

Planar Support for Non-piercing Regions and Applications

Rajiv Raman

IIIT-Delhi, Delhi, India

rajiv@iiitd.ac.in

Saurabh Ray

Department of Computer Science, NYU Abu Dhabi, United Arab Emirates

saurabh.ray@nyu.edu

Abstract

Given a hypergraph $\mathcal{H} = (X, \mathcal{S})$, a planar support for \mathcal{H} is a planar graph G with vertex set X , such that for each hyperedge $S \in \mathcal{S}$, the sub-graph of G induced by the vertices in S is connected. Planar supports for hypergraphs have found several algorithmic applications, including several packing and covering problems, hypergraph coloring, and in hypergraph visualization.

The main result proved in this paper is the following: given two families of regions R and B in the plane, each of which consists of connected, non-piercing regions, the *intersection hypergraph* $\mathcal{H}_R(B) = (B, \{B_r\}_{r \in R})$, where $B_r = \{b \in B : b \cap r \neq \emptyset\}$ has a planar support. Further, such a planar support can be computed in time polynomial in $|R|, |B|$, and the number of vertices in the arrangement of the regions in $R \cup B$. Special cases of this result include the setting where either the family R , or the family B is a set of points.

Our result unifies and generalizes several previous results on planar supports, PTASs for packing and covering problems on non-piercing regions in the plane and coloring of intersection hypergraph of non-piercing regions.

2012 ACM Subject Classification Mathematics of computing \rightarrow Approximation algorithms

Keywords and phrases Geometric optimization, packing and covering, non-piercing regions

Digital Object Identifier 10.4230/LIPIcs.ESA.2018.69

Acknowledgements Part of the work was done when the first author was visiting NYU Abu Dhabi, and the author thanks the institution for its hospitality.

1 Introduction

Hypergraphs arise naturally in several applications and it is often helpful to capture the structure of the hypergraph using a *sparse graph*. One popular way to capture the structure of a hypergraph is to construct a *planar support*, which is a planar graph G with the same vertex set as the hypergraph such that every hyperedge induces a connected subgraph of G .

In this paper, we study hypergraphs that arise in several, primarily geometric settings. For example, a set of points P , and family of disks D in the plane define a hypergraph, $\mathcal{H}(P, D)$, where each disk $d \in D$ defines a hyperedge $P \cap d$. This is a widely studied hypergraph, and it is well known that the Delaunay graph of the points is a planar support for this hypergraph. Another hypergraph on the same objects is $\mathcal{H}(D, P)$, where the vertices are the disks, and each point $p \in P$ defines a hyperedge $\{d \in D : d \ni p\}$. We refer to $\mathcal{H}(P, D)$ as the *primal hypergraph*, and $\mathcal{H}(D, P)$, as the *dual hypergraph*. A hypergraph that generalizes both these hypergraphs is the following: Given a family R of *red* disks, and a family B of *blue* disks, we define the *intersection hypergraph* $\mathcal{H}(B, R)$ as follows: the disks B are the vertices, and



© Rajiv Raman and Saurabh Ray;
licensed under Creative Commons License CC-BY
26th Annual European Symposium on Algorithms (ESA 2018).

Editors: Yossi Azar, Hannah Bast, and Grzegorz Herman; Article No. 69; pp. 69:1–69:14



Leibniz International Proceedings in Informatics

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

each disk $r \in R$ defines a hyperedge $B_r = \{b \in B : r \cap b \neq \emptyset\}$. Both the primal and dual hypergraphs can be seen as special cases of this intersection hypergraph, by viewing the points in the arrangement as either blue disks, or red disks of radius zero. Our result implies that this intersection hypergraph has a planar support.

If instead of disks, we use topological generalizations like *pseudodisks* or *k-admissible regions*, the primal hypergraph defined above still admits a planar support [13]. However, it was not known if the same is true for the dual hypergraph. Our result implies that the intersection hypergraphs of *k-admissible regions* (which includes pseudodisks) admits a planar support which in particular means that the dual hypergraph has a planar support.

We go beyond *k-admissible regions*, and show that these results are true even if the regions have holes (i.e., they need not be simply connected¹). Previous proofs relied heavily on the assumption that the regions are simply connected. In addition, our results are constructive and give polynomial time algorithms. Thus, even for the primal hypergraph of *k-admissible regions*, where a planar support was known to exist, we make progress by giving a polynomial time construction.

The existence of a planar support for the primal hypergraph for *k-admissible regions* has a surprising algorithmic consequence: it implies that local search yields a PTAS for the hitting set problem (the set cover problem for the dual hypergraph) [12] for *k-admissible regions*. Since a planar support for the dual hypergraph for *k-admissible regions* was not known, the authors in [1] had to construct another suitable graph in order to prove that local search yields a PTAS for the set cover problem for the primal hypergraph. Similarly, for the Dominating Set problem for *k-admissible regions*, an alternate graph construction was required in [1], which was a generalization of a previous construction of [7] for disks. We unify these results by considering a problem that generalizes all three problems: hitting set, set cover, and dominating set for which our result implies a PTAS. Our result is in fact stronger and is not implied by the previous results.

Prior to our result, Chan and Har-Peled [4] proved that an arrangement of *k-admissible regions* of depth two² admits a planar support. In fact, in this case, the intersection graph³ of the regions is itself the planar support. They used this result to obtain a PTAS for the independent set problem for *k-admissible regions* [4]. Our result also implies PTASs for generalized versions of packing problems considered in [1] and [6]. For some of the problems, we obtain a PTAS where only a constant factor approximation was known.

In a different line of work, Keszegh [10], recently proved the following interesting result: the intersection hypergraph of two families of pseudodisks is four colorable. This result follows immediately from our result since planar graphs are four colorable. Keszegh's paper has several other results but the above result is the central tool to which most of the paper is devoted and from which the other results follow.

The result of Keszegh extends the result of Keller and Smorodinsky [9] where the hypergraph is defined by a single family of pseudodisks, the vertex set is the set of pseudodisks and the hyperedges consists of open or closed neighborhood of each pseudodisk in the intersection graph.

Before describing our results and other related results in more detail, we introduce the necessary definitions and notation.

¹ A region is simply connected if any loop can be continuously shrunk to a point while staying within the region, which is true for a disk, but not for an annulus.

² i.e., no more than two intersect at any point in the plane

³ An intersection graph on a set of regions is a graph whose vertices are the regions, and two vertices are adjacent if their corresponding regions intersect.

Definitions and Notation

We will use the term *region* to refer to a set γ in the plane that can be described as $\gamma = \bar{\gamma} \setminus \text{int}(H_1 \cup \dots \cup H_k)$, where H_1, \dots, H_k are disjoint, compact, simply connected regions contained in the compact, simply connected region $\bar{\gamma}$. Essentially, γ is a compact connected region with holes. We will refer to the region $\bar{\gamma}$ as the *filled region* corresponding to γ , and the regions H_1, \dots, H_k as the holes in γ . We will refer to the boundary of $\bar{\gamma}$ as the *outer boundary* of γ .

We say that a set of regions are in *general position* if a) the boundaries of any two of the regions intersect a finite number of times, and cross at these points, b) the boundaries of three (or more) of the regions do not intersect at a common point, and c) any region can be expanded by a small non-zero amount without changing the arrangement of the regions combinatorially (see Definition 8 in Section 2 for a formal definition of expansion of a region). Furthermore, we say that a set of points and a set of regions are together in general position if none of the points lie on the boundary of any of the regions. In the rest of the paper, we will always assume general position. Two regions γ_1 and γ_2 in the plane are said to be *non-piercing* if the sets $\gamma_1 \setminus \gamma_2$, and $\gamma_2 \setminus \gamma_1$ are both connected. A family Γ of regions is said to be non-piercing if the regions in Γ are pairwise non-piercing.

► **Remark.** The term *non-piercing* has been used previously to refer to a family of regions defined exactly as we have except that the regions are required to be simply connected (i.e., not containing holes). The term *k-admissible* refers to such non-piercing families where in addition the boundaries of each pair of regions intersect at most k times. The term *pseudodisks* is used to refer to a 2-admissible family of regions. To be more consistent with the literature, we should be using a term like “non-piercing regions with holes”. However, for better readability we stick to using a shorter term.

Given a set Γ of non-piercing regions, and a set P of points in the plane, we define the *primal* hypergraph $\mathcal{H}(P, \Gamma)$ as the hypergraph in which the vertex set is P , and there is a hyperedge $\gamma \cap P$ corresponding to each $\gamma \in \Gamma$. We define the *dual* hypergraph $\mathcal{H}(\Gamma, P)$ as the hypergraph in which the vertex set is Γ , and corresponding to each $p \in P$, there is a hyperedge $\{\gamma \in \Gamma : \gamma \ni p\}$. Finally, given two families of non-piercing regions R and B , we define their *intersection hypergraph* $\mathcal{H}(B, R)$ as the hypergraph in which the vertex set is B , and corresponding to each region $r \in R$ there is a hyperedge $B_r = \{b \in B : r \cap b \neq \emptyset\}$.

1.1 Our Results and Implications

The main result we prove in this paper is the following:

► **Theorem 1.** *Given two families R and B of non-piercing regions, the intersection hypergraph $\mathcal{H}(B, R)$ admits a planar support.*

We now describe the implications of Theorem 1. Due to shortage of space in this extended abstract, we do not prove the claimed consequences of Theorem 1, as they follow in a straightforward manner from standard arguments.

Generalized Set Cover Problem for non-piercing regions

Given a family R of red non-piercing regions, and another family B of blue non-piercing regions, such that each $r \in R$ is intersected by at least one $b \in B$, find the smallest subset $B' \subseteq B$ such that each $r \in R$ is intersected by at least one $b \in B'$.

The following theorem below follows in a straightforward manner from the framework in [12] using Theorem 1.

► **Theorem 2.** *There is a PTAS for the Generalized Set Cover problem for non-piercing regions.*

When B is a set of points in the plane, this problem is equivalent to the *hitting set* problem for non-piercing regions: given a set of points and a family of non-piercing regions, find the smallest subset of points such that each region contains at least one point from our chosen subset of points. When R is a set of points in the plane, this problem is equivalent to the *set cover* problem for non-piercing regions in the plane: given a set of points and non-piercing regions in the plane, find the smallest subset of the regions so that each input point is covered by one of the chosen subset of regions. When both B and R are identical families of non-piercing regions, the problem is equivalent to the *dominating set* problem for non-piercing regions: given a family of non-piercing regions, find the smallest subset of regions so that each of the other regions intersects at least one of the chosen regions. Thus a PTAS for the generalized set cover problem for non-piercing regions implies a PTAS for all three problems: hitting set, set cover and dominating set. However, the reverse is not true: the PTASs for these three problems together do not imply a PTAS for the generalized set cover problem.

A PTAS for the hitting set problem for simply connected non-piercing regions in the plane, and halfspaces in \mathbb{R}^3 was given in [12]. For the set cover problem for disks in the plane a PTAS follows from the result for halfspaces, via a standard lifting to three dimensions. For the dominating set problem for disks in the plane, a PTAS was given in [7]. Generalizations of the PTASs for the set cover and dominating set problems for disks to simply connected non-piercing regions in the plane were given in [1]. None of the earlier results work in the setting where the regions are allowed to have holes.

Weighted Covering problems

In the weighted variant of the generalized set cover problem, each region $b \in B$ has a non-negative weight, and the goal is to minimize the total weight of the chosen set $B' \subseteq B$. Chan et. al. [3], building on the work of Varadarajan [14], obtained constant-factor approximation algorithms for set systems with linear *shallow-cell complexity*⁴. Note that the number of hyperedges of size 2 in $\mathcal{H}(B, R)$ is $O(|B|)$ since each such hyperedge corresponds to an edge in the planar support which cannot have more than $3|B| - 6$ edges. Since this is true for the projection of the hypergraph on any subset of B , a standard probabilistic argument due to Clarkson and Shor [5] implies that the number of hyperedges in $\mathcal{H}(B, R)$ of size at most k is $O(kn)$ implying that the shallow cell complexity of $\mathcal{H}(B, R)$ is linear. Our result thus implies a constant factor approximation for the weighted variant of the generalized set cover problem via the framework of Chan et. al. [3].

► **Theorem 3.** *There is an $O(1)$ -approximation algorithm for the weighted Generalized Set Cover problem for non-piercing regions.*

Generalized Set Packing problem for non-piercing regions

Given a set R of red non-piercing regions, and a set B of blue non-piercing regions where each red region has a capacity bounded above by a constant $C > 0$, find a maximum cardinality subset $B' \subseteq B$, such that the number of blue regions in B' intersecting any $r \in R$ does not exceed the capacity of the region r .

⁴ See [3] for the definition of Shallow Cell Complexity.

The theorem below follows in a straightforward manner from the framework in [1] using Theorem 1.

► **Theorem 4.** *There is a PTAS for the generalized set packing problem for non-piercing regions, when each region has a capacity bounded above by a constant.*

The generalized set packing problem specializes to the *region packing* problem when the set R is a set of points: Given a family B of non-piercing regions, and a set R of points, each with a capacity bounded above by a constant C find a maximum subset $B' \subseteq B$, such that no point $r \in R$ is covered by more than its capacity. When $C = 1$, note that the region packing problem is the discrete independent set problem which itself generalizes the independent set problem in the intersection graph of the regions in B . In [4], Chan and Har-Peled gave a PTAS for the independent set problem for a set of simply connected, non-piercing regions in the plane. When the set B is a set of points, the generalized set packing problem specializes to the *point packing problem*: Given a set of points B , and a family R of non-piercing regions, each with capacity upper bounded by a constant C , find a maximum cardinality subset of points $B' \subseteq B$, such that the number of points of B' in any $r \in R$ is at most the capacity r .


In [1], a PTAS was given for the region packing problem, when the regions were assumed to be simply connected non-piercing regions where the boundaries of each pair of regions intersect at most a constant number of times. For the point packing problem, again [1] gave a PTAS when $C = 1$. For larger values of C , only constant-factor approximation algorithms were known [6]. Our result thus extends the PTAS in [1] for any constant C . Furthermore, our PTAS for all the above problems works for non-piercing regions with holes, which generalizes the earlier results, but does not follow from them.

► **Remark.** All the PTASs mentioned above follow a local search framework, which requires the construction of a suitable graph for analysis. In earlier work, the graph construction for all the problems above relied on the fact that the regions were simply connected, and non-piercing, and did not extend to regions with holes. As far as we are aware, the above problems were not studied for non-piercing regions with holes. Further, each problem required a different graph construction. It is satisfying to finally have a unified view of all these problems.

Hypergraph Coloring

Recently, Keszegh [10], proved the following result which generalizes the result of Keller and Smorodinsky [9]: Let R and B be two families of pseudodisks. Then, the intersection hypergraph $\mathcal{H}(B, R)$ admits a coloring with 4 colors, and a *conflict-free* coloring with $O(\log n)$ colors. Keszegh's result follows from our result due to the fact that the planar support of $\mathcal{H}(B, R)$ is four colorable, and a valid coloring of the planar support is a valid coloring of the hypergraph. Our result thus extends Keszegh's result to non-piercing regions. In order to prove his result, Keszegh proves that the *Delaunay graph*, $G = (B, E)$ where the vertex set is B and E is the set of hyperedges of size 2 in $\mathcal{H}(B, R)$, is planar. Observe that our result is stronger since every edge in the Delaunay graph must be in the planar support.

► **Theorem 5.** *Given two families R and B of non-piercing regions, the intersection hypergraph $\mathcal{H}(B, R)$ can be colored with at most 4 colors.*



■ **Figure 1** The assumption that there is an $\alpha, \beta, \alpha, \beta$ subsequence in the sequence of labels along ∂C leads to a contradiction.

2 Cell Bypassing

In this section, we define *cell bypassing*⁵, a basic operation that is used to simplify a given arrangement of non-piercing regions. Let Γ denote a set of non-piercing regions, and let \mathcal{A} denote the arrangement of these regions. We will show that whenever there is a point contained in at least 3 regions, we can simplify the arrangement by modifying a region $\gamma \in \Gamma$, while maintaining key properties required to construct a planar support.

For a region γ , we let $\partial\gamma$ denote the boundary of γ , and $\text{int}(\gamma)$ to denote the interior of γ . For a family of regions Γ , define $\partial\Gamma = \bigcup_{\gamma \in \Gamma} \partial\gamma$ where $\partial\gamma$ denotes the boundary of the region γ . The closure of each connected component of $\mathbb{R}^2 \setminus \partial\Gamma$ defines a *cell*. For a cell C , we define the following: Γ_C denotes the set of regions containing the cell C , $\text{depth}(C)$ denotes $|\Gamma_C|$, and ∂C denotes the boundary of C . If an arc of $\partial\gamma$ lies on ∂C , then the region γ is said to *contribute* to ∂C . We define $\text{degree}(C)$ as the number of arcs on the boundary of C . Note that the same region may contribute multiple arcs to ∂C .

We say that two cells are adjacent if their boundaries share an arc of positive length. We define the *cell adjacency graph*⁶ G_Γ of Γ as the graph in which vertices correspond to the cells and two vertices are adjacent in the graph if the corresponding cells are adjacent. Clearly, G_Γ is a planar graph. Observe that the degree of a cell in G_Γ is equal to the number of arcs on its boundary. Also note that the depth of adjacent cells C and C' in an arrangement of regions differ by exactly 1. If $\text{depth}(C) < \text{depth}(C')$, then $\Gamma_{C'} = \Gamma_C \cup \{\rho\}$, where ρ is the unique region such that the arc $C \cap C' \subseteq \partial\rho$. We say that a cell C is *maximal* if its depth is more than the depth of any cell C' adjacent to it.

► **Lemma 6.** *All regions in Γ contributing to the boundary of a maximal cell C in \mathcal{A} contain the cell C .*

Proof. For contradiction, assume that there is an arc α of $\partial\gamma$ that appears on ∂C , but $C \not\subseteq \gamma$. Let C' be the cell adjacent to C along α . Then, $\text{depth}(C') > \text{depth}(C)$ as all regions containing C also contain C' , and C' is additionally contained in γ . This contradicts the maximality of C . ◀

⁵ The notion of cell bypassing is similar in spirit to the notion of *lens bypassing* in [1] and to the notion of *core decomposition* in [11]. However, the technical difference is critical for the applications in this paper.

⁶ The cell adjacency graph is the geometric dual of the arrangement graph in which the intersection points of the boundaries of the regions are the vertices and two vertices are adjacent in the graph if they appear consecutively along the boundary of some region.

► **Lemma 7.** *Let C be a maximal simply connected cell in \mathcal{A} , whose boundary arcs are labelled by the regions in Γ contributing them. Let σ be the cyclic sequence of the labels of the arcs in counterclockwise order along ∂C . Then, σ is a Davenport Schinzel sequence of order 2 i.e., it does not contain a subsequence of the form $\alpha, \beta, \alpha, \beta$.*

Proof. For contradiction, let a_1, b_1, a_2, b_2 be four arcs appearing in cyclic order along ∂C , where a_1 and a_2 have the label α , and b_1 and b_2 have the label β . Since C is a maximal cell, note that $C \subseteq \alpha \cap \beta$. Let p_1 and p_2 be points in the interior of the arcs a_1 and a_2 respectively. Similarly, let q_1 and q_2 be points in the interior of the arcs b_1 and b_2 respectively. See Figure 1. Since q_1 and q_2 lie on the boundary of $\alpha \setminus \beta$ which by assumption is connected, there is a curve τ joining q_1 and q_2 whose interior lies in $\alpha \setminus \beta$. Similarly, there is a curve τ' joining p_1 and p_2 whose interior lies in $\beta \setminus \alpha$. Note that the interiors of neither τ , nor τ' intersect C , since C does not intersect either $\alpha \setminus \beta$ or $\beta \setminus \alpha$. Since p_1, q_1, p_2, q_2 appear along ∂C in that order, and since C does not have any holes (i.e., it is simply connected), τ and τ' must intersect at a point outside C and in the interior of both the curves. This is a contradiction since $\alpha \setminus \beta$ and $\beta \setminus \alpha$ are disjoint sets. ◀

► **Definition 8** (ϵ -expansion). For any region R , define an ϵ -expansion of R , denoted R_ϵ as the Minkowski sum of R and a ball of an arbitrarily small radius ϵ centered at the origin.

► **Remark.** The ϵ -expansion of a region is necessary for technical reasons. The parameter ϵ is always chosen to be small enough so that combinatorial structure of the arrangement does not change if a region R is replaced by R_ϵ in the family Γ . Due to general position assumptions, such an ϵ always exists. The choice of ϵ thus depends on the other regions in the family Γ , but we suppress this dependency for better readability.

► **Definition 9** (Good region). For a maximal cell C in \mathcal{A} , a region $\gamma \in \Gamma$ is said to be a *good* region for C if the following conditions hold: i) γ contributes to the boundary of C and ii) $(\Gamma \setminus \{\gamma\}) \cup \{\gamma'\}$, where $\gamma' = \gamma \setminus \text{int}(C_\epsilon)$, is a non-piercing family.

► **Lemma 10.** *For any maximal simply connected cell C in \mathcal{A} , there is a region $\gamma \in \Gamma$ that contributes exactly one arc to the boundary of C . Furthermore, any such region is good for C .*

Proof. First we argue that there is a region that contributes exactly one arc to ∂C . If every arc on ∂C has a distinct label, we are done. Otherwise, consider two arcs a_1 and a_2 having the same label α that are closest to each other in counterclockwise order along ∂C . Let b be any arc lying between a_1 and a_2 . There is such an arc since consecutive arcs along ∂C cannot have the same label. Let γ be the label of the arc b . By the choice of a_1 and a_2 and Lemma 7 there cannot be any other arc on ∂C with label γ . Thus, there is at least one region γ that contributes exactly one arc to ∂C .

We now show that any such region γ is a good region for C . The fact that γ contributes exactly one arc to ∂C , along with the fact that C is simply connected implies that $\gamma' = \gamma \setminus \text{int}(C_\epsilon)$ is connected. For any other region $\nu \in \Gamma$, we now argue that both $\nu \setminus \gamma'$ and $\gamma' \setminus \nu$ are connected. Suppose first that ν does not contain C . Then $\nu \setminus \gamma' = \nu \setminus \gamma$ which is connected. Also $\gamma' \setminus \nu = (\gamma \setminus \nu) \setminus \text{int}(C_\epsilon)$, which is connected since the boundaries of $\gamma \setminus \nu$ and C intersect only on one arc along ∂C . Now suppose that ν contains C . Then note that $\gamma' \setminus \nu$ is almost the same as $\gamma \setminus \nu$. In fact if ν does not contribute to the boundary of C , then $\gamma' \setminus \nu = \gamma \setminus \nu$. Otherwise, $\gamma' \setminus \nu$ is obtained by shaving off a thin strips of width ϵ from the boundary of $\gamma \setminus \nu$. Thus $\gamma' \setminus \nu$ is connected. Also $\nu \setminus \gamma' = (\nu \setminus \gamma) \cup C_\epsilon$. Since $\nu \setminus \gamma$ and C_ϵ are both connected, and have a non-empty intersection, their union is also connected. ◀

► **Lemma 11.** *Any maximal cell C in \mathcal{A} of depth at least two, and containing a hole, has exactly one hole H . Exactly one region $\gamma \in \Gamma$ contributes to the boundary of H ⁷. This region γ does not contribute any other arc to the boundary of C , and is a good region for C .*

Proof. First, observe that the boundary of any hole in C can be contributed to by at most one region. To see this assume to the contrary that two or more regions contribute to the boundary of a hole in C . If the boundaries of two of these regions intersect on the boundary of the hole, then the boundaries of those regions intersect the interior of the cell, which is impossible. Otherwise the hole belongs to at least two distinct regions, which violates the general position assumption.

Let H_γ be a hole in C whose boundary is contributed to by the region γ . We will show that H_γ is the only hole in C . If γ contributed to another hole, say H'_γ in C , then for any other region β (such a region exists, since $\text{depth}(C) \geq 2$) containing C intersects $\text{int}(H_\gamma)$, as well as $\text{int}(H'_\gamma)$, and thus $\beta \setminus \gamma$ cannot be connected. Thus γ cannot contribute to any other hole in C .

Let ρ be any other region containing C . Then, note that ρ intersects $\text{int}(H_\gamma)$ and since C is contained in both γ and ρ , we must have $\rho \setminus \gamma \subset \text{int}(H_\gamma)$, as otherwise $\rho \setminus \gamma$ would not be connected. This implies that $\rho \subset \text{int}(\gamma \cup H_\gamma)$. In particular, this means that $\rho \subset \text{int}(\tilde{\gamma})$, where $\tilde{\gamma}$ is the outer boundary of the region γ . If ρ contributed to the boundary of another hole H_ρ in C , then by the same argument, we would have $\gamma \subset \text{int}(\bar{\rho})$, a contradiction. This means that H_γ is the only hole in C .

Since the depth of C is at least two, there is at least one other region ρ containing C . However, as argued before $\rho \subset \text{int}(\gamma \cup H_\gamma)$. Since $C \subseteq \rho$ the boundary of γ cannot contribute any other arc (other than the boundary of H_γ) to the boundary of C .

We now argue that γ is a good region for C . Let $\gamma' = \gamma \setminus C_\epsilon$. The region γ' is connected since γ' is obtained from γ by replacing the hole H_γ by a larger hole $H = C_\epsilon \cup H_\gamma \subset \text{int}(\tilde{\gamma})$ ⁸ containing H_γ . H is simply connected since H_γ is the only hole in C . Also note that no other hole of γ intersects H . Let ν be any other region in Γ . We will show that both $\nu \setminus \gamma'$, and $\gamma' \setminus \nu$ are connected. Suppose first that $\nu \cap C = \emptyset$. Then, $\nu \setminus \gamma' = \nu \setminus \gamma$, which by assumption, is connected. Also, note that $\gamma' \setminus \nu$ is obtained by replacing the hole H_γ in $\gamma \setminus \nu$ by the larger hole H which contains H_γ , is contained in $\text{int}(\tilde{\gamma})$, and does not intersect any other holes in $\gamma \setminus \nu$. Thus $\gamma' \setminus \nu$ is connected. Now suppose that $C \subset \nu$. Then $\gamma' \setminus \nu$ is almost the same as $\gamma \setminus \nu$ and is obtained by shaving off a thin strip of width ϵ from the boundary of $\gamma \setminus \nu$. Thus $\gamma' \setminus \nu$ is connected. Also, $\nu \setminus \gamma' = (\nu \setminus \gamma) \cup C_\epsilon$. Since $\nu \setminus \gamma$ intersects C_ϵ near the boundary of the hole H_γ , $\nu \setminus \gamma'$ is connected. ◀

Lemmas 10 and 11 imply the following lemma.


► **Lemma 12.** *Any maximal cell C of depth at least two in the arrangement \mathcal{A} has a good region.*

► **Definition 13 (Cell Bypassing).** Let C be a maximal cell and γ be a good region for C . Then, by cell-bypassing of C by γ , we mean the modification of γ to $\gamma' = \gamma \setminus \text{int}(C_\epsilon)$. See Figures 2a and 2b.


The following observations are immediate.

⁷ In other words H is a hole of γ .

⁸ Recall that γ does not contribute to the outer boundary of C .



(a) Case 1: C is simply connected. The portion of the boundary of γ' distinct from γ is shown in red.



(b) Case 2: C is not simply connected. The hole H_γ of γ is expanded to H , so that it now contains C .

■ **Figure 2** Bypassing of a cell C by a good region γ .

► **Observation 14.** If a maximal cell C is identical to a region γ , then γ is good for C and cell-bypassing replaces γ by an empty region, effectively removing γ from the family of regions.

► **Observation 15.** When a cell C of depth d is bypassed by a region γ , the number of cells of depth d decreases by 1. The number of cells of lower depth may increase. In fact, there are at most $O(\Delta)$ newly created cells of depth either $d - 2$ or $d - 3$, where Δ is the degree of the cell C .

3 Construction of a planar support

We first give the construction of a planar support for the dual hypergraph defined by a set of non-piercing regions, and all points in the plane. We then use this to construct a planar support for the intersection hypergraph of two non-piercing families.

3.1 Planar Support for Dual Hypergraph

In the following, let Γ be a family of non-piercing regions and let \mathcal{A} denote their arrangement. We construct a planar support for the hypergraph $\mathcal{H}(\Gamma, \mathbb{R}^2)$ ⁹, which immediately implies a planar support for any $P \subseteq \mathbb{R}^2$. For a set of $S \subseteq \Gamma$ of the regions and any point $p \in \mathbb{R}^2$, we denote the set of regions in S that contain p by S_p . Similarly for a cell C , S_C denotes the set of regions in S containing C . For a graph G and a subset U of the vertices of G , we denote the subgraph of G induced by U by $G[U]$.

► **Lemma 16.** Let C be a maximal cell in \mathcal{A} with $\text{depth}(C) \geq 3$. Let Γ' be the arrangement obtained by cell bypassing of C by a good region γ of C . Then, a planar support for $\mathcal{H}(\Gamma, \mathbb{R}^2)$ can be constructed from a planar support for $\mathcal{H}(\Gamma', \mathbb{R}^2)$.

Proof. Suppose ∂C consists of a single arc. Then, C is identical to some region $\gamma \in \Gamma$. Since $\text{depth}(C) \geq 3$, γ is completely contained in another region $\rho \in \Gamma$. Bypassing C in this case results in the arrangement $\Gamma' = \Gamma \setminus \{\gamma\}$. Given a planar support G' for $\mathcal{H}(\Gamma', \mathbb{R}^2)$, we can obtain a planar support G for $\mathcal{H}(\Gamma, \mathbb{R}^2)$ by simply adding a new vertex for γ and an edge between γ and ρ . Since we obtain G by adding a vertex of degree 1 to the planar graph G' , G is planar. To see that G is a support, consider any point p in the plane. If $p \notin \gamma$ then $\Gamma_p = \Gamma'_p$ and therefore $G[\Gamma_p] = G'[\Gamma'_p]$ which is connected. On the other hand, if $p \in \gamma$,

⁹ While there are infinitely many points, the hypergraph is still finite, since all points in the same cell of the arrangement \mathcal{A} define the same hyperedge.

then $\Gamma_p = \Gamma'_p \cup \{\gamma\}$. Since G' is a planar support for Γ' , $G[\Gamma'_p] = G'[\Gamma'_p]$ is connected. Since $\rho \in \Gamma'(p)$ and there is an edge between γ and ρ in G , it follows that $G[\Gamma(p)]$ is connected.

Now, suppose ∂C contains at least two arcs. Let $\rho \neq \gamma$ be a region that contributes to ∂C . Then, in the arrangement of Γ , there is a cell X adjacent to C such that $\Gamma_X = \Gamma_C \setminus \{\rho\}$. After γ is modified to bypass C , in the arrangement of Γ' , there is still a cell X' such that $\Gamma'_{X'} = \Gamma_X$ and there is a cell C' s.t. $\Gamma'_{C'} = \Gamma_C \setminus \{\gamma\}$. Since $|\Gamma_C| \geq 3$, the sets $\Gamma'_{X'}$ and $\Gamma'_{C'}$ intersect. Also note that their union is Γ_C . Since $G'[\Gamma'_{X'}]$ and $G'[\Gamma'_{C'}]$ are connected, it follows that $G'[\Gamma_C]$ is connected. Thus $G = G'$ is a planar support for $\mathcal{H}(\Gamma, \mathbb{R}^2)$. ◀

The following lemma follows from [4]. The intersection graph of a set of regions is a graph in which the vertices correspond to the regions and edges correspond to pairs of intersecting regions.

► **Lemma 17** ([4]). *If Γ is a family of non-piercing regions so that no cell in their arrangement has depth more than 2, then the intersection graph of Γ is planar, and is a support for $\mathcal{H}(\Gamma, \mathbb{R}^2)$.*

► **Remark.** Even though the proof presented in [4] is for k -admissible regions, essentially the same proof works for non-piercing regions (with holes).

► **Theorem 18.** *Let Γ be a family of non-piercing regions. Then, there exists a planar support for $\mathcal{H}(\Gamma, \mathbb{R}^2)$.*

Proof. For any set of regions Γ , define d_Γ to be the maximum depth of any cell in the arrangement of Γ , and n_Γ to be the number of cells in the arrangement with depth d_Γ . We use induction on the pair (d_Γ, n_Γ) . If $d_\Gamma \leq 2$, then by Lemma 17, the intersection graph of Γ is the required planar support. Given a set of regions with $d_\Gamma \geq 3$, let Γ' be the set of regions obtained after bypassing a cell of maximum depth in the arrangement of Γ . Then, by Observation 15, the pair $(d_{\Gamma'}, n_{\Gamma'})$ is lexicographically smaller than (d_Γ, n_Γ) . Inductively, a planar support for $\mathcal{H}(\Gamma', \mathbb{R}^2)$ exists, which by Lemma 16, is also a planar support for $\mathcal{H}(\Gamma, \mathbb{R}^2)$. ◀


3.2 Planar Support for the Intersection Hypergraph

Given a family R of *red* non-piercing regions, and a family B of *blue* non-piercing regions, we prove that there exists a planar support for the intersection hypergraph $\mathcal{H}(B, R) = (B, E)$, where $E = \{B_r : r \in R\}$, and $B_r = \{b \in B : b \cap r \neq \emptyset\}$. However, it is not essential for the rest of the paper. Note that we can assume without loss of generality that any red region intersects at least one blue region, and we make this assumption throughout this section. This implies that any maximal cell in the arrangement of $R \cup B$ has depth at least 2.

Let $B_{|r} = \{b \cap r : b \in B_r\}$ be the set of regions obtained by intersecting the regions in B with the region r . A region in $B_{|r}$ may have multiple connected components, but we still treat it as a single region. Let $G_{|r}(B)$ be the intersection graph of these regions. We start with the following simple observation.

► **Lemma 19.** *If for each red region $r \in R$, the graph $G_{|r}(B)$ is connected, then the planar support for the hypergraph $\mathcal{H}(B, \mathbb{R}^2)$, is also a planar support for the hypergraph $\mathcal{H}(B, R)$.*

Proof. Let G be the planar support of the hypergraph $\mathcal{H}(B, \mathbb{R}^2)$ defined by the blue regions. Recall that the graph G guarantees that for each $p \in \mathbb{R}^2$, the set of blue regions containing p induce a connected subgraph of G . Consider any red region $r \in R$. Since, by assumption, the graph $G_{|r}(B)$ is connected, for any pair of blue regions $s, t \in B_r$ intersecting r , there is



■ **Figure 3** The construction of β_C , for a non-maximal cell C not contained in a blue region in B .

a sequence $s = b_1, \dots, b_k = t$, of regions in B_r such that adjacent regions b_i, b_{i+1} intersect at a point $q_i \in r$. This means that there is a path in G between b_i and b_{i+1} via the regions containing q_i , i.e., via regions in B_r . This in turn implies that there is a path between s and t in G via regions in B_r . ◀

The construction of the planar support in the general setting is a reduction to the setting in Lemma 19.

► **Lemma 20.** *Given a set of red and blue non-piercing regions R and B , we can obtain a modified set of red regions R' , such that (i) the set of regions R' is non-piercing, (ii) in the arrangement of $R' \cup B$, each maximal cell is contained in some blue region in B , and (iii) the intersection hypergraph $\mathcal{H}(B, R')$ is isomorphic to $\mathcal{H}(B, R)$.*

Proof. Consider a maximal cell C in the arrangement of $R \cup B$ that is not contained in any blue region. Then, C is also a maximal cell in the arrangement of R . Since R is a non-piercing family, by Lemma 12 we can bypass C . This does not change the intersection hypergraph of the red and blue regions, since no red-blue intersection is lost or gained as a consequence of bypassing the cell C . We repeat this process until each maximal cell is in some blue region. The modified set of regions thus obtained satisfy the conditions of the lemma. ◀

► **Lemma 21.** *Given two families of red and blue non-piercing regions R and B , such that each maximal cell in the arrangement \mathcal{A} of $R \cup B$ is contained in some region $b \in B$, we can add a fake blue region β_C corresponding to each cell C in \mathcal{A} that is not contained in any blue region in B , such that R, B' and B'' , where B'' is the set of fake blue regions and $B' = B \cup B''$, satisfy: (i) B' is a family of non-piercing regions, (ii) each β_C intersects only those regions in R that contain C , and (iii) for each $r \in R$, the intersection graph $G|_r(B')$ is connected.*

Proof. Let C be a cell not contained in any of the blue regions. We define the fake blue region for this cell as $\beta_C = C_\epsilon \setminus \bigcup_{r \in R, C \not\subset r} \text{int}(r_\delta)$, where $\delta < \epsilon$ and ϵ is sufficiently small. See Figure 3. Intuitively, β_C is roughly the same as C but we modify its boundary slightly so that it intersects all the blue regions that contribute to the boundary C but does not intersect any red regions not containing C . Defining β_C this way ensures that property (ii) in the statement of the Lemma is satisfied. Choosing ϵ to be sufficiently small also ensures that property (i) is satisfied. Finally, choosing $\delta < \epsilon$ also ensures that each red region in R is covered by the union of the blue regions in B' which implies property (iii). To see this, consider a point p contained in $r \in R$, and let D be the cell in \mathcal{A} containing p . If p is not contained in a blue region in B , then since D does not lie in any blue region, we add a fake blue region β_D corresponding to D . If p lies at a distance of at least δ from any arc of ∂D ,

contributed by a red region in R not containing D , then note that p lies in β_D . Otherwise, if p lies within distance δ of some arc α contributed by a red region not containing D , the cell D' adjacent to D sharing the arc α is also not contained in any blue region in B . In this case, since $\epsilon > \delta$, $\beta_{D'}$ contains p . This implies that $G_{|r}(B)$ is connected for each $r \in R$. ◀

We are now ready to prove Theorem 1.

Proof of Theorem 1. If $G_{|r}(B)$ is connected for each $r \in R$, then we obtain a planar support by Lemma 19. If not, let R' be the set of modified red regions obtained by applying Lemma 20. Let B'' be the set of fake blue regions added by applying Lemma 21 to the blue regions B , and the modified red regions R' . Now, by Lemma 19, a planar support for $\mathcal{H}(B', \mathbb{R}^2)$ is a planar support for $\mathcal{H}(B', R')$. Let G' be this planar support. We show that we can obtain a planar support G for $\mathcal{H}(B, R')$, by suitably modifying G' . G is also a support for $\mathcal{H}(B, R)$ since $\mathcal{H}(B, R)$ is isomorphic to $\mathcal{H}(B, R')$. In the following, we refer to a vertex in the support graph by the corresponding region. Let \mathcal{A} be the arrangement of $B \cup R'$, and let \mathcal{B}' denote the arrangement of the regions in B' . Let C be a cell in \mathcal{A} not contained in any blue region in B . By Lemma 20, C is not a maximal cell. Thus, there is a fake blue region $\beta_C \in B'$ corresponding to C . Since C is not maximal, we can pick a cell C' (arbitrarily chosen in case of ties) adjacent to C in \mathcal{A} , such that $\text{depth}(C') = \text{depth}(C) + 1$. Let b be a fake or real blue region defined as follows: if C' is not contained in any blue region in B then $b = \beta_{C'}$, otherwise b is the unique blue region in B containing C' .

Note that b forms a depth 2 intersection with β_C in \mathcal{B}' (i.e. there is a point in the plane contained in just these two regions), and therefore, β_C and b are adjacent in G' . We orient the corresponding edge in G' from β_C to b . Thus, for each fake region in G' , there is exactly one outgoing oriented edge incident on it. Note that not all edges in G' are oriented.

By Property (ii) of Lemma 21, any red region intersecting β_C contains C . Such a region also contains C' since C and C' are adjacent cells in \mathcal{A} , and $\text{depth}(C') = \text{depth}(C) + 1$. Thus all red regions intersecting β_C also intersect b . In other words, the set of red regions in R' intersecting β_C is a subset of the red regions in R' intersecting b . This is a key property we will use later. By Lemma 20, each maximal cell in \mathcal{A} is contained in a blue region in B , and therefore there is a unique directed path starting from a fake region β_C and ending at a real blue region $\tilde{b} \in B$. Crucially, the set of red regions intersecting any fake region β_C is a subset of the red regions intersecting \tilde{b} due to the key property mentioned earlier.

The set of oriented edges in G' form a spanning forest, where each arc is oriented towards the root of a tree, and the root of each tree corresponds to a real region. We obtain G by contracting the edges in the forest (i.e., all oriented edges in G), effectively merging all nodes in a sub-tree to its root. Since edge contraction preserves planarity, it follows that G is planar. To see that G is a support for $\mathcal{H}(B, R')$, let $b, b' \in B$ be a pair of blue regions intersecting a red region $r \in R'$. Since G' is a planar support for $\mathcal{H}(B', R')$, there is a path $b = b_1, \dots, b_k = b'$ such that each region b_i in the path intersects r . Since each fake region on the path is merged with a real region that also intersects r , it follows that there is a path in G between b and b' , such that all regions along the path intersect r .

Finally, by Lemma 20, since $\mathcal{H}(B, R')$ is isomorphic to $\mathcal{H}(B, R)$, G is the desired planar support. ◀

3.3 Algorithms

The proofs for the existence of the planar supports given in Section 3 are constructive and can be converted to an algorithm with running time $O(n^3 + m)$ where n is the number of regions and m is the number of vertices in the arrangement of the regions. Note that if the

boundaries of every pair of regions intersect at most $O(n)$ times, then m is $O(n^3)$. While the applications mentioned in this paper require only the existence of a planar support, the algorithmic question is natural and may have applications in hypergraph visualization (see [2] and [8]). In the proofs of the existence of the planar support, we modify the arrangement of the regions by doing cell bypassing. For algorithmic purposes, instead of maintaining the geometric shapes of the regions, we can maintain the dual arrangement graph. In addition, we orient the edges from a cell of lower depth to a cell of higher depth and label the edge by the index of the region whose boundary separates the cells. The vertex corresponding to a cell also stores the depth of the cell and pointers to neighboring vertices. The size of the graph is $O(m)$. We assume that the initial dual arrangement graph H of the input set of region Γ is given. It is not at all apparent that these algorithms run in polynomial time, since the number of cells in the arrangement can increase after each cell-bypassing. However, using a more detailed combinatorial analysis, we can define a suitable potential function on the arrangement of the regions, so that each time we do a cell bypassing operation the potential goes down by an amount proportional to the time spent in the cell-bypassing operation. The initial potential can be shown to be $O(n^3 + m)$ which implies the upper bound on the running time. The details of the analysis are not included due to lack of space.

References

- 1 Aniket Basu Roy, Sathish Govindarajan, Rajiv Raman, and Saurabh Ray. Packing and covering with non-piercing regions. *Discrete & Computational Geometry*, Mar 2018. doi: 10.1007/s00454-018-9983-2.
- 2 Kevin Buchin, Marc Van Kreveld, Henk Meijer, Bettina Speckmann, and Kevin Verbeek. On planar supports for hypergraphs. In *International Symposium on Graph Drawing*, pages 345–356. Springer, 2009.
- 3 Timothy M. Chan, Elyot Grant, Jochen Könemann, and Malcolm Sharpe. Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '12*, pages 1576–1585, 2012.
- 4 Timothy M. Chan and Sarel Har-Peled. Approximation algorithms for maximum independent set of pseudo-disks. *Discrete & Computational Geometry*, 48(2):373–392, 2012.
- 5 Kenneth L Clarkson and Peter W Shor. Applications of random sampling in computational geometry, ii. *Discrete & Computational Geometry*, 4(5):387–421, 1989.
- 6 Alina Ene, Sarel Har-Peled, and Benjamin Raichel. Geometric packing under non-uniform constraints. In *Proceedings of the Twenty-eighth Annual Symposium on Computational Geometry, SoCG '12*, pages 11–20, New York, NY, USA, 2012. ACM.
- 7 Matt Gibson and Imran A. Pirwani. Algorithms for dominating set in disk graphs: Breaking the $\log n$ barrier. In *Algorithms – ESA 2010 - 18th Annual European Symposium, Liverpool, United Kingdom, September 6–8, 2010, Proceedings*, pages 243–254, 2010.
- 8 David S Johnson and Henry O Pollak. Hypergraph planarity and the complexity of drawing venn diagrams. *Journal of graph theory*, 11(3):309–325, 1987.
- 9 Chaya Keller and Shakhar Smorodinsky. Conflict-free coloring of intersection graphs of geometric objects. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 2397–2411, 2018. doi:10.1137/1.9781611975031.154.
- 10 Balázs Keszegh. Coloring intersection hypergraphs of pseudo-disks. In *Proceedings of the Thirty-fourth International Symposium on Computational Geometry*, 2018. URL: <http://arxiv.org/abs/1711.05473>.

69:14 Planar Support for Non-piercing Regions and Applications

- 11 Nabil H. Mustafa, Rajiv Raman, and Saurabh Ray. Quasi-polynomial time approximation scheme for weighted geometric set cover on pseudodisks and halfspaces. *SIAM Journal on Computing*, 44(6):1650–1669, 2015.
- 12 Nabil H. Mustafa and Saurabh Ray. Improved results on geometric hitting set problems. *Discrete & Computational Geometry*, 44(4):883–895, 2010.
- 13 Evangelia Pyrga and Saurabh Ray. New existence proofs for ϵ -nets. In *Proceedings of the Twenty-fourth Annual Symposium on Computational Geometry*, SoCG '08, pages 199–207, New York, NY, USA, 2008. ACM.
- 14 Kasturi Varadarajan. Weighted geometric set cover via quasi-uniform sampling. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, pages 641–648, 2010.