direction. 0 is from left to right; 1 is reverse;
5  $L_2 \leftarrow \emptyset$ ;  $S_2 \leftarrow \emptyset$  // unmatched lines and segments
6 step A: if there are any 3-intersection left then
7     sweep along the direction indicated by  $SD$ ;
8     When a line and segment start at the same time, the segment is seen first.
9     if the event is a line  $l_1$  then
10         // Do not consider the line as a part of a 3-hitter.;
11         // Send to the final edge-cover problem;
12         step B: ;
13          $L \leftarrow L - \{l_1\}$  ;
14          $L_2 \leftarrow L_2 \cup \{l_1\}$  ;
15          $H[SD]++$ ;
16         toggle  $SD$ ;
17     else
18         // the event is a segment  $e_1$ ;
19         if  $e_1$  crosses some line( $s$ ) then
20              $l \leftarrow$  the line hitting  $e_1$  that is closest along  $SD$ ;
21             else
22                 // Do not consider  $e_1$  as a part of a 3-hitter.;
23                 // Send to the final edge-cover problem;
24                  $S \leftarrow S - \{e_1\}$  ;
25                  $S_2 \leftarrow S_2 \cup \{e_1\}$  ;
26                 go to step A;
27             end
28             if  $l$  is critical then
29                 //For example, if  $SD$  is 0, look at the right endpoints of segments
30                 //crossed by  $l$ . Pick the one with the leftmost right endpoint
31                  $e_2 \leftarrow$  the segment crossing  $l$  with the closest endpoint to  $l$  along  $SD$ ;
32                 // put the intersection of  $e_2$  and  $l$  into  $I_3$  ;
33                  $I_3 \leftarrow I_3 \cup (l, e_2)$  ;
34                  $S \leftarrow S - \{e_2\}$  ;
35                  $L \leftarrow L - \{l\}$  ;
36                 go to step A;
37             else
38                 go to step B;
39             end
40         end
41     else
42         The remaining problem  $L \cup L_2$  and  $S \cup S_2$  can be solved optimally as an edge
43         cover problem;
44     end

```

Algorithm 1: Bidirectional sweeping algorithm for HRVL.

$(l_1, e_1), (l_2, e_2), \dots, (l_j, e_j)$. Let t be the size of the maximum matching in G at the time line l is tested for criticality. That is, the maximum matching size drops to $t - 1$ when line l is removed. When line l_h is tested for criticality, the size of the maximum matching is $t - j - 1$ whether l_h is included or not. Let M be a maximum matching when l_h is not included. We can augment M to a matching of size t that does not include line l : add the j pairings from the 3-hitters between l and l_h and then add (l_h, e) . Segment e is not part of

matching M since e is removed from the set of active segments in line 34 of Algorithm 1. This contradicts the criticality of line l . \square

Lemma 3 lh_i cannot be on the right side of lh'_i , $i = 1, 2, \dots, k$.

Proof The proof is by induction.

We prove the base case by contradiction. Figure 7 illustrates the following argument. When $i = 1$, we assume that lh_1 is on the right side of lh'_1 . This means that in S , line lh'_1 is a 3-hitter. Suppose e_1 is the corresponding segment hit by lh'_1 in S . We know in S' line lh'_1 does not hit e_1 (since line lh'_1 is not a 3-hitter in S'). Thus e_1 must be hit by a different 3-hitter in S' , say l_3 (otherwise, lh' can 3-hit e_1 to increase the number of 3-hitters, contradicting the choice of S'). Again in S , l_3 hits e_2 , which means in S' , e_2 must be hit by another line l_4 . Claim 4 guarantees that all of the lines l_i and the segments e_j involved in the tracing process are on the left side of lh_1 . This tracing will stop eventually, because there are only a finite number of lines to the left of line lh_1 . This gives a contradiction.

Now we assume that i is the smallest integer such that lh_i is to the right of lh'_i . We again start tracing from lh'_i . The tracing process can only end at lh_j ($j < i$); otherwise, a contradiction exists, as in the base case. Let the tracing sequence be $lh'_i, e_1, l_3, e_2, l_4, \dots, lh_j$. Since lh'_j is on the left side of lh'_i , so is lh_j . In S' , we replace 3-hitters of S' in the tracing sequence by 3-hitters of S in the sequence. Now, in S' , lh_j becomes an l-hitter instead of lh'_i . This modified S' has improved with the replacement of an l-hitter with a “more capable” (left) l-hitter, while keeping the number of 3-hitters the same. Now, we start a new trace with the new S' .

In summary, if the tracing ends with an l-hitter in S , we improve S' and resume tracing. This process must end in a contradiction because there are only a finite number of lines, and each revised S' moves an l-hitter strictly left. \square

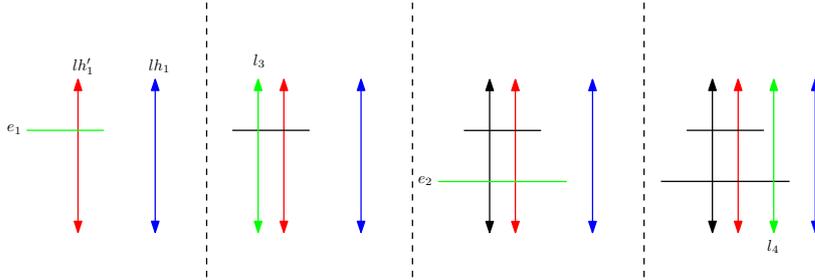


Fig. 7 The tracing sequence for the base case of the proof of Lemma 3. The order items are visited in the tracing is $lh_1, lh'_1, e_1, l_3, e_2, l_4$. The original l-hitter lh_1 is always blue and the original l-hitter lh'_1 is always red. Then for each step of the tracing, the new line or segment is green, with all older elements in black.

An immediate result from this lemma is

$$E(lh_i, S') \leq E(lh'_i, S'). \quad (10)$$

Given a solution P and a line l , let $C(l, P)$ denote the number of segments on the left side of l that have been 3-hit in P . Let $N(l)$ be the total number of segments on the left of line l . The following lemma shows that the segments that S leaves to be double hit are the segments that are easier to double-hit.

Lemma 4 $C(lh_i, S) \geq C(lh_i, S')$, $i = 1, 2, \dots, k$.

Proof We showed in Claim 4 that if a segment is 3-hit to the left of an l-hitter l , then the segment ends before reaching line l . If $C(lh_i, S) < C(lh_i, S')$, then we replace the part of S that is on the left side of lh_i with the corresponding part of S' . This gives us a solution that has more 3-hitters than S has, contradicting the assumption that S has the maximum set of 3-hitters. \square

Therefore we obtain

$$\begin{aligned} E(lh_i, S) &= N(lh_i) - C(lh_i, S) \\ &\leq N(lh_i) - C(lh_i, S') = E(lh_i, S') \leq E(lh'_i, S'). \end{aligned}$$

5 Hitting Lines and Segments

5.1 Hardness

Theorem 5 *Hitting set for horizontal unit segments and vertical lines is NP-complete.*

Proof The reduction is from 3SAT. Consider a 3SAT instance with n variables and m clauses. See Figure 8.

Each variable is represented by a collinear connected set of $2m$ horizontal unit segments, and each clause is represented by a red vertical line that intersects appropriate pairs of horizontal variable segments (if that variable occurs in a clause) or just single segments (in case a variable does not occur in a clause). Setting appropriate parities for the literals in a clause is achieved by appropriate horizontal shifting of the segments, as shown in the figure. This results in a construction in which the only place where three of the elements (segments or lines) can be hit involves a vertical line representing a clause, corresponding to literals occurring in the respective clauses. (These are indicated by magenta circles in the figure.) Let $N = 4mn$ be the number of horizontal segments. This includes the $2m$ per variable just described, and $2m$ more per variable at the bottom of the instance as shown in Figure 8. The N horizontal segments are called variable components.

We show that any feasible hitting set with exactly $N/2$ points induces a truth assignment and vice versa. There is no point that hits more than two of

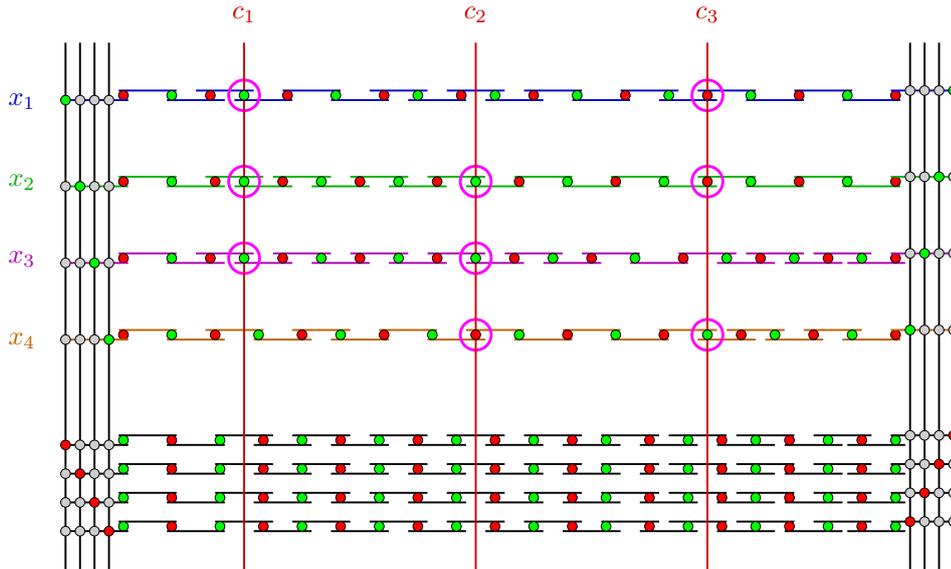


Fig. 8 A set of horizontal unit segments and vertical lines that represents the 3SAT instance $I = (x_1 \vee x_2 \vee x_3) \wedge (x_2 \vee x_3 \vee \bar{x}_4) \wedge (\bar{x}_1 \vee \bar{x}_2 \vee x_4)$. For better visibility, collinear segments are slightly shifted vertically, with red and green points indicating overlapping segments. In an optimal hitting set, the point covering a *labeled* horizontal segment induces a truth value for the corresponding variable: selecting one of its grey points (e.g., in the indicated green manner) assigns a value of “true”; selecting a red point, a value of “false”. Overall, truth assignments for each variable correspond to a set of green or red points, respectively. (Note that there are several equivalent choices from the grey points, which all correspond to the same truth assignments.) Literals occurring in clauses are indicated by magenta circles; these are the only places where a point can hit three segments or lines at once.

the variable components at once. Therefore, stabbing all N of them requires at least $N/2$ points, and any solution consisting of exactly $N/2$ points must hit each variable component exactly once. Consequently, the black vertical lines and the black collinear sets of connected horizontal segments in Figure 8 may be reordered so that, without loss of generality, a solution with exactly $N/2$ points does not pick any of the gray points. Eliminating the gray points results in a natural partition of the instance into point-disjoint even-cardinality loops of variable components for each variable, where the points in each loop alternate between red and green. Thus any solution of size $N/2$ hitting the variable components must select all red or all green points from each variable’s loop, corresponding to a truth assignment. We get an overall feasible hitting set if and only if the points also stab the vertical clause lines, corresponding to a satisfying truth assignment.

□

After appropriate vertical scaling, we can replace the vertical lines by vertical unit segments, immediately giving the following corollary.

Corollary 1 *Deciding if there exists a set of k points in the plane that hit a given set S of unit-length axis-parallel segments is NP-complete.*

We now show APX-hardness for the all-segment case.

Theorem 6 *Computing a minimum hitting set of axis-parallel segments is APX-hard.*

Proof We give a reduction from MAX-2SAT(3), maximum 2-satisfiability in which each variable appears in at most three clauses. MAX-2SAT(3) is known to be APX-hard [3]. In our reduction, a clause is represented by a vertical segment. A variable gadget is a “loop” consisting of at most 8 horizontal segments and exactly 2 vertical segments at the far left/right, linking a chain of an odd number (3, 5, or 7, depending if the variable appears in 1, 2, or 3 clauses) of collinear horizontal segments on the upper portion of the gadget to a single horizontal segment closing the loop along the bottom portion of the gadget. Refer to Figure 9. In total, a variable loop consists of an even number (6, 8, or 10) of segments, whose intersection graph is an even cycle (no three of them intersect). We place red and green points, each representing an edge of the cycle that is the intersection graph, alternating around the cycle. These green/red points occur at crossings with the left/right vertical segments of the loop, or at overlap points along the top chain of horizontal segments of the loop. Membership of a variable x_i in a clause c_j is represented by having the clause segment pass through a green or red point (according to whether the variable or its negation appears in the clause) along the variable loop for x_i , creating a 3-intersection point at the crossing.

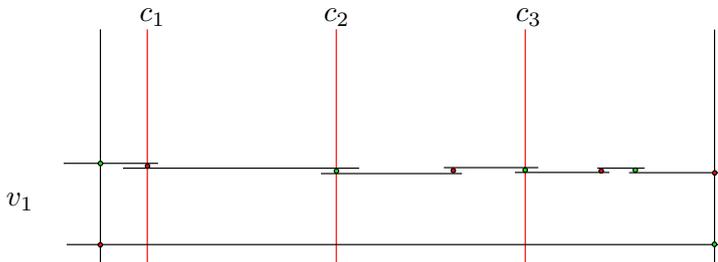


Fig. 9 In this example, clause c_1 includes the literal \bar{v}_1 ; clauses c_2 and c_3 each include the literal v_1 .

Let m and n be the numbers of clauses and variables, respectively, in an instance of MAX-2SAT(3). Then, $\frac{n}{2} \leq m \leq \frac{3n}{2}$. Let k be the number of points in an optimal solution of the corresponding hitting set problem. If the hitting set does not contain any 3-intersection, we know that $k \leq 5n + m \leq 11m$, since all of the segments in each variable loop can be hit using at most 5 hit points.

Suppose \mathcal{A} is an approximation algorithm for the minimum hitting point problem on axis-parallel line segments, and that \mathcal{A} guarantees an approximation factor of $1 + \varepsilon$. For any hitting set H (of size $|H| \leq (1 + \varepsilon)k$) produced by \mathcal{A} , we obtain a solution for the corresponding MAX-2SAT(3) instance, as follows.

Consider the variable loop for x_i . In H , the number, h_i , of 3-intersection hit points along the variable loop could be 0, 1, 2, or 3.

If $h_i = 0$, then we replace the hit points of H along the loop with an optimal set of hit points along the loop – either the green or the red points. This sets the truth value of x_i (green is “true”, red is “false”). Further, this exchange has not caused the number of hit points to go up.

If $h_i = 1$, with one 3-intersection (green or red) point p , we replace the hit points of H along the loop with the (optimal) set of all green or all red hit points along the loop. This sets the truth value of x_i (green is “true”, red is “false”), and this exchange has not caused the number of hit points to go up.

If $h_i = 2$, then the two 3-intersection points might “agree” (be the same color) or “disagree” (be different colors). If they agree, we set the truth value of the variable accordingly, and use an optimal set of hit points of the appropriate color along the loop. If they disagree, then we know that the number of hit points of H used in hitting the variable loop is suboptimal (by at least 1), since H does not use all-red or all-green hit points. Thus, we set the variable either way (use an optimal hitting set of all-green or all-red), and we have a leftover point of H , which we use to hit the clause line that became unhit in the process of setting the hit point set to be monochromatic. There was no increase in the number of hit points.

If $h_i = 3$, then if the three points all agree (are of the same color), we set the truth value of x_i accordingly. Otherwise, we know that the set of points of H used to hit segments in this variable loop is suboptimal; we set the truth value of x_i according to the majority color among the three 3-intersections, and use the one saved hit point to hit the clause that was previously hit by the (minority color) point of H , but now is not.

In this way, we have now transformed H into a set $H_v \cup H_c$, with $|H_v \cup H_c| = |H|$, where H_v is an optimal hitting set for variable loops (using hit points of a single color around each loop), plus a (disjoint) set H_c of additional points to hit clause segments. Let alg be the number of clauses satisfied by the variable setting determined by H_v . Then, we know that $alg \geq m - |H_c|$.

Given an optimal truth assignment for the MAX-2SAT(3) instance, achieving opt satisfied clauses, one way to construct a hitting set for all of the segments in the construction is the following: Optimally place hitting points within each variable loop, according to the truth assignment (and using exactly $|H_v|$ hit points), and then hit the remaining $m - opt$ clause lines, not yet hit by 3-intersections within variable loops, using $m - opt$ separate hit points. Since k is the optimal number of hit points for the whole construction, we know that $k \leq |H_v| + (m - opt)$.

Since we are assuming that \mathcal{A} is a $(1 + \varepsilon)$ -approximation, we have that

$$|H| = |H_v| + |H_c| \leq (1 + \varepsilon)k,$$

which implies that

$$|H_c| \leq (1 + \varepsilon)k - |H_v|.$$

Then, putting these together, and using the facts that $opt \geq \frac{m}{3}$ and that $k \leq 11m$, we have

$$\begin{aligned} alg &\geq m - |H_c| \geq m - (1 + \varepsilon)k + |H_v| \\ &\geq m - (1 + \varepsilon)k + k - m + opt \geq (1 - 33\varepsilon)opt. \end{aligned}$$

This implies that our hitting set problem is APX-hard (a PTAS for the minimum hitting set of axis-parallel segments would imply a PTAS for MAX-2SAT(3)). \square

5.2 Approximation

We give a $5/3$ -approximation for hitting a set V of vertical lines and a set H of horizontal segments. We start by looking at the lower bounds: $v = |V|$ is the number of vertical lines. It is a lower bound. Let h be the lower bound on hitting horizontal segments only. We can compute h exactly; it is the minimum number of hit points for the horizontal segments (computed on each horizontal line). At any stage of the algorithm, we let h and v be the current values of these lower bounds for hitting the current (remaining unhit) sets H and V .

In stage 1, we place two kinds of points:

(a) We place hitting points on vertical lines that reduce h (and v) by one. These points are “maximally productive” since no single hitting point can do more than to reduce h and v each by one. As vertical lines are hit, we remove them from V . Similarly, as horizontal segments are hit, we remove them from H .

(b) Look for pairs (if any) of points, on the same horizontal line and on two vertical lines (from among the current set V), that decrease h by one.

Let k_1 and k_2 be the number of type(a) and type(b) points placed in this stage, respectively. Therefore, for the remaining instance, the lower bound h decreases by $k_1 + k_2/2$, and v decreases by $k_1 + k_2$.

In stage 2, we now have a set of vertical lines V and horizontal segments H such that no single point at the intersection of a vertical line and a horizontal segment (or segments) reduces h , and no pair of points on two distinct vertical lines reduces h .

Lemma 5 *For such sets V and H as in stage 2, an optimal hitting set has size at least $v + h$, where $v = |V|$ and h is the minimum number of points to hit H .*

Proof The hit points we place on V (one per line) might conceivably decrease h . We claim that this cannot happen. Assume to the contrary that it happens. Let $\{q_1, q_2, \dots, q_K\}$ be a minimum-cardinality set such that each of them is on some line of V from left to right and h is decreased by 1 after placing the set. Since the set is minimum, the points in it should be on a horizontal line L .

Since we have found all productive points and pairs of points in stage 1, K should be at least 3. Consider the hit point q_2 . The segments on L that are not hit by q_2 are either completely left or right of q_2 ; let H_l and H_r be the corresponding sets. Points to the left of q_2 do not hit H_r , and points to the right of q_2 do not hit H_l . If adding q_1 decreases H , that means q_1 and q_2 is a productive pair, which should be found in stage 1; otherwise this means that the point q_1 is unnecessary, contradicting the minimality of K . \square

Theorem 7 *There is a polynomial-time $5/3$ -approximation algorithm for geometric hitting set for a set of vertical lines and horizontal segments.*

Proof The total number of points selected by our algorithm is $k_1 + k_2$ from the first stage and $h - k_1 - k_2/2 + v - k_1 - k_2$ from the second stage. By Lemma 5, the points chosen in stage 2 is a lower bound on the cost of an optimal solution:

$$h - k_1 - k_2/2 + v - k_1 - k_2 \leq OPT. \quad (11)$$

We also have $h \leq OPT$ and $v \leq OPT$. There are two cases.

- (i) $k_1 + k_2 \leq 2/3 \cdot OPT$: In this case we select at most $2/3 \cdot OPT$ points in Stages 1, and we use (11) to bound the number of points selected in Stage 2. We conclude that our algorithm selects at most $5/3 \cdot OPT$ points.
- (ii) $k_1 + k_2 > 2/3 \cdot OPT$: The total number of points selected by our algorithm is $h - k_1 - k_2/2 + v \leq 2 \cdot OPT - (k_1 + k_2/2)$. Since $k_1 + k_2/2 \geq k_1/2 + k_2/2 > 1/3 \cdot OPT$, we obtain a $5/3$ -approximation in this case as well. \square

Theorem 8 *There is a polynomial-time $5/3$ -approximation algorithm for geometric hitting set for a set of vertical (downward) rays and horizontal segments.*

Proof The 2-stage approximation algorithm described above works for this case as well. The key observation is that among any set of collinear downward rays, we may remove all but the one with the lowest apex from the instance. Therefore after Stage 1, the hitting points we place on the rays not yet hit will not decrease h . The argument is analogous to that in Lemma 5. \square

6 Hitting Pairs of Segments

We consider now the hitting set problem for inputs that are *unions* of two segments, one horizontal and one vertical. While we are motivated by pairs

(and larger sets) of segments that form paths, our methods apply to general pairs of segments, which might meet to form an “L” shape, a “+”, or a “T” shape, or they may be disjoint. This hitting set problem is NP-hard, since it generalizes the case of horizontal and vertical segments.

Theorem 9 *For objects that are unions of a horizontal and a vertical segment, the hitting set problem has a polynomial-time 4-approximation.*

Proof For ease of discussion, we call the union of two segments an “L”. We use a method similar to those used in [11,23].

Briefly, we do the following. Solve the natural set-cover linear programming (LP) relaxation. Create two new problems: one that has only the horizontal piece of some of the Ls and another that has only the vertical pieces of the remaining Ls. Place an L into the vertical problem if the LP vertical segment has value at least $1/2$, and into the horizontal problem otherwise. Solve the two new problems in polynomial time using the combinatorial method for the 1D problem, or solving the LPs, which are totally unimodular, and thus will return integer solutions. Take all the points selected by either new problem. We prove that these points are a 4-approximation.

In more detail, suppose we have l unions of segments as described above, and let P be the set of points serving as our potential hitters. We assume that $|P|$ is polynomial in l by preprocessing the instance, if necessary, so that we only consider points at endpoints and crossings of segments. For each such union i , we let S_i be the set of points covering the union, while H_i and V_i are the sets of points covering the horizontal segment and vertical segment respectively. We employ the standard set cover linear program (LP) relaxation specialized to our problem:

$$\begin{aligned} \min \quad & \sum_{p \in P} x_p \\ \sum_{p \in H_i} x_p + \sum_{q \in V_i} x_q & \geq 1, \quad \forall 1 \leq i \leq l \\ 0 \leq x_p & \leq 1, \quad \forall p \in P. \end{aligned} \tag{12}$$

We use an optimal LP solution, x^* , to construct a new instance of the problem in which each union contains either a vertical segment or a horizontal segment, but not both. This new instance is easier to approximate but no longer provides a lower bound on the original optimum value, OPT ; however, we show that it provides a lower bound that is within a constant factor of OPT .

For each union of segments i , we set

$$S'_i = \begin{cases} H_i, & \text{if } \sum_{p \in H_i} x_p^* \geq 1/2 \\ V_i, & \text{otherwise.} \end{cases} \tag{13}$$

Now each S'_i corresponds to either a horizontal or vertical segment. Let $H' = \{i \mid S'_i \text{ represents a horizontal segment}\}$, and let $V' = \{1, \dots, l\} \setminus H'$. Our algorithm is as follows:

1. Solve the LP, and let x^* be an optimal solution.
2. Construct S'_i for each union of segments i as described above.
3. Solve the hitting set problems for all the horizontal segments, H' , and all the vertical segments, V' , independently. Return the union of the points, X selected by optimal solutions to each instance.

This algorithm returns a feasible solution since it selects some point in $S'_i \subseteq S_i$ for each union of segments. The first two steps run in polynomial time. Hitting segments of a single orientation is solvable in polynomial time; in fact the corresponding set cover LP relaxation in this case has the consecutive ones property and is totally unimodular, hence the optimum LP value equals the optimum integer solution value.

To see that it is a 4-approximation, let $y_p^* = \min\{2x_p^*, 1\}$ for all p . By (12) and (13) we see that the fractional vector y^* is feasible for the LP instance defined by the segments corresponding to the S'_i . Now we modify the latter LP instance by taking each point p and replacing the variable x_p with variables $x_{p,h}$ and $x_{p,v}$, where $x_{p,h}$ appears only in horizontal segment constraints where x_p formerly appeared, and $x_{p,v}$ appears only in such vertical segment constraints. The resulting LP decouples the horizontal and vertical segments and captures precisely the problem from Step 3 of the algorithm. Since this LP is totally unimodular, we have that the number of chosen points, $|X|$, is at most the cost of any feasible fractional solution. In particular we see that the fractional vector z^* with $z_{p,h}^* = z_{p,v}^* = y_p^*$ is feasible for the decoupled LP, and so:

$$|X| \leq \sum_q z_q^* = 2 \sum_p y_p^*. \quad (14)$$

To obtain our desired result we note that $\sum_p y_p^* \leq 2 \sum_p x_p^*$ by the definition of y^* , yielding $|X| \leq 4 \sum_p x_p^* \leq 4 \cdot OPT$ by (14). \square

The above idea naturally extends to a 4-approximation for the weighted version of the problem. For unions consisting of at most k segments drawn from r orientations, the approach yields a $(k \cdot r)$ -approximation.

The LP-rounding technique in the proof above was introduced by Carr et al. [11] to obtain a 2.1-approximation for the weighted edge-dominating set problem. A similar idea was introduced independently by Gaur et al. [23] to obtain a 2-approximation for stabbing axis-aligned rectangles with horizontal and vertical lines. By using the approach above in conjunction with our approximation algorithm for Theorem 7, we obtain an improved approximation factor in the case that the vertical segments are lines. Before describing this result, we need a slightly stronger version of Theorem 7:

Lemma 6 *There is a polynomial-time $5/3$ -approximation algorithm for hitting a set of vertical lines and horizontal segments that always returns a solution of cost within $5/3$ that of an optimal solution to the natural set cover LP relaxation.*

Proof Given an instance of geometric hitting set over vertical lines and horizontal segments, let LP^* be the optimum value achieved by the natural set

cover LP relaxation. We show that the algorithm used to establish Theorem 7 satisfies stronger versions of the bounds used in the proof of Theorem 7:

$$h \leq LP^*, v \leq LP^*, \text{ and } h - k_1 - k_2/2 + v - k_1 - k_2 \leq LP^*.$$

Since the vertical lines are disjoint, by summing the corresponding LP constraints, we see that $\sum_p x_p \geq v$ for any feasible x . Taking x to be an optimal solution, x^* , we have that $LP^* = \sum_p x_p^* \geq v$. As noted before, the natural set cover LP relaxation is totally unimodular in the case of hitting only horizontal segments. Thus by dropping the constraints corresponding to the lines from the LP, we conclude that $LP^* \geq h$.

For the final bound, we need to show that for the type of instance obtained by our 5/3-approximation in Stage 2, LP^* is equal to OPT' , the optimum size of a hitting set. Lemma 5 shows that for such instances, $OPT' = v' + h'$, where v' and h' are the individual vertical and horizontal lower bounds for the instance.

Consider a collection of collinear horizontal segments from a Stage-2 instance, and remove all points that lie on some vertical line along with all the horizontal segments hit by such points. The proof of Lemma 5 shows that such a deletion does not increase the optimal number of points required to hit such an instance. Hence, appealing to the integrality of such LP instances when dropping the vertical line constraints, we have that $\sum_{p \in P' \setminus P'_V} x_p \geq h'$, where P' is the set of points of a Stage-2 instance, and P'_V is the set of points that lie on some vertical line. Considering only the vertical line constraints, as above, gives us $\sum_{p \in P'_V} x_p \geq v'$. Together, these inequalities yield the desired bound, $\sum_{p \in P'} x_p \geq h' + v'$.

We substitute these bounds in the proof of Theorem 7 to conclude that our algorithm selects at most $5/3 \cdot LP^*$ points instead of $5/3 \cdot OPT$ points. \square

Using similar methods and the above lemma, we also have the following:

Theorem 10 *For objects that are unions of a horizontal segment and a vertical line, the hitting set problem has a polynomial-time 10/3-approximation.*

Proof Our algorithm is essentially the same as the 4-approximation of Theorem 9, with a different last step:

1. Solve the LP, and let x^* be an optimal solution.
2. Construct S'_i for each union i of a horizontal segment and a vertical line.
3. Now each S'_i is either a horizontal segment or a vertical line, and we find a feasible solution X for this instance using our 5/3-approximation. \square

We construct y^* just as in the proof of our 4-approximation; however, now we observe that y^* is feasible for the set cover LP relaxation for the instance defined by the S'_i . This is just an instance of hitting horizontal segments and vertical lines, and so:

$$|X| \leq 5/3 \cdot LP^* \leq 5/3 \cdot \sum_p y_p^*.$$

Since $\sum_p y_p^* \leq 2 \sum_p x_p^*$ as before, we have that $|X| \leq 10/3 \cdot \sum_p x_p^* \leq 10/3 \cdot OPT$ as desired. \square

7 Hitting Triangle-Free Sets of Segments

We consider now the problem in which the n input segments S are allowed to cross or to share endpoints, but not to overlap (i.e., the intersection of any two input segments is not a non-zero length segment – it is either empty or a single point).

Let $G = (V, E)$ denote the (planar) *arrangement graph*, $G(S)$, induced by the segments S ; thus, G has vertex set V equal to the set of all endpoints or crossing points of S and has edge set E joining each pair of vertices that appear consecutively along a segment of S .

We assume that G is *triangle-free*, meaning that it has no cycle of length 3 (i.e., its *girth* is at least 4). It is well known that a planar triangle-free graphs must have a vertex of degree at most 3. (For completeness, we provide the proof: In a triangle-free planar graph (having n nodes, e edges, and f faces), each face has at least 4 edges bounding it. The sum of the number of edges bounding each of the faces is simply $2e$, and in a triangle-free graph must be at least $4f$; thus, $2e \geq 4f$. By Euler's formula ($f - e + n = 1 + c$, for $c \geq 1$ connected components), we get $2e \geq 4(1 + c + e - n) \geq 4(2 + e - n)$, implying that $e \leq 2n - 4$. The sum of the vertex degrees is exactly $2e$ and is thus at most $4n - 8$; thus, not all vertices have degree 4 or more – there must be a vertex of degree at most 3.)

In this section we give a linear-time 3-approximation algorithm for computing a hitting set of points that hit all of the segments of S , assuming that the arrangement graph $G(S)$ is triangle-free. Our approximation factor of 3 matches that obtained recently by Joshi and Narayanaswamy [28]; however, their algorithm employs linear programming, while ours is a simple, combinatorial linear-time ($O(m)$) algorithm.

Our algorithm is the following clipping/shortening process:

- (i) Pick a vertex $v \in V$ of degree at most 3 (it will necessarily be a segment endpoint); such a vertex must exist, by the triangle-free property.
- (ii) Remove the vertex v , and shrink the incident segments with endpoint v to the next adjacent vertex. (In particular, if v is a T-junction, where two of the edges incident to v lie on a common segment, then only the one segment with endpoint at v is shrunk, leaving the other two edges connected.)
- (iii) When shortening a segment s results in segment s becoming a single point (vertex), u , establish a hitting point at u and remove all segments that pass through u .

The following invariants hold at any stage of the process:

- (1) There is at most one remaining subsegment of an input segment (i.e., the portion of an original segment s that remains is connected).
- (2) All segments that have been removed are hit by the hitting points that have been established.
- (3) Any hitting set of the remaining segments, together with the established hitting points already found, forms a hitting set for the original set of input segments.

(4) The graph G remains triangle-free during the process.

The invariants imply that the set of points computed by the algorithm is a valid hitting set. The following lemma establishes the approximation factor:

Lemma 7 *The number of hitting points established by the algorithm is at most 3 times the number, $|H|$, of points in any hitting set H for S .*

Proof Place tokens on the vertices H and consider running the clipping/shortening process on G , with the following actions on the tokens.

When there is a token on the vertex v that is about to be clipped, replace the token with at most 3 clones of it, one on each of the segments that meet at v , allowing each clone to slide along with the endpoint of a clipped segment s as the segment is shrunk, leaving the clone at a new vertex u , the new endpoint of segment s . (There might also be a token at u already; we allow two or more tokens/clones to accumulate at a vertex.) We never clone a clone; if a clone associated with a segment s exists at a vertex v that is being clipped, it remains on segment s , and slides along it as it shrinks. Thus, associated with each point of H there is either a single token or up to 3 clones of the token (but not both).

This ensures that the tokens/clones continue to hit all segments, at all stages of the clipping/shortening process. (Here, we are using the degree-3 property, which allows us to make sure that two edges incident on v that lie on the same segment s are not cut apart at v in our process; thus, a point of H that lies on s continues to hit the shrunk version of segment s . If we had split s at v , with no point of H at v , then no clones are generated at v , and the point(s) of H on segment s may no longer be a valid hitting set for the new arrangement after splitting s at v .) In particular, when a segment shrinks to a point u , there is at least one token/clone present there. Thus, the number of hitting points established by our algorithm is at most $3|H|$, for any hitting set H of S . Letting H be an optimal hitting set, we get that the number of hitting points produced by the algorithm is at most 3 times optimal. \square

Theorem 11 *The algorithm yields a 3-approximation and runs in time $O(m)$, where m is the number of edges in the original (planar) arrangement graph G .*

Proof Immediate, since we only have to maintain the graph G in a standard planar network data structure (e.g., the Doubly Connected Edge List (DCEL) [4]) that allows us to know vertex degrees and perform elementary operations in constant time. \square

8 Conclusion

We have given a variety of new hardness and approximation results for geometric hitting sets involving lines, rays, and segments from a small number of discrete orientations. We have also given a linear-time combinatorial algorithm that yields a 3-approximation for hitting triangle-free sets of non-overlapping

segments in the plane, matching the approximation factor recently obtained by [28] using linear programming methods.

We note that our methods apply as well to yield the same results (lower bounds, upper bounds) for the more general setting in which “segments”, “rays”, and “lines” are given as subsets of families of disjoint pseudoline curves, with each disjoint family playing the role of an “orientation” of lines. (The pseudoline property requires that any two pseudoline curves that intersect (are not “parallel”) do so in a single point of intersection, where they cross.)

Natural open questions ask if any of these approximation bounds can be improved. Notably, we believe that the trivial 2-approximation for hitting segments of two orientations can be improved. Another direction for future research is fixed parameter tractable (FPT) algorithms. Finally, we are interested in optimal coverage versions of these problems in which, e.g., one desires a smallest cardinality set of line segments, rays, or lines, from a small number of orientations, in order to cover a given set of points.

Acknowledgements This work is supported by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-AC04-94AL85000. J. Mitchell acknowledges support from the US-Israel Binational Science Foundation (grant 2010074) and the National Science Foundation (CCF-1018388, CCF-1526406).

References

1. Alon, N.: A non-linear lower bound for planar epsilon-nets. *Discrete & Computational Geometry* **47**(2), 235–244 (2012)
2. Aronov, B., Ezra, E., Sharir, M.: Small-size ϵ -nets for axis-parallel rectangles and boxes. *SIAM Journal on Computing* **39**(7), 3248–3282 (2010)
3. Ausiello, G., Crescenzi, P., Gambosi, G., Kann, V., Marchetti-Spaccamela, A., Protasi, M.: Complexity and approximation: Combinatorial optimization problems and their approximability properties. Springer Science & Business Media (2012)
4. de Berg, M., Cheong, O., van Kreveld, M., Overmars, M.: *Computational Geometry: Algorithms and Applications*, 3rd edn. Springer-Verlag, Berlin, Germany (2008)
5. Brimkov, V.E.: Approximability issues of guarding a set of segments. *International Journal of Computer Mathematics* **90**(8), 1653–1667 (2013)
6. Brimkov, V.E., Leach, A., Mastroianni, M., Wu, J.: Experimental study on approximation algorithms for guarding sets of line segments. In: *Advances in Visual Computing*, pp. 592–601. Springer (2010)
7. Brimkov, V.E., Leach, A., Mastroianni, M., Wu, J.: Guarding a set of line segments in the plane. *Theoretical Computer Science* **412**(15), 1313–1324 (2011)
8. Brimkov, V.E., Leach, A., Wu, J., Mastroianni, M.: Approximation algorithms for a geometric set cover problem. *Disc. App. Math* **160**, 1039–1052 (2012)
9. Brodén, B., Hammar, M., Nilsson, B.J.: Guarding lines and 2-link polygons is APX-hard. In: *Proc. 13th Canadian Conf. Comput. Geometry*, pp. 45–48 (2001). URL <http://www.cccg.ca/proceedings/2001/mikael-2351.ps.gz>
10. Brönnimann, H., Goodrich, M.T.: Almost optimal set covers in finite VC-dimension. *Discrete Comput. Geom.* **14**, 263–279 (1995)
11. Carr, R.D., Fujito, T., Konjevod, G., Parekh, O.: A $2 \frac{1}{10}$ -approximation algorithm for a generalization of the weighted edge-dominating set problem. In: *Algorithms-ESA 2000*, pp. 132–142. Springer (2000)

12. Chvátal, V.: A greedy heuristic for the set-covering problem. *Mathematics of operations research* **4**(3), 233–235 (1979)
13. Clarkson, K.L.: Algorithms for polytope covering and approximation. In: *Algorithms and Data Structures*, pp. 246–252. Springer (1993)
14. Clarkson, K.L., Varadarajan, K.: Improved approximation algorithms for geometric set cover. *Discrete & Computational Geometry* **37**(1), 43–58 (2007)
15. Dinur, I., Steurer, D.: Analytical approach to parallel repetition. In: *Proceedings of the 46th Annual ACM Symposium on Theory of Computing, STOC '14*, pp. 624–633. ACM, New York, NY, USA (2014). DOI 10.1145/2591796.2591884. URL <http://doi.acm.org/10.1145/2591796.2591884>
16. Dom, M., Fellows, M.R., Rosamond, F.A.: Parameterized complexity of stabbing rectangles and squares in the plane. In: *WALCOM: Algorithms and Computation*, pp. 298–309. Springer (2009)
17. Duh, R.c., Fürer, M.: Approximation of k-set cover by semi-local optimization. In: *Proceedings of the Twenty-ninth Annual ACM Symposium on Theory of Computing, STOC '97*, pp. 256–264. ACM, New York, NY, USA (1997). DOI 10.1145/258533.258599. URL <http://doi.acm.org/10.1145/258533.258599>
18. Dumitrescu, A., Jiang, M.: On the approximability of covering points by lines and related problems. arXiv preprint arXiv:1312.2549 (2013)
19. Dumitrescu, A., Jiang, M.: On the approximability of covering points by lines and related problems. *Comput. Geom.* **48**(9), 703–717 (2015). DOI 10.1016/j.comgeo.2015.06.006. URL <http://dx.doi.org/10.1016/j.comgeo.2015.06.006>
20. Even, G., Levi, R., Rawitz, D., Schieber, B., Shahar, S.M., Sviridenko, M.: Algorithms for capacitated rectangle stabbing and lot sizing with joint set-up costs. *ACM Transactions on Algorithms (TALG)* **4**(3), 34 (2008)
21. Even, G., Rawitz, D., Shahar, S.M.: Hitting sets when the vc-dimension is small. *Information Processing Letters* **95**(2), 358–362 (2005)
22. Gaur, D.R., Bhattacharya, B.: Covering points by axis parallel lines. In: *Proc. 23rd European Workshop on Computational Geometry*, pp. 42–45. Citeseer (2007)
23. Gaur, D.R., Ibaraki, T., Krishnamurti, R.: Constant ratio approximation algorithms for the rectangle stabbing problem and the rectilinear partitioning problem. In: *Algorithms-ESA 2000*, pp. 211–219. Springer (2000)
24. Giannopoulos, P., Knauer, C., Rote, G., Werner, D.: Fixed-parameter tractability and lower bounds for stabbing problems. *Comp. Geometry* **46**, 839–860 (2013)
25. Hassin, R., Megiddo, N.: Approximation algorithms for hitting objects with straight lines. *Discrete Applied Mathematics* **30**(1), 29–42 (1991)
26. Heednacram, A.: The NP-hardness of covering points with lines, paths and tours and their tractability with FPT-algorithms. Griffith University (2010)
27. Hochbaum, D.S., Maas, W.: Approximation schemes for covering and packing problems in image processing and VLSI. *J. ACM* **32**, 130–136 (1985)
28. Joshi, A., Narayanaswamy, N.: Approximation algorithms for hitting triangle-free sets of line segments. In: *SWAT*, pp. 357–367. Springer (2014)
29. Kovaleva, S., Spieksma, F.C.: Approximation algorithms for rectangle stabbing and interval stabbing problems. *SIAM J. Discrete Math.* **20**(3), 748–768 (2006)
30. Kratsch, S., Philip, G., Ray, S.: Point line cover: the easy kernel is essentially tight. In: *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1596–1606. SIAM (2014)
31. Kumar, V.A., Arya, S., Ramesh, H.: Hardness of set cover with intersection 1. In: *Automata, Languages and Programming*, pp. 624–635. Springer (2000)
32. Langerman, S., Morin, P.: Covering things with things. *Discrete & Computational Geometry* **33**(4), 717–729 (2005)
33. Megiddo, N., Tamir, A.: On the complexity of locating linear facilities in the plane. *Operations Research Letters* **1**(5), 194–197 (1982)
34. Mustafa, N.H., Ray, S.: Improved results on geometric hitting set problems. *Discrete & Computational Geometry* **44**(4), 883–895 (2010)
35. O'Rourke, J.: *Art Gallery Theorems and Algorithms*. The International Series of Monographs on Computer Science. Oxford University Press, New York, NY (1987)
36. Pach, J., Tardos, G.: Tight lower bounds for the size of epsilon-nets. *Journal of the American Mathematical Society* **26**(3), 645–658 (2013)

-
37. Urrutia, J., et al.: Art gallery and illumination problems. *Handbook of Computational Geometry* **1**(1), 973–1027 (2000)