Stabbing Rectangles by Line Segments – How Decomposition Reduces the Shallow-Cell Complexity

Timothy M. Chan
University of Illinois at Urbana-Champaign, U.S.A.
tmc@illinois.edu

Thomas C. van Dijk[1]
Universität Würzburg, Germany
https://orcid.org/0000-0001-6553-7317

Krzysztof Fleszar[2]
Max-Planck-Institut für Informatik, Saarbrücken, Germany
https://orcid.org/0000-0002-1129-3289

Joachim Spoerhase[3]
Aalto University, Espoo, Finland
Universität Würzburg, Germany
joachim.spoerhase@uni-wuerzburg.de
https://orcid.org/0000-0002-2601-6452

Alexander Wolff
Universität Würzburg, Germany
https://orcid.org/0000-0001-5872-718X

Abstract

We initiate the study of the following natural geometric optimization problem. The input is a set of axis-aligned rectangles in the plane. The objective is to find a set of horizontal line segments of minimum total length so that every rectangle is *stabbed* by some line segment. A line segment stabs a rectangle if it intersects its left and its right boundary. The problem, which we call Stabbing, can be motivated by a resource allocation problem and has applications in geometric network design. To the best of our knowledge, only special cases of this problem have been considered so far.

Stabbing is a weighted geometric set cover problem, which we show to be NP-hard. While for general set cover the best possible approximation ratio is $\Theta(\log n)$, it is an important field in geometric approximation algorithms to obtain better ratios for geometric set cover problems. Chan et al. [SODA'12] generalize earlier results by Varadarajan [STOC'10] to obtain sub-logarithmic performances for a broad class of *weighted* geometric set cover instances that are characterized by having low *shallow-cell complexity*. The shallow-cell complexity of Stabbing instances, however, can be high so that a direct application of the framework of Chan et al. gives only logarithmic bounds. We still achieve a constant-factor approximation by decomposing general instances into what we call *laminar* instances that have low enough complexity.

Our decomposition technique yields constant-factor approximations also for the variant where rectangles can be stabbed by horizontal and vertical segments and for two further geometric set cover problems.

[1] Supported by DFG grant DI 2161/2-1.
[2] This research was partially supported by Conicyt Grant PII 20150140 and by Millennium Nucleus Information and Coordination in Networks RC130003.
[3] Supported by European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant number 759557) and by Academy of Finland (grant number 310415).

29th International Symposium on Algorithms and Computation (ISAAC 2018).
Editors: Wen-Lian Hsu, Der-Tsai Lee, and Chung-Shou Liao; Article No. 61; pp. 61:1–61:13
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

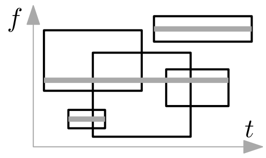

Figure 1 An instance of STABBING (rectangles) with an optimal solution (gray line segments).

2012 ACM Subject Classification Theory of computation → Packing and covering problems, Theory of computation → Computational geometry

Keywords and phrases Geometric optimization, NP-hard, approximation, shallow-cell complexity, line stabbing

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2018.61

Related Version A full version is available at `https://arxiv.org/abs/1806.02851`.

1 Introduction

In this paper, we study the following geometric optimization problem, which we call STABBING. The input is a set R of n axis-aligned rectangles in the plane. The objective is to find a set S of horizontal line segments of minimum total length $\|S\|$, where $\|S\| = \sum_{s \in S} \|s\|$, such that each rectangle $r \in R$ is stabbed by some line segment $s \in S$. Here, we say that s stabs r if s intersects the left and the right edge of r (see Fig. 1). The length of a line segment s is denoted by $\|s\|$. Throughout this paper, rectangles are assumed to be axis-aligned and segments are horizontal line segments (unless explicitly stated otherwise).

Our problem can be viewed as a resource allocation problem. Consider a server that receives a number of communication requests. Each request r is specified by a time window $[t_1, t_2]$ and a frequency band $[f_1, f_2]$. In order to satisfy the request r, the server has to open a communication channel that is available in the time interval $[t_1, t_2]$ and operates at a fixed frequency within the frequency band $[f_1, f_2]$. Therefore, the server has to open several channels over time so that each request can be fulfilled. Requests may share the same channel if their frequency bands and time windows overlap. Each open channel incurs a fixed cost per time unit and the goal is to minimize the total cost. Consider a t–f coordinate system. A request r can be identified with a rectangle $[t_1, t_2] \times [f_1, f_2]$. An open channel corresponds to horizontal line segments and the operation cost equals its length. Satisfying a request is equivalent to stabbing the corresponding rectangle.

To the best of our knowledge, general STABBING has not been studied, although it is a natural problem. Finke et al. [10] consider the special case of the problem where the left sides of all input rectangles lie on the y-axis. They derive the problem from a practical application in the area of batch processing and give a polynomial time algorithm that solves this special case of STABBING to optimality. Das et al. [6] describe an application of STABBING in geometric network design. They obtain a constant-factor approximation for a slight generalization of the special case of Finke et al. in which rectangles are only constrained to *intersect* the y-axis. This result constitutes the key step for an $O(\log n)$-approximation algorithm to the GENERALIZED MINIMUM MANHATTAN NETWORK problem.

We also consider the following variant of our problem, which we call CONSTRAINED STABBING. Here, the input additionally consists of a set F of horizontal line segments of which any solution S must be a subset.

Related Work. STABBING can be interpreted as a weighted geometric set cover problem where the rectangles play the role of the elements, the potential line segments correspond to the sets and a segment s "contains" a rectangle r if s stabs r. The weight of a segment s equals its length $\|s\|$. SET COVER is one of the classical NP-hard problems. The greedy algorithm yields a $\ln n$-approximation (where n is the number of elements) and this is known to be the best possible approximation ratio for the problem unless $\mathsf{P} = \mathsf{NP}$ [9, 7]. It is an important research direction of computational geometry to surpass the lower bound known for general SET COVER in geometric settings. In their seminal work, Brönniman and Goodrich [3] gave an $O(\log \mathrm{OPT})$-approximation algorithm for *unweighted* SET COVER, where OPT is the size of an optimum solution, for the case when the underlying *VC-dimension* is constant. This holds in many geometric settings. Numerous subsequent works have improved upon this result in specific geometric settings. For example, Aronov et al. [1] obtained an $O(\log \log \mathrm{OPT})$-approximation algorithm for the problem of piercing a set of axis-aligned rectangles with the minimum number of points (HITTING SET for axis-aligned rectangles) by means of so-called *ε-nets*. Mustafa and Ray [17] obtained a PTAS for the case of piercing pseudo-disks by points. A limitation of these algorithms is that they only apply to *unweighted* geometric SET COVER; hence, we cannot apply them directly to our problem. In a break-through, Varadarajan [18] developed a new technique, called *quasi-uniform sampling*, that gives sub-logarithmic approximation algorithms for a number of *weighted* geometric set cover problems (such as covering points with weighted fat triangles or weighted disks). Subsequently, Chan et al. [5] generalized Varadarajan's idea. They showed that quasi-uniform sampling yields a sub-logarithmic performance if the underlying instances have low shallow-cell complexity. Bansal and Pruhs [2] presented an interesting application of Varadarajan's technique. They reduced a large class of scheduling problems to a particular geometric set cover problem for anchored rectangles and obtained a constant-factor approximation via quasi-uniform sampling. Recently, Chan and Grant [4] and Mustafa et al. [16] settled the APX-hardness status of all natural weighted geometric SET COVER problems where the elements to be covered are points in the plane or space.

Gaur et al. [12] considered the problem of stabbing a set of axis-aligned rectangles by a minimum number of axis-aligned *lines*. They obtain an elegant 2-approximation algorithm for this NP-hard problem by rounding the standard LP-relaxation. Kovaleva and Spieksma [14] considered a generalization of this problem involving weights and demands. They obtained a constant-factor approximation for the problem. Even et al. [8] considered a *capacitated* variant of the problem in arbitrary dimension. They obtained approximation ratios that depend linearly on the dimension and extended these results to approximate certain lot-sizing inventory problems. Giannopoulos et al. [13] investigated the fixed-parameter tractability of the problem where given translated copies of an object are to be stabbed by a minimum number of lines (which is also the parameter). Among others, they showed that the problem is W[1]-hard for unit-squares but becomes FPT if the squares are disjoint.

Our Contribution. We are the first to investigate STABBING in this general form: horizontal line segments stabbing axis-aligned rectangles without further restrictions. We examine the complexity and the approximability of this problem.

We rule out the possibility of efficient exact algorithms by showing that STABBING is NP-hard; see Section 4. Another negative result is that STABBING instances can have high shallow-cell complexity so that a direct application of the quasi-uniform sampling method yields only the same logarithmic bound as for arbitrary set cover instances; see Section 2.2.

Our main result is a constant-factor approximation algorithm for Stabbing; see Section 2. Our algorithm is based on the following three ideas. First, we show a simple decomposition lemma that implies a constant-factor approximation for (general) set cover instances whose set family can be decomposed into two disjoint sub-families each of which admits a constant-factor approximation. Second, we show that Stabbing instances whose segments have a special *laminar* structure have low enough shallow-cell complexity so that they admit a constant-factor approximation by quasi-uniform sampling. Third, we show that an arbitrary instance can be transformed in such a way that it can be decomposed into two disjoint laminar families. Together with the decomposition lemma, this establishes the constant-factor approximation.

Another (this time more obvious) application of the decomposition lemma gives also a constant-factor approximation for the variant of Stabbing where we allow horizontal and vertical stabbing segments. Also in this case, a direct application of quasi-uniform sampling gives only a logarithmic bound as there are laminar families of horizontal *and* vertical segments that have high shallow-cell complexity. This and two further applications of the decomposition lemma are sketched in Section 3.

The above results provide two natural examples for the fact that the property of having low shallow-cell complexity is *not* closed under the union of the set families. In spite of this, constant-factor approximations are still possible. Our results also show that the representation as a union of low-complexity families may not be obvious at first glance. We therefore hope that our approach helps to extend the reach of quasi-uniform sampling beyond the concept of low shallow-cell complexity also in other settings. Our results for Stabbing may also lead to new insights for other related geometric problems such as the Generalized Minimum Manhattan Network problem [6].

Due to space constraints, we refer the reader for further results such as the APX-hardness of Constrained Stabbing and the relationship of Stabbing to well-studied geometric set cover (or equivalently hitting set) problems to the full version of our paper (see page 2).

2 A Constant-Factor Approximation Algorithm for Stabbing

In this section, we present a constant-factor approximation algorithm for Stabbing. First, we model Stabbing as a set cover problem, and we revisit the standard linear programming relaxation for set cover and the concept of shallow-cell complexity; see Sections 2.1 and 2.2. Then, we observe that there are Stabbing instances with high shallow-cell complexity. This limiting fact prevents us from obtaining any constant approximation factor if applying the generalization of Chan et al. [5] in a direct way; see Section 2.2. In order to bypass this limitation, we decompose any Stabbing instance into two disjoint families of low shallow-cell complexity. Before describing the decomposition in Section 2.5, we show how to merge solutions to these two disjoint families in an approximation-factor preserving way; see Section 2.3. Then, in Section 2.4, we observe that these families have sufficiently small shallow-cell complexity to admit a constant-factor approximation.

2.1 Set Cover and Linear Programming

An instance $(U, \mathcal{F}, c)$ of weighted Set Cover is given by a finite universe U of n elements, a family $\mathcal{F}$ of subsets of U that covers U, and a cost function $c\colon \mathcal{F} \to \mathbb{Q}^+$. The objective is to find a sub-family $\mathcal{S}$ of $\mathcal{F}$ that also covers U and minimizes the total cost $c(\mathcal{S}) = \sum_{S \in \mathcal{S}} c(S)$.

An instance (R, F) of Constrained Stabbing, given by a set R of rectangles and a set F of line segments, can be seen as a special case of weighted Set Cover where the rectangles in R are the universe U, the line segments in F form the sets in $\mathcal{F}$, and a line segment $s \in F$ "covers" a rectangle r if and only if s stabs r. Unconstrained Stabbing can be modeled by Set Cover as follows. We can, without loss of generality, consider only feasible solutions where the end points of any line segment lie on the left or right boundaries of rectangles and where each line segment touches the top boundary of some rectangle. Thus, we can restrict ourselves to feasible solutions that are subsets of a set F of $O(n^3)$ candidate line segments. This shows that Stabbing is a special case of Constrained Stabbing and, hence, of Set Cover.

The standard LP relaxation $\mathrm{LP}(U, \mathcal{F}, c)$ for a Set Cover instance $(U, \mathcal{F}, c)$ is as follows:

$$\begin{array}{lll} \text{Minimize} & \displaystyle\sum_{S \in \mathcal{F}} c(S) z_S & \\ \text{subject to} & \displaystyle\sum_{S \in \mathcal{F}, S \ni e} z_S \geq 1 & \text{for all } e \in U, \\ & z_S \geq 0 & \text{for all } S \in \mathcal{F}. \end{array}$$

The optimum solution to this LP provides a lower bound on OPT. An algorithm is called *LP-relative α-approximation algorithm* for a class Π of set cover instances if it rounds any feasible solution $\mathbf{z} = (z_S)_{S \in \mathcal{F}}$ to the above standard LP relaxation for some instance $(U, \mathcal{S}, c)$ in this class to a feasible integral solution $\mathcal{S} \subseteq \mathcal{F}$ of cost $c(\mathcal{S}) \leq \alpha \sum_{S \in \mathcal{F}} c(S) z_s$.

2.2 Shallow-Cell Complexity

We define the shallow-cell complexity for classes that consist of instances of weighted Set Cover. Informally, the shallow-cell complexity is a bound on the number of equivalent classes of elements that are contained in a small number of sets. Here is the formal definition.

▶ **Definition 1** (Chan et al. [5]). Let $f(m, k)$ be a function non-decreasing in m and k. An instance $(U, \mathcal{F}, c)$ of weighted Set Cover has shallow-cell complexity f if the following holds for every k and m with $1 \leq k \leq m \leq |\mathcal{F}|$, and every sub-family $\mathcal{S} \subseteq \mathcal{F}$ of m sets: All elements that are contained in at most k sets of $\mathcal{S}$ form at most $f(m, k)$ equivalence classes (called *cells*), where two elements are equivalent if they are contained in precisely the same sets of $\mathcal{S}$. A class of instances of weighted Set Cover has shallow-cell complexity f if all its instances have shallow-cell complexity f.

Chan et al. proved that if a set cover problem has low shallow-cell complexity then quasi-uniform sampling yields an LP-relative approximation algorithm with good performance.

▶ **Theorem 2** (Chan et al. [5]). *Let $\varphi(m)$ be a non-decreasing function, and let Π be a class of instances of weighted* Set Cover. *If Π has shallow-cell complexity $m\varphi(m)k^{O(1)}$, then Π admits an LP-relative approximation algorithm (based on quasi-uniform sampling) with approximation ratio $O(\max\{1, \log \varphi(m)\})$.*

Unfortunately, there are instances of Stabbing (and its constrained variants) that have high shallow-cell complexity, so we cannot directly obtain a sub-logarithmic performance via Theorem 2. These instances can be constructed as follows; see Fig. 2a. Let m be an even positive integer. For $i = 1, \ldots, m$, define the point $p_i = (i, i)$. For each pair i, j with $1 \leq i \leq m/2 < j \leq m$, let r_{ij} be the rectangle with corners p_i and p_j. Now, consider the following set $\mathcal{S}$ of m line segments. For $i = 1, \ldots, m/2$, the set $\mathcal{S}$ contains the segment s_i with endpoints p_i and (m, i). For $i = m/2 + 1, \ldots, m$, the set $\mathcal{S}$ contains the segment s_i with endpoints $(1, i)$ and p_i. We want to count the number of rectangles that are stabbed by at most two segments in $\mathcal{S}$. Consider any i and j satisfying $1 \leq i \leq m/2 < j \leq m$. Observe that

(a) for Stabbing **(b)** for Horizontal–Vertical Stabbing

Figure 2 Instances with high shallow-cell complexity.

the rectangle r_{ij} is stabbed precisely by the segments s_i and s_j in $\mathcal{S}$. Hence, according to Definition 1, our instance consists of at least $m^2/4$ equivalence classes for $k = 2$. Thus, if our instance has shallow cell-complexity f for some suitable function f, we have $f(m, 2) = \Omega(m^2)$. Since f is non-decreasing, we also have $f(m, k) = \Omega(m^2)$ for $k \geq 2$. Hence, Theorem 2 implies only an $O(\log n)$-approximation algorithm for Stabbing (and its constrained variants) where we use the above-mentioned fact (see Section 2.1) that we can restrict ourselves to $m = O(n^3)$ many candidate segments.

2.3 Decomposition Lemma for Set Cover

Our trick is to decompose general instances of Stabbing (which may have high shallow-cell complexity) into partial instances of low complexity with a special, laminar structure. We use the following simple decomposition lemma, which holds for arbitrary set cover instances.

▶ **Lemma 3.** *Let Π, Π_1, Π_2 be classes of* Set Cover *where Π_1 and Π_2 admit LP-relative α_1- and α_2-approximation algorithms, respectively. The class Π admits an LP-relative $(\alpha_1 + \alpha_2)$-approximation algorithm if, for every instance $(U, \mathcal{F}, c) \in \Pi$, the family $\mathcal{F}$ can be partitioned into $\mathcal{F}_1, \mathcal{F}_2$ such that, for any partition of U into U_1, U_2 where U_1 is covered by $\mathcal{F}_1$ and U_2 by $\mathcal{F}_2$, the instances $(U_1, \mathcal{F}_1, c)$ and $(U_2, \mathcal{F}_2, c)$ are instances of Π_1 and Π_2, respectively.*

Proof. Let $\mathbf{z} = (z_S)_{S \in \mathcal{F}}$ be a feasible solution to $\text{LP}(U, \mathcal{F}, c)$. Let $U_1, U_2 = \emptyset$ initially. Consider an element $e \in U$. Because of the constraint $\sum_{S \in \mathcal{F}, S \ni e} z_S \geq 1$ in the LP relaxation and because of $\mathcal{F} = \mathcal{F}_1 \cup \mathcal{F}_2$, at least one of the two cases $\sum_{S \in \mathcal{F}_1, S \ni e} z_S \geq \alpha_1/(\alpha_1 + \alpha_2)$ and $\sum_{S \in \mathcal{F}_2, S \ni e} z_S \geq \alpha_2/(\alpha_1 + \alpha_2)$ occurs. If the first case holds, we add e to U_1. Otherwise, the second case holds and we add e to U_2. We execute this step for each element $e \in U$.

Now, consider the instance $(U_1, \mathcal{F}_1, c)$. For each $S \in \mathcal{F}_1$, set $z^1_S := \min\{z_S(\alpha_1 + \alpha_2)/\alpha_1, 1\}$. Since $\sum_{S \in \mathcal{F}_1, S \ni e} z_S \geq \alpha_1/(\alpha_1 + \alpha_2)$ for all $e \in U_1$, we have that $\mathbf{z}^1 = (z^1_S)_{S \in \mathcal{F}_1}$ forms a feasible solution to $\text{LP}(U_1, \mathcal{F}_1, c)$. Next, we apply the LP-relative α_1-approximation algorithm to this instance to obtain a solution $\mathcal{S}_1 \subseteq \mathcal{F}_1$ that covers U_1 and whose cost is at most $\alpha_1 \sum_{S \in \mathcal{F}_1} c(S) z^1_S \leq (\alpha_1 + \alpha_2) \sum_{S \in \mathcal{F}_1} c(S) z_S$. Analogously, we can compute a solution $\mathcal{S}_2 \subseteq \mathcal{F}_2$ to $(U_2, \mathcal{F}_2, c)$ of cost at most $(\alpha_1 + \alpha_2) \sum_{S \in \mathcal{F}_2} c(S) z_S$.

To complete the proof, note that $\mathcal{S}_1 \cup \mathcal{S}_2$ is a feasible solution to $(U, \mathcal{F}, c)$ of cost at most $(\alpha_1 + \alpha_2) \sum_{S \in \mathcal{F}_1 \cup \mathcal{F}_2} c(S) z_S$. Hence, our algorithm is an LP-relative $(\alpha_1 + \alpha_2)$-approximation algorithm. ◀

2.4 x-Laminar Instances

▶ **Definition 4.** An instance of Constrained Stabbing is called *x-laminar* if the projection of the segments in this instance onto the x-axis forms a laminar family of intervals. That is, any two of these intervals are either interior-disjoint or one is contained in the other.

We remark that for an x-laminar instance of Constrained Stabbing the corresponding instance $(U, \mathcal{F}, c)$ of Set Cover does not necessarily have a laminar set family $\mathcal{F}$.

▶ **Lemma 5.** *The shallow-cell complexity of an x-laminar instance of* Constrained Stabbing *can be upper bounded by $f(m,k) = mk^2$. Hence, such instances admit a constant-factor LP-relative approximation algorithm.*

Proof. To prove the bound on the shallow-cell complexity, consider a set $\mathcal{S}$ of m segments. Let $1 \le k \le m$ be an integer. Consider an arbitrary rectangle r that is stabbed by at most k segments in $\mathcal{S}$. Let $\mathcal{S}_r$ be the set of these segments. Consider a shortest segment $s \in \mathcal{S}_r$. By laminarity, the projection of any segment in $\mathcal{S}_r$ onto the x-axis contains the projection of s onto the x-axis. Let $C_s = (s_1, \ldots, s_\ell)$ be the sequence of *all* segments in $\mathcal{S}$ whose projection contains the projection of s, ordered from top to bottom. The crucial point is that the set $\mathcal{S}_r$ forms a contiguous sub-sequence $s_i, \ldots, s_{i+|\mathcal{S}_r|-1}$ of C_s that contains $s = s_j$ for some $i \le j \le i + |\mathcal{S}_r| - 1$. Hence, $\mathcal{S}_r$ is uniquely determined by the choice of $s \in \mathcal{S}$ (for which there are m possibilities), the choice of s_i with $i \in \{j-k, \ldots, j\}$ within the sequence C_s (for which there are at most k possibilities), and the cardinality of $\mathcal{S}_r$ (for which there are at most k possibilities). This implies that $\mathcal{S}_r$ is one of mk^2 many sets that define a cell. This completes our proof since r was picked arbitrarily. ◀

2.5 Decomposing General Instances into Laminar Instances

▶ **Lemma 6.** *Given an instance I of (unconstrained)* Stabbing *with rectangle set R, we can compute an instance $I' = (R, F)$ of* Constrained Stabbing *with the following properties. The set F of segments in I' has cardinality $O(n^3)$, it can be decomposed into two disjoint x-laminar sets F_1 and F_2, and $\mathrm{OPT}_{I'} \le 6 \cdot \mathrm{OPT}_I$.*

Proof. Let F' be the set of $O(n^3)$ candidate segments as defined in Sec. 2.1: For every segment s of F', the left endpoint of s lies on the left boundary of some rectangle, the right endpoint of s lies on the right boundary of some rectangle, and s contains the top boundary of some rectangle. Recall that F' contains the optimum solution.

Below, we stretch each of the segments in F' by a factor of at most 6 to arrive at a set F of segments having the claimed properties. By scaling the instance we may assume that the longest segment in F' has length 1/3.

For any $i, j \in \mathbb{Z}$ with $i \ge 0$, let I_{ij} be the interval $[j/2^i, (j+1)/2^i]$. Let $\mathcal{I}_1$ be the family of all such intervals I_{ij}. We say that I_{ij} has level i. Note that $\mathcal{I}_1$ is an x-laminar family of intervals (segments). Let $\mathcal{I}_2$ be the family of intervals that arises if each interval in $\mathcal{I}_1$ is shifted to the right by the amount of 1/3. That is, $\mathcal{I}_2$ is the family of all intervals of the form $I_{ij} + 1/3 := [j/2^i + 1/3, (j+1)/2^i + 1/3]$ (for any $i, j \in \mathbb{Z}$ with $i \ge 0$). Clearly, $\mathcal{I}_2$ is x-laminar, too.

We claim that any arbitrary interval $J = [a, b]$ of length at most 1/3 is contained in an interval I that is at most 6 times longer than J and that is contained in $\mathcal{I}_1$ or in $\mathcal{I}_2$. This completes the proof of the lemma since then any segment in F' can be stretched by a factor of at most 6 so that its projection on the x-axis lies in $\mathcal{I}_1$ (giving rise to the segment set F_1) or in $\mathcal{I}_2$ (giving rise to the segment set F_2). Setting $F = F_1 \cup F_2$ completes the construction of the instance $I' = (R, F)$.

To show the above claim, let s be the largest non-negative integer with $b - a \le 1/(3 \cdot 2^s)$. If J is contained in the interval I_{sj} for some integer j, we are done because $b - a > 1/(6 \cdot 2^s)$ by the choice of s. If J is not contained in any interval I_{sj}, then there exists some integer j such that $j/2^s \in J = [a, b]$ and thus $a \in I_{s,j-1}$. Since $b - a \le 1/(3 \cdot 2^s)$, we have that J is completely contained in the interval $I' := I_{s,j-1} + 1/(3 \cdot 2^s)$ and in the interval $I'' := I_{s,j} - 1/(3 \cdot 2^s)$.

We complete the proof by showing that one of the intervals I', I'' is actually contained in $\mathcal{I}_2$. To this end, note that $1/3 = \sum_{\ell=1}^{\infty}(-1)^{\ell-1}/2^{\ell}$. Hence, if s is even, the interval $I' - 1/3$ lies in $\mathcal{I}_1$, and if s is odd, the interval $I'' - 1/3$ lies in $\mathcal{I}_1$. ◀

Applying the decomposition lemma to Lemmas 5 and 6 yields our main result. We do not give an explicit approximation factor due to our reliance on the result by Chan et al. [5]. We also cannot apply a decomposition technique similar to CONSTRAINED STABBING since Lemma 6 requires a free choice of the set F of stabbing line segments.

▶ **Theorem 7.** STABBING *admits a constant-factor LP-relative approximation algorithm.*

Complementing Lemmas 5 and 6, Fig. 2a shows that the union of two x-laminar families of segments may have shallow-cell complexity with quadratic dependence on m. Hence, the property of having low shallow-cell complexity is not closed under taking unions.

3 Further Applications of the Decomposition Lemma

Here we show that our decomposition technique can be applied in other settings, too.

Horizontal–Vertical Stabbing. In this new variant of STABBING, a rectangle may be stabbed by a horizontal or by a vertical line segment (or by both). Using the results of Section 2.5 and the decomposition lemma where we decompose into horizontal and vertical segments, we immediately obtain the following result.

▶ **Corollary 8.** HORIZONTAL–VERTICAL STABBING *admits an LP-relative constant-factor approximation algorithm.*

Figure 2b shows that a laminar family of horizontal segments and vertical segments may have a shallow-cell complexity with quadratic dependence on m. Thus, Corollary 8 is another natural example where low shallow-cell complexity is not closed under union and where the decomposition lemma gives a constant-factor approximation although the shallow-cell complexity is high.

Stabbing 3D-Boxes by Squares. In the 3D-variant of STABBING, we want to stab 3D-boxes with axis-aligned squares, minimizing the sum of the areas or the sum of the perimeters of the squares. Here, "stabbing" means "completely cutting across". By combining the same idea with shifted quadtrees – the 2D-equivalent of laminar families of intervals – we obtain a constant-factor approximation for this problem. It is an interesting question if our approach can be extended to handle also arbitrary rectangles but this seems to require further ideas.

Covering Points by Anchored Squares. Given a set P of points that need to be covered and a set A of anchor points, we want to find a set of axis-aligned squares such that each square contains at least one anchor point, the union of the squares covers P, and the total area or the total perimeter of the squares is minimized. Again, with the help of shifted quadtrees, we can apply the decomposition lemma. In this case, we do not even need to apply the machinery of quasi-uniform sampling; instead, we can use dynamic programming on the decomposed instances. This yields a deterministic algorithm with a concrete constant approximation ratio ($4 \cdot 6^2$, without polishing).

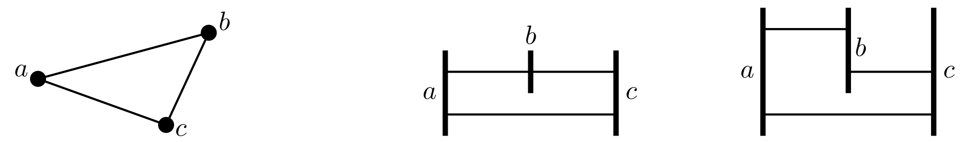

(a) a Planar Vertex Cover instance **(b)** a visibility representation **(c)** no coinciding edges

Figure 3 Obtaining a visibility representation from a Planar Vertex Cover instance.

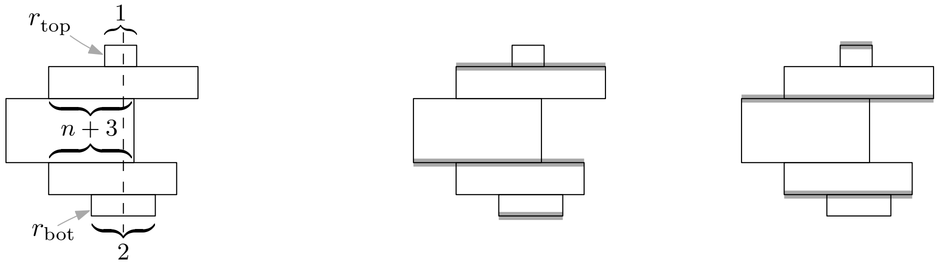

(a) All rectangles of R_v are intersected by the (dashed) vertex segment of v. **(b)** R_v stabbed by S^v_{act} **(c)** R_v stabbed by S^v_{ina}

Figure 4 The vertex gadget R_v of vertex v.

4 NP-Hardness of Stabbing

To show that Stabbing is NP-hard, we reduce from Planar Vertex Cover: Given a planar graph G and an integer k, decide whether G has a vertex cover of size at most k. This problem is NP-hard [11]. Omitted proofs can be found in the full version of the paper.

▶ **Theorem 9.** Stabbing *is* NP*-hard, even for interior-disjoint rectangles.*

Let $G = (V, E)$ be a planar graph with n vertices, and let k be a positive integer. Our reduction will map G to a set R of rectangles and k to another integer $k^\star$ such that (G, k) is a yes-instance of Planar Vertex Cover if and only if $(R, k^\star)$ is a yes-instance of Stabbing. Consider a *visibility representation* of G, which represents the vertices of G by non-overlapping vertical line segments (called *vertex segments*), and each edge of G by a horizontal line segment (called *edge segment*) that touches the vertex segments of its endpoints; see Figs. 3a and 3b. Any planar graph admits a visibility representation on a grid of size $O(n) \times O(n)$, which can be found in polynomial time [15]. We compute such a visibility representation for G. Then we stretch the vertex segments and vertically shift the edge segments so that no two edge segments coincide (on a vertex segment); see Fig. 3c. The height of the visibility representation remains linear in n.

In the next step, we create a Stabbing instance based on this visibility representation, using the edge segments and vertex segments as indication for where to put our rectangles. All rectangles will be interior-disjoint, have positive area and lie on an integer grid that we obtain by scaling the visibility representation by a sufficiently large factor (linear in n). A vertex segment will intersect $O(n)$ rectangles (lying above each other since they are disjoint), and each rectangle will have width $O(n)$. The precise number of rectangles and their sizes will depend on the constraints formulated below. Our construction will be polynomial in n.

For each edge e in G, we introduce an *edge gadget* r_e, which is a rectangle that we place such that it is stabbed by the edge segment of e in the visibility representation.

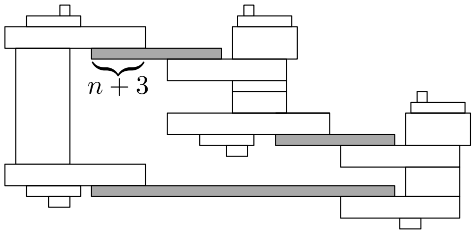

Figure 5 The Stabbing instance that encodes the Planar Vertex Cover instance of Fig. 3; edge gadgets are shaded gray.

For each vertex v in G, we introduce a *vertex gadget* R_v as shown in Fig. 4a. It consists of an odd number of rectangles that are (vertically) stabbed by the vertex segment of v in the visibility representation. Any two neighboring rectangles share a horizontal line segment. Its length is exactly $n+3$ if neither of the rectangles is the top-most rectangle r_{top} or the bottom-most rectangle r_{bot}. Otherwise, the intersection length equals the width of the respective rectangle r_{top} or r_{bot}. We set the widths of r_{top} and r_{bot} to 1 and 2, respectively. A vertex gadget R_v is called *incident* to an edge gadget r_e if v is incident to e.

Before we describe the gadgets and their relation to each other in more detail, we construct, in two steps, a set S^v of line segments for each vertex gadget R_v. First, let S^v be the set of line segments that correspond to the top and bottom edges of the rectangles in R_v. Second, replace each pair of overlapping line segments in S^v by its union. Then number the line segments in S^v from top to bottom starting with 1. Let S^v_{ina} be the set of the odd-numbered line segments, and let S^v_{act} be the set of the even-numbered ones; see Figs. 4b and 4c. By construction, S^v_{act} and S^v_{ina} are feasible stabbings for R_v. Furthermore, $|S^v_{\text{ina}}| = |S^v_{\text{act}}|$ as $|R_v|$ is odd and, hence, $|S^v|$ is even. Given the difference in the widths of r_{top} and r_{bot}, we have that $\|S^v_{\text{act}}\| = \|S^v_{\text{ina}}\| + 1$. Note that this equation holds regardless of the widths of the rectangles in $R_v \setminus \{r_{\text{top}}, r_{\text{bot}}\}$.

The rectangles of all gadgets together form a Stabbing instance R. They meet two further constraints: First, no two rectangles of different vertex gadgets intersect. We can achieve this by scaling the visibility representation by an appropriate factor linear in n. Second, each edge gadget r_e intersects exactly two rectangles, one of its incident left vertex gadgets, R_v, and one of its incident right vertex gadgets, R_u. The top edge of r_e touches a segment of S^v_{act} and the bottom edge of r_e touches a segment of S^u_{act}. The length of each of the two intersections is exactly $n+3$; see Fig. 5. Thus, we have $|R_v| = O(\deg(v)) = O(n)$.

Let S be a feasible solution to the instance R. We call a vertex gadget R_v *active* in S if $\{s \cap \bigcup R_v \mid s \in S\} = S^v_{\text{act}}$, and *inactive* in S if $\{s \cap \bigcup R_v \mid s \in S\} = S^v_{\text{ina}}$. We will see that in any optimum solution each vertex gadget is either active or inactive. Furthermore, we will establish a direct correspondence between the Planar Vertex Cover instance G and the Stabbing instance R: Every optimum solution to R *covers* each edge gadget by an active vertex gadget while minimizing the number of active vertex gadgets.

Let OPT_G denote the size of a minimum vertex cover for G, let OPT_R denote the length of an optimum solution to R, let $\text{width}(r)$ denote the width of a rectangle r, and finally let $c = \sum_{e \in E}(\text{width}(r_e) - n - 3) + \sum_{v \in V} \|S^v_{\text{ina}}\|$. To show NP-hardness of Stabbing, we prove that $\text{OPT}_G \leq k$ if and only if $\text{OPT}_R \leq c + k$. We show the two directions separately.

▶ **Lemma 10.** $\text{OPT}_G \leq k$ *implies that* $\text{OPT}_R \leq c + k$.

Proof sketch. Set each vertex gadget to active if it corresponds to a vertex in the given vertex cover, otherwise to inactive. Stab each edge gadget by prolonging one of the line segments that it touches. Using $\|S^v_{\text{act}}\| = \|S^v_{\text{ina}}\| + 1$, the bound follows. ◀

Next we show the other, more challenging direction. Consider an optimum solution S_{OPT} to R and choose $k \le n$ such that $\mathrm{OPT}_R \le c+k$ is satisfied. Let R_v be any vertex gadget, let r_{top} and r_{bot} be its top- and bottom-most rectangles, respectively, and let $S^v_{\mathrm{OPT}} = \{s \cap \bigcup R_v \mid s \in S_{\mathrm{OPT}}\}$. In the following, we prove that S^v_{OPT} equals either S^v_{ina} or S^v_{act}.

▶ **Lemma 11.** *If $S^v_{\mathrm{ina}} \not\subseteq S^v_{\mathrm{OPT}}$ and $S^v_{\mathrm{act}} \not\subseteq S^v_{\mathrm{OPT}}$, then $\|S^v_{\mathrm{OPT}}\| > \|S^v_{\mathrm{act}}\| + n$.*

Proof sketch. Consider all pairs of neighboring rectangles in R_v that are stabbed by the same line segment of S^v_{OPT}. Let P be a maximum-cardinality subset of these pairs such that every rectangle appears at most once. Thus, $\sum_{r \in R_v} \mathrm{width}(r) - \sum_{(r_1,r_2)\in P} \mathrm{width}(r_1 \cap r_2)$ is a lower bound of $\|S^v_{\mathrm{OPT}}\|$. Observe that the lower bound is minimized if the total intersection length of the rectangles in P is maximized. This happens (even with tightness) if and only if $S^v_{\mathrm{OPT}} = S^v_{\mathrm{ina}}$. Given that $|R_v|$ is odd, there is at least one rectangle not in P. If $S^v_{\mathrm{ina}} \not\subseteq S^v_{\mathrm{OPT}}$ and $S^v_{\mathrm{act}} \not\subseteq S^v_{\mathrm{OPT}}$, there is a rectangle r not in P that is neither r_{top}, r_{bot} nor a neighbor of those. Thus, r contributes $n+3$ to the total intersection length in S^v_{ina} but nothing in S^v_{OPT}. The difference of the total intersection lengths implies the lemma. ◀

▶ **Lemma 12.** *Exactly one of the following three statements holds:*

(i) $S^v_{\mathrm{OPT}} = S^v_{\mathrm{ina}}$, *or*
(ii) $S^v_{\mathrm{OPT}} = S^v_{\mathrm{act}}$, *or*
(iii) $\|S^v_{\mathrm{OPT}}\| > \|S^v_{\mathrm{ina}}\| + n$.

Proof sketch. If $S^v_{\mathrm{ina}} \subsetneq S^v_{\mathrm{OPT}}$, there is a line segment $s \in S^v_{\mathrm{OPT}} \setminus S^v_{\mathrm{ina}}$ that stabs a rectangle in $R_v \setminus \{r_{\mathrm{top}}, r_{\mathrm{bot}}\}$. By construction, its length is at least $n+3$. Hence, $\|S^v_{\mathrm{OPT}}\| > \|S^v_{\mathrm{ina}}\| + n$. The same holds if $S^v_{\mathrm{act}} \subsetneq S^v_{\mathrm{OPT}}$. ◀

Now, we show that S_{OPT} forces each vertex gadget to be either active or inactive.

▶ **Lemma 13.** *In S_{OPT}, each vertex gadget is either active or inactive.*

Proof. Suppose that there is a vertex gadget R_u that is neither active nor inactive in S_{OPT}. This implies $\mathrm{OPT}_R > c+n$ and contradicts our previous assumption $\mathrm{OPT}_R \le c+k \le c+n$.

To this end, we give a lower bound on OPT_R. Since R_u is neither active nor inactive, $S^u_{\mathrm{OPT}} > \|S^u_{\mathrm{ina}}\| + n$ by Lemma 12. Thus, $\sum_{v\in V} \|S^v_{\mathrm{OPT}}\| > \sum_{v \in V} \|S^v_{\mathrm{ina}}\| + n$. Let $S^{\mathrm{out}}_{\mathrm{OPT}}$ be the set of all segment fragments of S_{OPT} lying outside of $\bigcup_{v\in V} S^v_{\mathrm{OPT}}$. Each edge gadget r_v contains a segment fragment from $S^{\mathrm{out}}_{\mathrm{OPT}}$ of length at least $\mathrm{width}(r_v) - n - 3$ since, by construction, it can share a line segment with only one of its incident vertex gadgets. Since all edge gadgets are interior-disjoint, we have $\|S^{\mathrm{out}}_{\mathrm{OPT}}\| \ge \sum_{e\in E} \mathrm{width}(r_v) - n - 3$. Hence,

$$\mathrm{OPT}_R \;\ge\; \|S^{\mathrm{out}}_{\mathrm{OPT}}\| + \sum_{v\in V} \|S^v_{\mathrm{OPT}}\| > \sum_{e\in E} (\mathrm{width}(r_e) - n - 3) + \sum_{v \in V} \|S^v_{\mathrm{ina}}\| + n \;=\; c+n\,. \blacktriangleleft$$

▶ **Lemma 14.** *For each edge gadget, one of its incident vertex gadgets is active in S_{OPT}.*

Proof. Suppose that for an edge gadget r_e both vertex gadgets are not active in S_{OPT}. By Lemma 13, they are inactive. Without loss of generality, the line segment s stabbing r_e lies on the top or bottom edge of r_e. Then s intersects a vertex gadget to the left or right, say R_v, and hence $S^v_{\mathrm{OPT}} \ne S^v_{\mathrm{ina}}$ and $S^v_{\mathrm{OPT}} \ne S^v_{\mathrm{act}}$. A contradiction to Lemma 13. ◀

▶ **Lemma 15.** $\mathrm{OPT}_R = c + k'$, *where k' is the number of active vertex gadgets in S_{OPT}.*

Proof sketch. Every edge gadget r_e is stabbed by a line segment s that also stabs a rectangle r of an incident active vertex gadget R_v. Hence, $\|s\| = \mathrm{width}(r) + \mathrm{width}(r_e) - n - 3$. By $\|S^v_{\mathrm{act}}\| = \|S^v_{\mathrm{ina}}\| + 1$, $\mathrm{OPT}_R = \sum_{e\in E} (\mathrm{width}(r_e) - n - 3) + \sum_{v\in V} \|S^v_{\mathrm{ina}}\| + k' = c + k'$. ◀

Given S_{OPT}, we put exactly those vertices in the vertex cover whose vertex gadgets are active. By Lemma 14, this yields a vertex cover of G. By Lemma 15, the size of the vertex cover is exactly $\mathrm{OPT}_R - c$, which is bounded from above by k given that $\mathrm{OPT}_R \leq c + k$.

▶ **Lemma 16.** $\mathrm{OPT}_R \leq c + k$ *implies that* $\mathrm{OPT}_G \leq k$.

By our construction, we represent R on a grid of size polynomial in n, hence, all numerical values are upperbounded by a polynomial in n. Our construction is polynomial. With Lemmas 10 and 16, we conclude that Stabbing is NP-hard.

References

1 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.

2 Nikhil Bansal and Kirk Pruhs. The Geometry of Scheduling. *SIAM J. Computing*, 43(5):1684–1698, 2014. doi:10.1137/130911317.

3 Hervé Brönnimann and Michael T. Goodrich. Almost Optimal Set Covers in Finite VC-Dimension. *Discrete Comput. Geom.*, 14(4):463–479, 1995. doi:10.1007/BF02570718.

4 Timothy M. Chan and Elyot Grant. Exact Algorithms and APX-hardness Results for Geometric Packing and Covering Problems. *Comput. Geom. Theory Appl.*, 47(2):112–124, 2014. doi:10.1016/j.comgeo.2012.04.001.

5 Timothy M. Chan, Elyot Grant, Jochen Könemann, and Malcolm Sharpe. Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling. In *Proc. 23th Annu. ACM-SIAM Symp. Discrete Algorithms (SODA'12)*, pages 1576–1585, 2012. URL: http://portal.acm.org/citation.cfm?id=2095241&CFID=63838676&CFTOKEN=79617016.

6 Aparna Das, Krzysztof Fleszar, Stephen G. Kobourov, Joachim Spoerhase, Sankar Veeramoni, and Alexander Wolff. Approximating the Generalized Minimum Manhattan Network Problem. *Algorithmica*, 80(4):1170–1190, 2018. doi:10.1007/s00453-017-0298-0.

7 Irit Dinur and David Steurer. Analytical approach to parallel repetition. In *Proc. Symp. Theory Comput. (STOC'14)*, pages 624–633, 2014. doi:10.1145/2591796.2591884.

8 Guy Even, Retsef Levi, Dror Rawitz, Baruch Schieber, Shimon Shahar, and Maxim Sviridenko. Algorithms for capacitated rectangle stabbing and lot sizing with joint set-up costs. *ACM Trans. Algorithms*, 4(3), 2008. doi:10.1145/1367064.1367074.

9 Uriel Feige. A Threshold of $\ln n$ for Approximating Set Cover. *J. ACM*, 45(4):634–652, 1998. doi:10.1145/285055.285059.

10 Gerd Finke, Vincent Jost, Maurice Queyranne, and András Sebö. Batch processing with interval graph compatibilities between tasks. *Discrete Appl. Math.*, 156(5):556–568, 2008. doi:10.1016/j.dam.2006.03.039.

11 Michael R. Garey, David S. Johnson, and Larry Stockmeyer. Some simplified NP-complete problems. In *Proc. 6th Annu. ACM Symp. Theory Comput. (STOC'74)*, pages 47–63, 1974. doi:10.1145/800119.803884.

12 Daya Ram Gaur, Toshihide Ibaraki, and Ramesh Krishnamurti. Constant Ratio Approximation Algorithms for the Rectangle Stabbing Problem and the Rectilinear Partitioning Problem. *J. Algorithms*, 43(1):138–152, 2002. doi:10.1006/jagm.2002.1221.

13 Panos Giannopoulos, Christian Knauer, Günter Rote, and Daniel Werner. Fixed-parameter tractability and lower bounds for stabbing problems. *Comput. Geom. Theory Appl.*, 46(7):839–860, 2013. doi:10.1016/j.comgeo.2011.06.005.

14 Sofia Kovaleva and Frits C. R. Spieksma. Approximation Algorithms for Rectangle Stabbing and Interval Stabbing Problems. *SIAM J. Discrete Math.*, 20(3):748–768, 2006. doi:10.1137/S089548010444273X.

15 Ching-Chi Lin, Hsueh-I Lu, and I-Fan Sun. Improved Compact Visibility Representation of Planar Graph via Schnyder's Realizer. *SIAM J. Discrete Math.*, 18(1):19–29, 2004. `doi:10.1137/S0895480103420744`.

16 Nabil H. Mustafa, Rajiv Raman, and Saurabh Ray. Quasi-Polynomial Time Approximation Scheme for Weighted Geometric Set Cover on Pseudodisks and Halfspaces. *SIAM J. Computing*, 44(6):1650–1669, 2015. `doi:10.1137/14099317X`.

17 Nabil H. Mustafa and Saurabh Ray. Improved Results on Geometric Hitting Set Problems. *Discrete Comput. Geom.*, 44(4):883–895, 2010. `doi:10.1007/s00454-010-9285-9`.

18 Kasturi Varadarajan. Weighted geometric set cover via quasi-uniform sampling. In *Proc. 42nd ACM Symp. Theory Comput. (STOC'10)*, pages 641–648, 2010. `doi:10.1145/1806689.1806777`.