



Coloring Intersection Hypergraphs of Pseudo-Disks

Balázs Keszegh¹

Received: 13 July 2018 / Revised: 29 September 2019 / Accepted: 9 October 2019 /
Published online: 30 October 2019
© The Author(s) 2019

Abstract

We prove that the intersection hypergraph of a family of n pseudo-disks with respect to another family of pseudo-disks admits a proper coloring with four colors and a conflict-free coloring with $O(\log n)$ colors. Along the way we prove that the respective Delaunay-graph is planar. We also prove that the intersection hypergraph of a family of n regions with linear union complexity with respect to a family of pseudo-disks admits a proper coloring with constantly many colors and a conflict-free coloring with $O(\log n)$ colors. Our results serve as a common generalization and strengthening of many earlier results, including ones about proper and conflict-free coloring points with respect to pseudo-disks, coloring regions of linear union complexity with respect to points and coloring disks with respect to disks.

Keywords Geometric hypergraph coloring · Intersection hypergraph · Delaunay-graph · Pseudo-disks · Regions with linear union complexity

1 Introduction

Proper colorings of hypergraphs and graphs defined by geometric regions are studied extensively due to their connections to, among others, conflict-free colorings and to cover-decomposability problems, both of which have real-life motivations in cellular networks, scheduling, etc. For applications and history we refer to the surveys [25,28]. Here we restrict our attention to the (already long list of) results that directly precede our results. We discuss further directions and connections to other problems in Sect. 6.

Dedicated to the memory of Ricky Pollack.

Editor in Charge: János Pach

Research supported by the Lendület program of the Hungarian Academy of Sciences (MTA), under the Grant LP2017-19/2017 and by the National Research, Development and Innovation Office—NKFIH under the Grant K 116769.

Balázs Keszegh
keszegh@renyi.hu

¹ Alfréd Rényi Institute of Mathematics, Budapest, Hungary

Our central family of regions is the family of pseudo-disks. Families of pseudo-disks have been regarded in many settings for a long time due to being the natural way of generalizing disks while retaining many of their topological and combinatorial properties. Problems ranged from classic algorithmic questions like finding maximum size independent (disjoint) subfamilies [6] to classical combinatorial geometric questions like the Erdős–Szekeres problem [12]. Probably the most important pseudo-disk family is the family of homothets of a convex region. Pseudo-disks are also central in the area of geometric hypergraph coloring problems as we will see soon in this section.

Before we state our results and their precursors, we need to introduce some basic definitions.

Definition 1.1 Given a hypergraph \mathcal{H} , a *proper coloring* of its vertices is a coloring in which every hyperedge H of size at least 2 of \mathcal{H} contains two differently colored vertices.

Throughout the paper hypergraphs can contain hyperedges of size at least 2 only, in other words we do not allow (or automatically delete, when necessary) hyperedges of size 1.

Definition 1.2 Given a hypergraph \mathcal{H} , a *conflict-free coloring* of its vertices is a coloring in which every hyperedge H contains a vertex whose color differs from the color of every other vertex of H .

Definition 1.3 Given a family \mathcal{B} of regions and a family \mathcal{F} of regions, the *intersection hypergraph* of \mathcal{B} with respect to (wrt. in short) \mathcal{F} is the simple (that is, it has no multiple hyperedges) hypergraph $\mathcal{I}(\mathcal{B}, \mathcal{F})$ which has a vertex v_B corresponding to every $B \in \mathcal{B}$ and for every $F \in \mathcal{F}$ it has a hyperedge $H = \{v_B : B \cap F \neq \emptyset\}$ (if this set has size at least 2) which corresponds to F . The *Delaunay-graph* of \mathcal{B} wrt. \mathcal{F} is the graph on the same vertex set containing only the hyperedges of $\mathcal{I}(\mathcal{B}, \mathcal{F})$ of size 2.

Note that if \mathcal{B} is finite, then even if \mathcal{F} is infinite, $\mathcal{I}(\mathcal{B}, \mathcal{F})$ has finitely many hyperedges. In particular, a hyperedge H can correspond to multiple members of \mathcal{F} .

We also define the following subgraph of the Delaunay-graph:

Definition 1.4 The *restricted Delaunay-graph* of \mathcal{B} wrt. \mathcal{F} is the subgraph of the Delaunay-graph containing only those (hyper)edges $H_F = \{v_{B_1}, v_{B_2}\}$ for which the corresponding $F \in \mathcal{F}$ intersects B_1 and B_2 in disjoint regions, that is, $F \cap B_1 \neq \emptyset$, $F \cap B_2 \neq \emptyset$, $F \cap B_1 \cap B_2 = \emptyset$ and $F \cap B = \emptyset$ for every $B \in \mathcal{B} \setminus \{B_1, B_2\}$.

Observation 1.5 ([27]) *If \mathcal{B} and \mathcal{F} are families for which every hyperedge of $\mathcal{I}(\mathcal{B}, \mathcal{F})$ contains a hyperedge of size 2, then a proper coloring of the Delaunay-graph is also a proper coloring of $\mathcal{I}(\mathcal{B}, \mathcal{F})$.*

In the literature hypergraphs that have the property assumed in Observation 1.5 are sometimes called *rank two* hypergraphs (e.g., in [27]).

Definition 1.6 We say that \mathcal{B} can be properly/conflict-free/etc. colored wrt. \mathcal{F} with $f(n)$ colors if for any \mathcal{B}' subfamily of \mathcal{B} of size n , the hypergraph $\mathcal{I}(\mathcal{B}', \mathcal{F})$ can be properly/conflict-free/etc. colored with $f(n)$ colors.

We note that one could further generalize this notion so that instead of intersection we regard inclusion/reverse-inclusion as the relation that defines the hyperedges, as it is done, e.g., in [18] for hypergraphs defined by intervals. For problems related to such variants see Sect. 6.

In the majority of earlier results, coloring points wrt. a family of regions or coloring a family of regions wrt. points were regarded. While they were not defined as intersection hypergraphs, they do fit into our general definition of intersection hypergraphs, choosing either \mathcal{B} or \mathcal{F} to be the family of points of the plane. Notice that whenever every point is (or can added to be) a member of the family \mathcal{F} , e.g., for disks, for homothets of a convex region and for pseudo-disks¹, coloring intersection hypergraphs of any finite subfamily \mathcal{F}' of \mathcal{F} wrt. to \mathcal{F} is a common generalization both of coloring any finite set of points wrt. \mathcal{F} and of coloring any finite subfamily \mathcal{F}' of \mathcal{F} wrt. points.

There are only a few earlier results regarding this general setting. In [18] intersection (and also inclusion and reverse-inclusion) hypergraphs of intervals of the line were considered. In [10,14] they considered intersection hypergraphs (and graphs) of (unit) disks, pseudo-disks, squares and axis-parallel rectangles.

1.1 Results Related to Pseudo-Disks

It is well known that the Delaunay-graph of points wrt. disks is planar and thus proper 4-colorable, and Observation 1.5 implies that the respective hypergraph is also proper 4-colorable.

In [27] Smorodinsky developed a general framework (based on the framework presented in [9]): using proper colorings of the subhypergraphs of the hypergraph with constantly many colors it can build a conflict-free coloring with $O(\log n)$ colors (where n is the number of vertices of the hypergraph). The results mentioned from now on use this framework to get a conflict-free coloring once there is a proper coloring. First, using the precursor of this framework, it was proved by Even et al. [9] that points wrt. disks admit a conflict-free coloring with $O(\log n)$ colors:

Theorem 1.7 ([9]) *Let \mathcal{D} be the family of disks in the plane and P a finite set of points, $\mathcal{I}(P, \mathcal{D})$ admits a conflict-free coloring with $O(\log n)$ colors, where $n = |P|$.*

Definition 1.8 A *Jordan region* is a (simply connected) closed bounded region whose boundary is a closed simple Jordan curve.

A family of Jordan regions is called a family of *pseudo-disks* if the boundaries of every pair of the regions intersect in at most two points.

Although we could not find it in the literature, it is well known that if instead of a family of disks we take a family of pseudo-disks, the bound of Theorem 1.7 still holds:

Theorem 1.9 (Folklore) *Let \mathcal{F} be a family of pseudo-disks and P a finite set of points, $\mathcal{I}(P, \mathcal{F})$ admits a proper coloring with four colors and a conflict-free coloring with $O(\log n)$ colors, where $n = |P|$.*

¹ See Corollary 2.3 for a precise formulation of this property for pseudo-disks.

The proof is that the Delaunay-graph of points with respect to pseudo-disks is also a planar graph (implied by, e.g., Lemma 2.6, see later) and thus proper 4-colorable and then we can apply Observation 1.5 (we can assume that the hypergraph is rank two by, e.g., Corollary 2.3) to conclude that the respective hypergraph is also 4-colorable. Then the general framework of Smorodinsky mentioned above implies an $O(\log n)$ upper bound for a conflict-free coloring.

The dual of Theorem 1.7 was also proved by Even et al. and was generalized to pseudo-disks by Smorodinsky:

Theorem 1.10 ([9]) *Let P be the set of all points of the plane and \mathcal{B} a finite family of disks, $\mathcal{I}(\mathcal{B}, P)$ admits a proper coloring with four colors and a conflict-free coloring with $O(\log n)$ colors, where $n = |\mathcal{B}|$.*

Theorem 1.11 ([27]) *Let P be the set of all points of the plane and \mathcal{B} be a finite family of pseudo-disks, $\mathcal{I}(\mathcal{B}, P)$ admits a proper coloring with a constant number of colors and a conflict-free coloring with $O(\log n)$ colors, where $n = |\mathcal{B}|$.*

While for coloring pseudo-disks wrt. points there was no explicit upper bound, for the special case of homothets of a convex region², the constant was shown by Cardinal and Korman to be 4, just like for disks:

Theorem 1.12 ([5]) *Let P be the set of all points of the plane and \mathcal{B} be a finite family of homothets of a given convex region, $\mathcal{I}(\mathcal{B}, P)$ admits a proper coloring with four colors.*

Recently it was proved by Keller and Smorodinsky that disks wrt. disks can also be colored in such a way, a common generalization of Theorems 1.7 and 1.10 :

Theorem 1.13 ([14]) *Let \mathcal{D} be the family of all disks in the plane and \mathcal{B} a finite family of disks, $\mathcal{I}(\mathcal{B}, \mathcal{D})$ admits a proper coloring with six colors and a conflict-free coloring with $O(\log n)$ colors, where $n = |\mathcal{B}|$.*

While they did not prove the same for pseudo-disks, they solved two special cases (they stated these only as part of the proof of their main result):

Claim 1.14 ([14]) *Given a finite family \mathcal{F} of pseudo-disks and a subfamily \mathcal{B} of \mathcal{F} , if either \mathcal{B} or $\mathcal{F} \setminus \mathcal{B}$ contains only pairwise disjoint pseudo-disks, then $\mathcal{I}(\mathcal{B}, \mathcal{F})$ admits a proper coloring with a constant number of colors and a conflict-free coloring with $O(\log n)$ colors, where $n = |\mathcal{B}|$.*

In this paper we generalize Theorem 1.13 to the case of coloring a pseudo-disk family wrt. another pseudo-disk family, which is a common generalization of all the above results (Theorems 1.7, 1.10, 1.11, 1.12, 1.13 and Claim 1.14). Moreover we prove the optimal upper bound of four colors, which improves the bound of Theorem 1.11 for coloring pseudo-disks wrt. points from some constant number of colors to four colors and improves the bound of Theorem 1.13 for coloring disks wrt. disks from six colors to four colors. Furthermore, it provides an alternative proof for Theorems 1.10 and 1.12 (both of these were originally proved using dualization and solving equivalent problems about coloring points in the 3-dimensional space):

² For further results about homothets of a convex region see Sect. 6.

Theorem 1.15 *Given a family \mathcal{F} of pseudo-disks and a finite family \mathcal{B} of pseudo-disks, $\mathcal{I}(\mathcal{B}, \mathcal{F})$ admits a proper coloring with four colors.*

Along the way we prove that the respective Delaunay-graph is planar:

Theorem 1.16 *Given a family \mathcal{F} of pseudo-disks and a finite family \mathcal{B} of pseudo-disks, the Delaunay-graph of \mathcal{B} wrt. \mathcal{F} is a planar graph.*

This was not known even for the Delaunay-graph of pseudo-disks wrt. points. With standard methods (present also in the proof of Theorem 1.25) this already implies Theorem 1.15 with six colors instead of four. To achieve the optimal bound we will need some additional ideas.

We mention that if \mathcal{B} and \mathcal{F} are both families of pairwise disjoint simply connected regions, then $\mathcal{I}(\mathcal{B}, \mathcal{F})$ is a planar hypergraph³ and all planar hypergraphs can be generated this way. From this perspective Theorem 1.15 says that even with a much more relaxed definition of planarity of a hypergraph it remains 4-colorable.

Using the usual framework it easily follows from Theorem 1.15 that:

Corollary 1.17 *Given a family \mathcal{F} of pseudo-disks and a finite family \mathcal{B} of pseudo-disks, $\mathcal{I}(\mathcal{B}, \mathcal{F})$ admits a conflict-free coloring with $O(\log n)$ colors, where $n = |\mathcal{B}|$.*

We note that Claim 1.14 implies in a straightforward way the main result of [14] about conflict-free coloring the (open/closed) neighborhood hypergraphs of intersection graphs of pseudo-disks (for definitions and details see [14]). Thus, Theorem 1.15 also implies their main result.

Note that in Theorem 1.15, \mathcal{B} and \mathcal{F} are not related in any way, thus among others, even though two convex regions can intersect infinitely many times, it implies somewhat surprisingly the following:

Corollary 1.18 *We can proper color with four colors the family of homothets of a convex region A wrt. the family of homothets of another convex region B .*

Buzaglo et al. [4] have shown that the VC-dimension of the hypergraph $\mathcal{I}(P, \mathcal{F})$, defined by a finite point set P with respect to a family \mathcal{F} of pseudo-disks, is at most 3 (and this bound is tight) and used this to prove that the number of hyperedges of size at most k in such a hypergraph is $O(k^2n)$ (and this can be attained already by \mathcal{F} being a specific family of disks). More recently, Aronov et al. [3] proved independently from us with an almost identical proof the case of Lemma 2.7 when $\mathcal{B} \subset \mathcal{F}$ (for details see later, it is a very special case of Theorem 1.16) and showed that this implies that for a pseudo-disk family \mathcal{F} and a finite subfamily \mathcal{B} of \mathcal{F} , the intersection hypergraph $\mathcal{I}(\mathcal{B}, \mathcal{F})$ has VC-dimension at most 4 (and this bound is tight). Using methods similar to the ones in [4] they show that this implies that the number of hyperedges of size at most k in $\mathcal{I}(\mathcal{B}, \mathcal{F})$ in this case is $O(k^3n)$ (they do not give a construction with a matching lower bound, however). We show that using Lemma 2.7 instead of their weaker variant we get the following more general statements (where \mathcal{B} is not necessarily a subfamily of \mathcal{F}):

³ The most common definition of a planar hypergraph is that its bipartite incidence graph is planar.

Theorem 1.19 *Given a family \mathcal{F} of pseudo-disks and a finite family \mathcal{B} of pseudo-disks, the intersection hypergraph $\mathcal{I}(\mathcal{B}\mathcal{F})$ has VC-dimension at most 4 (and this bound is tight).*

Theorem 1.20 *Given a family \mathcal{F} of pseudo-disks and a family \mathcal{B} of n pseudo-disks, the number of hyperedges of size at most k in the intersection hypergraph $\mathcal{I}(\mathcal{B}\mathcal{F})$ is $O(k^3n)$.*

We note that in case of Theorem 1.20 we are again not aware of a matching lower bound (and in fact this upper bound might easily not be tight).

1.2 Results Related to Families of Linear Union Complexity

In [27] Theorem 1.11 was shown by proving a more general statement about coloring a family of regions that has linear union complexity wrt. points.

Definition 1.21 Let \mathcal{B} be a family of finitely many Jordan regions in the plane such that the boundaries of its members intersect in finite many points. The *vertices* of the *arrangement* of \mathcal{B} are the intersection points of the boundaries of regions in \mathcal{B} , the *edges* are the maximal connected parts of the boundaries of regions in \mathcal{B} that do not contain a vertex and the *faces* are the maximal connected parts of the plane which are disjoint from the edges and the vertices of the arrangement.

Definition 1.22 The *union complexity* $\mathcal{U}(\mathcal{B})$ of a family of Jordan regions \mathcal{B} is the number of edges of the arrangement \mathcal{B} that lie on the boundary of $\bigcup_{B \in \mathcal{B}} B$.

We say that a family of regions \mathcal{B} has (*c*)-*linear union complexity* if there exists a constant c such that for any subfamily \mathcal{B}' of \mathcal{B} the union complexity of \mathcal{B}' is at most $c|\mathcal{B}'|$.⁴

Theorem 1.23 ([13]) *Any finite family of pseudo-disks in the plane has a linear union complexity.*

Theorem 1.23 shows that the following result of Smorodinsky is indeed more general than Theorem 1.11.

Theorem 1.24 ([27]) *Let P be the set of all points of the plane and \mathcal{B} be a finite family of Jordan regions with linear union complexity. Then $\mathcal{I}(\mathcal{B}, P)$ admits a proper coloring with a constant number of colors and a conflict-free coloring with $O(\log n)$ colors, where $n = |\mathcal{B}|$.*⁵

We generalize this in the following way:

⁴ Sometimes in the literature, in the definition of union complexity they count vertices rather than edges, yet it is easy to see (see Lemma 5.1 later) that this does not affect the property of having linear union complexity when \mathcal{B} is a family of Jordan regions. Also note that linear union complexity is not always defined hereditarily, we defined it this way in order to simplify our statements.

⁵ When a statement is about a family with (*c*)-linear union complexity, by a constant we mean a constant that depends on c and the O notation similarly hides a dependence on c .

Theorem 1.25 *Given a family \mathcal{F} of pseudo-disks and a finite family \mathcal{B} of Jordan regions with linear union complexity, $\mathcal{I}(\mathcal{B}, \mathcal{F})$ admits a proper coloring with a constant number of colors.*

Note that using Theorem 1.23 we get that Theorem 1.25 implies Theorem 1.15 with a (non-explicit and worse) upper bound.

Corollary 1.26 *Given a family \mathcal{F} of pseudo-disks and a finite family \mathcal{B} of Jordan regions with linear union complexity, $\mathcal{I}(\mathcal{B}, \mathcal{F})$ admits a conflict-free coloring with $O(\log n)$ colors, where $n = |\mathcal{B}|$.*

We note that the proof of Theorem 1.25 works without modification also when instead of Jordan-regions (i.e., bounded regions whose boundary is a single closed Jordan curve), \mathcal{B} is a family of bounded regions whose boundary is a union of a finite number of disjoint Jordan curves (in particular the regions are not necessarily simple or connected). Observe that the definition of union complexity is well-defined for such families as well, yet in this case we do need to count edges and not vertices to have a true statement (recall that previously it did not matter). Thus Theorem 1.25 and Corollary 1.26 hold in this more general setting as well.

The rest of the paper is structured as follows. In Sect. 2 we prove Theorem 1.16. In Sect. 3 we prove Theorem 1.15. Implications about conflict-free colorings and VC-dimensions and the number of bounded size hyperedges in intersection hypergraphs are detailed in Sect. 4. In Sect. 5 we prove the results about regions of linear union complexity. We also give an example that shows that we cannot always color properly points wrt. a family of linear union complexity with constantly many colors. This shows that Theorem 1.25 is strongest possible in the sense that we cannot change \mathcal{F} from pseudo-disks to a family with linear union complexity, even if \mathcal{B} would be only a finite family of points in the plane. Finally, in Sect. 6 we discuss some related (open) problems.

2 The Delaunay-Graph of Pseudo-Disks wrt. Pseudo-Disks

In this section we prove Theorem 1.16. First we list some tools we need from the papers of Pinchasi [26] and Buzaglo et al. [4] about pseudo-disks:

Lemma 2.1 ([26]) *Let \mathcal{F} be a family of pseudo-disks. Let $D \in \mathcal{F}$ and let $x \in D$ be any point. Then D can continuously be shrunk to the point x so that at each moment \mathcal{F} is a family of pseudo-disks.*

Note that when shrinking this way, we can keep all shrunk copies of D in the family, it remains a pseudo-disk family as their boundaries are pairwise disjoint.

Definition 2.2 We say that a pseudo-disk family \mathcal{F} is *shrinking-closed* (for some family of regions \mathcal{B}) if we cannot add a new pseudo-disk F to \mathcal{F} which is contained in a pseudo-disk F' already in \mathcal{F} (i.e., $F \subset F'$) in a way that strictly increases the number of hyperedges in $\mathcal{I}(\mathcal{B}, \mathcal{F})$.

We can enlarge \mathcal{F} greedily until it becomes shrinking-closed (we need to add at most $2^{|\mathcal{B}|}$ pseudo-disks—in fact at most polynomial many in $|\mathcal{B}|$ according to Theorem 1.20):

Corollary 2.3 *Given a family \mathcal{F} of pseudo-disks and a finite family \mathcal{B} of regions, there exists a shrinking-closed (for \mathcal{B}) family of pseudo-disks \mathcal{F}' with $\mathcal{F}' \supseteq \mathcal{F}$ and consequently with $\mathcal{I}(\mathcal{B}, \mathcal{F}') \supseteq \mathcal{I}(\mathcal{B}, \mathcal{F})$. In particular, for every point p which is in some $F \in \mathcal{F}$, $H_p = \{v_B : p \in B \in \mathcal{B}\}$ is a hyperedge of $\mathcal{I}(\mathcal{B}, \mathcal{F}')$ (which we call the hyperedge that corresponds to p) if it is of size at least two.*

The second part of Corollary 2.3 holds as for a point p contained in some $F \in \mathcal{F}$ we can apply Lemma 2.1.

By Corollary 2.3 when proving our results, it will be enough to consider the case when \mathcal{F} is shrinking-closed as adding hyperedges to a hypergraph cannot remove edges from its Delaunay-graph and cannot decrease the number of colors needed for a proper (or conflict-free) coloring. The following corollary of Lemma 2.1 will be useful when we deal with a shrinking-closed pseudo-disk family:

Corollary 2.4 ([26]) *Let \mathcal{B} be a family of pairwise disjoint regions in the plane and let \mathcal{F} be a family of pseudo-disks. Let D be a member of \mathcal{F} and suppose that D intersects exactly k members of \mathcal{B} one of which is the set $B \in \mathcal{B}$. Then for every $2 \leq l \leq k$ there exists a set $D' \subset D$ such that D' intersects B and exactly $l - 1$ other regions from \mathcal{B} , and $\mathcal{F} \cup \{D'\}$ is again a family of pseudo-disks.*

Lemma 2.5 ([4]) *Let D_1 and D_2 be two pseudo-disks in the plane. Let x and y be two points in $D_1 \setminus D_2$. Let a and b be two points in $D_2 \setminus D_1$. Let e be any Jordan arc connecting x and y that is fully contained in D_1 . Let f be any Jordan arc connecting a and b that is fully contained in D_2 . Then e and f cross an even number of times.*

Lemma 2.6 ([26]) *Given a family \mathcal{F} of pseudo-disks and a finite family \mathcal{B} of pairwise disjoint connected sets in the plane, the Delaunay-graph of \mathcal{B} wrt. \mathcal{F} is a planar graph.*

Note that for a family \mathcal{B} of pairwise disjoint connected sets, the Delaunay-graph and the restricted Delaunay-graph are the same. Assuming also that \mathcal{B} is a family of Jordan regions, Lemma 2.6 becomes a special case of Theorem 1.16. Before proving Theorem 1.16 we prove another special case which for Jordan regions strengthens Lemma 2.6, while its proof remains relatively simple:

Lemma 2.7 *Given a family \mathcal{F} of pseudo-disks and a finite family \mathcal{B} of pseudo-disks such that each member $B \in \mathcal{B}$ contains a point which is in no other $C \in \mathcal{B}$, the Delaunay-graph of \mathcal{B} wrt. \mathcal{F} is a planar graph.*

Proof We will draw G in the plane in such a way that every pair of edges in G that do not share a common vertex cross an even number of times. The Hanani–Tutte Theorem [7,30] then implies the planarity of G . In our case the edges may have self-crossings too but the Hanani–Tutte theorem holds in this case as well as we can easily redraw the edges to avoid self-crossings (for further details see, e.g., [11]).

For every $v_B, B \in \mathcal{B}$ choose a point $p_B \in B$ which is in no other $C \in \mathcal{B}$. For an illustration of the rest of the proof see Fig. 1.

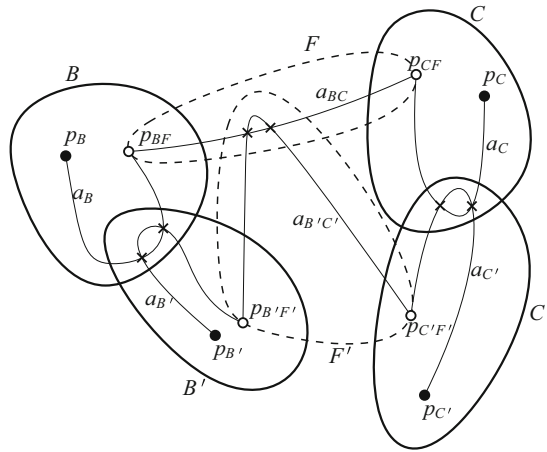


Fig. 1 The edges e_f and $e_{f'}$ intersect an even number of times

If v_B and v_C are connected by an edge f in G , then e_f , the drawing of the edge f , connecting p_B and p_C , is as follows. Let $F \in \mathcal{F}$ be a pseudo-disk corresponding to f . Draw an arc a_B inside B from p_B to a point $p_{BF} \in B \cap F$ ($p_{BF} \in C$ is allowed, also p_{BF} may coincide with p_B). Similarly draw an arc a_C inside C from p_C to a point $p_{CF} \in C \cap F$ ($p_{CF} \in B$ is allowed and p_{CF} may coincide with p_C). Finally, draw an arc a_{BC} inside F from p_{BF} to p_{CF} (p_{BF} and p_{CF} may also coincide in which case a_{BC} is of length zero). The concatenation of these three arcs is the drawing e_f of the edge f between the points p_B and p_C . Note that e_f may have self-crossings.

We are left to prove that in this drawing of G every pair of edges that do not share a vertex cross an even number of times.

Let $B, C, B', C' \in \mathcal{B}$ with edges f defined by F between v_B and v_C and f' defined by F' between $v_{B'}$ and $v_{C'}$. Suppose $e_f = a_B \cup a_{BC} \cup a_C$ and $e_{f'} = a_{B'} \cup a_{B'C'} \cup a_{C'}$ are the drawings of the two edges. Notice that the two endpoints of a_B are in $B \setminus B'$ and the two endpoints of $a_{B'}$ are in $B' \setminus B$. Thus using Lemma 2.5 we get that a_B with $a_{B'}$ intersects an even number of times. The same way we get that a_B with $a_{C'}$, a_C with $a_{B'}$ and a_C with $a_{C'}$ intersect an even number of times. As a_{BC} (resp. $a_{B'C'}$) is disjoint from B' and C' (resp. B and C), there is no intersection between a_{BC} and $a_{B'}$, $a_{C'}$ nor between $a_{B'C'}$ and a_B , a_C . Finally, the two endpoints of a_{BC} are in $F \setminus F'$ and the two endpoints of $a_{B'C'}$ are in $F' \setminus F$ thus again using Lemma 2.5 we get that a_{BC} and $a_{B'C'}$ intersect an even number of times. These together imply that indeed e_f and $e_{f'}$ intersect an even number of times, as required. \square

We note that Pach and Sharir [23] proved that (among others) pseudo-disks have linear union complexity using a similar approach as the proof of Lemma 2.7, connecting own-points of intersecting regions along their boundaries.

Corollary 2.8 *Given a family \mathcal{F} of pseudo-disks and a finite family \mathcal{B} of pseudo-disks, the restricted Delaunay-graph of \mathcal{B} wrt. \mathcal{F} is a planar graph.*

Proof We can delete a member of \mathcal{B} from \mathcal{B} which corresponds to a degree-0 vertex in the restricted Delaunay-graph and then the restricted Delaunay-graph of the new family contains the original restricted Delaunay-graph as a subgraph apart from the degree-0 vertex (which does not alter planarity). We can keep doing this until possible, thus we can assume that there are no degree-0 vertices. In this case for every $B \in \mathcal{B}$ we have a point $p_B \in B$ which is in no other $C \in \mathcal{B}$. Indeed, taking an edge f of G incident to v_B and the corresponding $F \in \mathcal{F}$, by definition of the restricted Delaunay-graph every point of $B \cap F$ is a point which is in B but in no other $C \in \mathcal{B}$. Thus we can apply Lemma 2.7 to conclude that the Delaunay-graph (and so the restricted-Delaunay-graph as well) is planar. \square

Observation 2.9 *If \mathcal{F} is shrinking-closed (for \mathcal{B}), then for every $F \in \mathcal{F}$ the corresponding hyperedge $H = \{v_B : B \cap F \neq \emptyset\}$ of $\mathcal{I}(\mathcal{B}, \mathcal{F})$ either contains an edge of the restricted Delaunay-graph or F contains a point contained in at least two members of \mathcal{B} .*

Corollary 2.8 shows that Lemma 2.7 takes care of the planarity of the restricted Delaunay-graph. From Observation 2.9 we may think that we are left to take care of hyperedges that contain a point that is contained in at least two members of \mathcal{B} . This intuition turns out to be good, in all of our main results Lemma 2.7 will essentially reduce the problem to regarding the intersection hypergraph of \mathcal{B} wrt. all points of the plane (instead of \mathcal{B} wrt. \mathcal{F}).

Before starting the proof of Theorem 1.16 we make some more preparations:

Definition 2.10 Given a family of regions \mathcal{B} , a point is k -deep if it is contained in exactly k members of \mathcal{B} . We denote by ∂B the boundary of some region B of \mathcal{B} . We call a point which is in B but no other $C \in \mathcal{B}$, an *own-point* of B .

Definition 2.11 A hypergraph \mathcal{H}' supports another hypergraph \mathcal{H} if they are on the same vertex set and for every hyperedge $H \in \mathcal{H}$ there exists a hyperedge $H' \in \mathcal{H}'$ such that $H' \subseteq H$.

Observation 2.12 *If \mathcal{H}'' supports \mathcal{H}' and \mathcal{H}' supports \mathcal{H} , then \mathcal{H}'' supports \mathcal{H} .*

Observation 2.13 *If a hypergraph \mathcal{H}' supports another hypergraph \mathcal{H} , then the Delaunay-graph of \mathcal{H} is a subgraph of the Delaunay-graph of \mathcal{H}' .*

Our last tool is the following theorem of Snoeyink and Hershberger (we write here a special case of what they called as the Sweeping theorem):

Theorem 2.14 ([29]) *Let Γ be a finite set of bi-infinite curves in the plane such that any pair of them intersects at most twice. Let d be a closed curve which intersects at most twice every curve in Γ . We can sweep d so that every member of the sweeping of d intersects at most twice every other curve of Γ .*

A sweeping of d in Theorem 2.14 is defined as a family $d(t)$, $t \in (-1, 1]$, of pairwise disjoint curves such that $d(0) = d$, their union contains all points of the plane and $d(\alpha)$ contains $d(\beta)$ for every $1 \leq \beta < \alpha < 1$ ($d(-1)$ is a degenerate curve consisting of a singular point). We note that Theorem 2.14 is stated in [29] for smooth curves but considering instead Jordan curves is not a significant difference.

Proof of Theorem 1.16 Using Corollary 2.3 we can assume that \mathcal{F} is shrinking-closed (for \mathcal{B}) and in particular, for every point p which is in some $F \in \mathcal{F}$, $H_p = \{v_B : p \in B \in \mathcal{B}\}$ is a hyperedge of $\mathcal{I}(\mathcal{B}, \mathcal{F})$ if it is of size at least two.

We can keep deleting members of \mathcal{B} from \mathcal{B} which correspond to a vertex with degree 0 or 1 in the Delaunay-graph, during this process the Delaunay-graph of the new family keeps containing the graph induced by the original Delaunay-graph on the new (reduced) vertex set. If we can prove that the new Delaunay-graph is planar, then adding back the degree 0 and 1 vertices in reverse order we see that the original Delaunay-graph is also planar. Thus we can assume that every vertex of the Delaunay-graph has degree at least two. In this case for every $B \in \mathcal{B}$ we have a point in B which is in at most one other $C \in \mathcal{B}$ (as there are no degree-0 vertices) and there are no two regions $B, C \in \mathcal{B}$ such that $B \subset C$ (as there are no degree-0 or degree-1 vertices).

Given \mathcal{B} and \mathcal{F} we will modify \mathcal{B} so that for the new family $\widehat{\mathcal{B}}$ of pseudo-disks every $B \in \widehat{\mathcal{B}}$ contains an own-point and furthermore we do this in a way that $\mathcal{I}(\widehat{\mathcal{B}}, \mathcal{F})$ supports $\mathcal{I}(\mathcal{B}, \mathcal{F})$. Using Lemma 2.7 we get that the Delaunay-graph of $\widehat{\mathcal{B}}$ wrt. \mathcal{F} is planar and then by Observation 2.13 we get that the Delaunay graph of \mathcal{B} wrt. \mathcal{F} is also planar.

We do this modification by repeating the below defined operation finitely many times, at each time decreasing by at least one the number of members of \mathcal{B} without an own-point. We now define the three steps of this operation.

Step 1: Preparation. Take an arbitrary $B \in \mathcal{B}$ that does not contain an own-point. We take a $C \in \mathcal{B}$ for which v_B is connected to v_C in the Delaunay-graph. Thus, there is a point p which is in $B \cap C$ but no other member of \mathcal{B} . We morph the plane so that $B \cap C$ becomes a square R such that the two intersection points of the boundary of B and C are on the horizontal halving line of the square (the upper part of ∂R belongs to ∂B and the lower part of ∂R belongs to ∂C) and no member of \mathcal{B} (and \mathcal{F}) intersects the vertical sides of the square. This morphing is easily doable and it leaves intact the intersection structure of \mathcal{B} and \mathcal{F} . Denote by p_l and p_r the intersection points of ∂B and ∂C , that is, the midpoints of the left and right side of R . This finishes the Preparation Step of our operation. See the left side of Fig. 2(a) for what we get after this step.

Step 2: Pulling apart $B \cap C$. Now we define the Pulling apart $B \cap C$ Step of the operation which keeps the intersection structure of \mathcal{B} and \mathcal{F} outside $B \cap C$ intact. First we morph the plane so that the square's height is doubled to get the rectangle R' while its horizontal halving line does not change (the part of $\partial R'$ above the halving line is part of ∂B and the part below is part of ∂C) and the drawing inside R remains untouched. We do this such that the intersections of members of \mathcal{B} (and of \mathcal{F}) with the horizontal sides of the square are stretched to become vertical lines in $R' \setminus R$. Next, for every $D \in \mathcal{B} \setminus \{B\}$ which intersects exactly one (horizontal) side of R' , we redraw the part of ∂D inside B to be a half-circle inside B that has the same endpoints. See the right side of Fig. 2(a), where these redrawn boundary parts are drawn with thin red strokes. We get a modified family \mathcal{B}' , while the topology of \mathcal{F} is unmodified.

It is easy to see that we modified \mathcal{B} in a way that \mathcal{B}' remains a pseudo-disk family. Next we show that $\mathcal{I}(\mathcal{B}', \mathcal{F})$ supports $\mathcal{I}(\mathcal{B}, \mathcal{F})$, that is for every $H \in \mathcal{I}(\mathcal{B}, \mathcal{F})$ there exists an $H' \in \mathcal{I}(\mathcal{B}', \mathcal{F})$ such that $H' \subseteq H$. If $F \in \mathcal{F}$ contains only depth-1 points, then it is disjoint from $B \cap C$ and so the hyperedge H_F remains in the intersection

hypergraph after pulling apart $B \cap C$. Otherwise, F contains a depth-2 point q and then H_F contains the hyperedge H_q . So it is enough to prove that for every H_q there is a hyperedge H' (also corresponding to some point) in $\mathcal{I}(\mathcal{B}', \mathcal{F})$ such that $H' \subseteq H_q$. For $q \notin B \cap C$, H_q remains in the intersection hypergraph after pulling apart $B \cap C$. Finally, for $q \in B \cap C$ the hyperedge H_q contains the hyperedge $\{v_B, v_C\}$ which is exactly the hyperedge corresponding to p in \mathcal{B}' (recall that this is the point that was only in B and C and no other member of \mathcal{B} , and this property remains true for \mathcal{B}' after the operation).

Step 3: Shrinking of C . Now we do the final step, the Shrinking of C Step of the operation. Let arc a_B (resp. a_C) be the part of $\partial R'$ above (resp. below) the halving line, which arc is also part of ∂B (resp. ∂C) after the first two steps of the operation. Let arc a^* (resp. arc a_*) be the part of ∂R which is above (resp. below) the halving line. We perturb the vertical parts of a^* and a_* slightly so that they do not overlap (nor intersect) with a_B and a_C . See right side of Fig. 2(a) for an illustration.

If either a^* or a_* contains a point p which is 2-deep (thus contained by B and C and no other member of \mathcal{B}'), then we redraw ∂C so that we change a_C to a^* or a_* (whichever contains the 2-deep point p) and then what we get remains a pseudo-disk family \mathcal{B}'' . This way the hyperedge $\{v_B, v_C\}$ is still in $\mathcal{I}(\mathcal{B}'', \mathcal{F})$ as it is corresponding to a point close to p . Thus $\mathcal{I}(\mathcal{B}'', \mathcal{F})$ supports $\mathcal{I}(\mathcal{B}', \mathcal{F})$, and so $\mathcal{I}(\mathcal{B}, \mathcal{F})$ as well. Furthermore, in \mathcal{B}'' there is a point close to p which is an own-point of B .

Unfortunately it can happen that none of a^* and a_* contains a 2-deep point and so they are not suitable for redrawing a_C . Nevertheless, we assumed that there exists a 2-deep point p in B before the operation, which remains 2-deep after the first two steps of the operation. In the following we find an arc a_p connecting p_l and p_r that goes inside R , goes through p and intersects exactly once every maximal part inside R of the boundary of a member of \mathcal{B}' , see Fig. 2(b). Assuming we have this arc a_p we can redraw ∂C such that we change a_C to a_p and what we get remains a pseudo-disk family \mathcal{B}'' . Also, the same way as in the previous case, we have that $\mathcal{I}(\mathcal{B}'', \mathcal{F})$ supports $\mathcal{I}(\mathcal{B}, \mathcal{F})$ and in \mathcal{B}'' there is a point close to p which is an own-point of B .

Step 3b: Drawing a_p . While the existence of such an a_p is intuitively not surprising, proving its existence is a somewhat technical application of Theorem 2.14. Take the maximal parts inside R of the boundaries of members of $\mathcal{B}' \setminus \{B, C\}$. Extending the top and bottom of these with vertical half-lines we get a family Γ of bi-infinite curves that pairwise intersect at most twice. Let $d = a_B \cup a^*$, it intersects every curve of Γ exactly twice. Now we can apply Theorem 2.14 to this family of curves Γ to sweep d . In particular, we get a closed curve d_p through p which intersects every member of Γ at most twice, see Fig. 2(b). Notice that it must intersect once every member of Γ above the top side of R , thus it can intersect at most once (and easy to see that actually exactly once) every member of Γ below the top side of R . Furthermore, d_p must intersect a_* at least twice. Take the maximal connected part d'_p of $d_p \setminus a_*$ which contains p . Take the curve a_p which goes from p_l along a_* to the closest endpoint of d'_p on a_* , continues with d'_p and then goes again along a_* to p_r . It is easy to see that a_p also intersects every member of Γ exactly once, as required.⁶

⁶ The existence of such an a_p can be proved in various ways. Instead of making the curves bi-infinite we can close them up so that they become boundaries of pseudo-disks all containing p_l . Then by a result of Agarwal

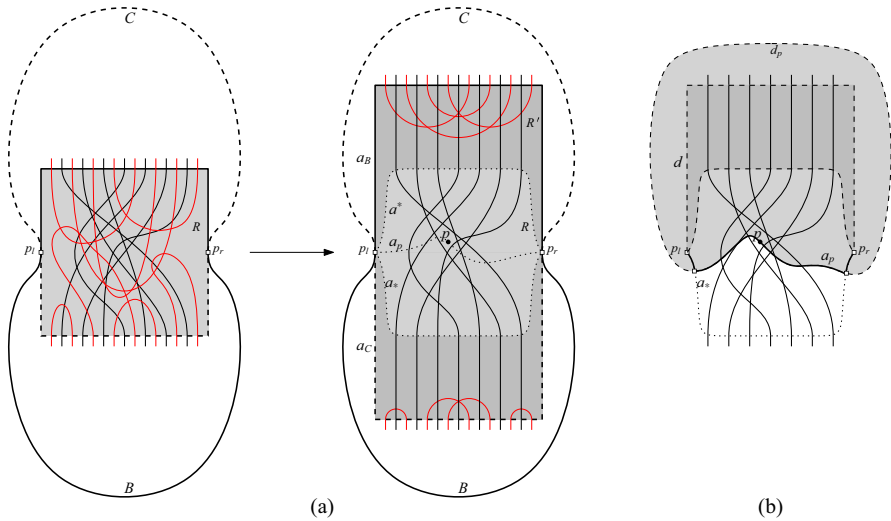


Fig. 2 Major steps of the operation used in the proof of Theorem 1.16: **a** pulling apart $B \cap C$, **b** drawing a_p

Thus, with the above 3-step operation we changed \mathcal{B} so that one more member, B , contains an own-point. Also it is easy to see that no other member could have lost its own-point, as we have redrawn members of \mathcal{B} only inside $B \cap C$, where there were no 1-deep points. So after finitely many times repeating this operation we get a family $\widehat{\mathcal{B}}$ in which every member has an own-point and whose intersection hypergraph $\mathcal{I}(\widehat{\mathcal{B}}, \mathcal{F})$ supports the intersection hypergraph $\mathcal{I}(\mathcal{B}, \mathcal{F})$ we started with, concluding the proof. \square

3 Proper Coloring Pseudo-Disks wrt. Pseudo-Disks

Note that Observation 1.5 can be rephrased equivalently so that if the Delaunay-graph of \mathcal{H} supports \mathcal{H} , then a proper coloring of the Delaunay-graph is a proper coloring of \mathcal{H} . More generally, it is true that:

Observation 3.1 *If a hypergraph \mathcal{H}' supports another hypergraph \mathcal{H} , then a proper coloring of \mathcal{H}' is a proper coloring of \mathcal{H} .*

Proof of Theorem 1.15 Given the pseudo-disk family \mathcal{F} and the finite pseudo-disk family \mathcal{B} , we want to color the vertices of $\mathcal{I}(\mathcal{B}, \mathcal{F})$ corresponding to the members of \mathcal{B} so that for every $F \in \mathcal{F}$ the hyperedge H_F in \mathcal{B} containing exactly those v_B for which $B \cap F \neq \emptyset$ is not monochromatic (assuming that $|H_F| \geq 2$). Using Corollary 2.3 and Observation 3.1 we can assume that \mathcal{F} is shrinking-closed (for \mathcal{B}).

Footnote 6 continued

et al. [2] these pseudo-disks are star-shaped (we omit the definition) which implies that there is a half-infinite curve (a ‘ray’) starting at p_l and going through p intersecting all pseudo-disk boundaries once, which in turn implies the existence of the required a_p as before.

Taking the direct product of a constant coloring provided by Theorem 1.11 and a 4-coloring of the restricted Delaunay-graph which is planar by Corollary 2.8, it is not hard to see that we get a required coloring with a constant number of colors (this idea appears also in the proof of Theorem 1.25). Furthermore, as we mentioned in the introduction, even the existence of a proper 6-coloring follows immediately from Theorem 1.16 using known frameworks (e.g., [14], we utilize this framework also in the fourth paragraph of the proof of Theorem 1.25). However, with the operation introduced in the proof of Theorem 1.16 in hand we can show even the existence of a proper 4-coloring, which is the optimal number of colors (see the end of this section).

We prove the existence of a 4-coloring by induction first on the size of \mathcal{B} , second on the size of the largest containment-minimal hyperedges and third on the number of largest containment-minimal hyperedges. We say that a hypergraph \mathcal{H}' is *better* than a hypergraph \mathcal{H} (on the same vertex set) and that \mathcal{H} is *worse* than \mathcal{H}' if either the size of the largest containment-minimal hyperedges is smaller in \mathcal{H}' than in \mathcal{H} or they are of the same size but there are more of them in \mathcal{H} .

If \mathcal{B} contains two pseudo-disks B, C such that $B \subset C$, then we can proper color $\mathcal{I}(\mathcal{B} \setminus \{B\}, \mathcal{F})$ by induction and then color v_B with a color different from the color of v_C to get a proper coloring of $\mathcal{I}(\mathcal{B}, \mathcal{F})$ as every hyperedge which contains v_B must contain v_C as well. Thus, we can assume that no two pseudo-disks in \mathcal{B} contain one another.

We can assume that \mathcal{F} is shrinking-closed (for \mathcal{B}) as adding non-maximal hyperedges cannot make a hypergraph worse (and also cannot decrease its chromatic number). In particular, we can assume that for every point p which is in some $F \in \mathcal{F}$, $\mathcal{I}(\mathcal{B}, \mathcal{F})$ contains a hyperedge $H_p = \{v_B : p \in B \in \mathcal{B}\}$ (if it is of size at least two).

Take one of the largest containment-minimal hyperedges, $H \in \mathcal{I}(\mathcal{B}, \mathcal{F})$.

If H is of size 2 that means that the Delaunay-graph supports $\mathcal{I}(\mathcal{B}, \mathcal{F})$ and then by Theorem 1.16 we can color it with four colors, which by Observation 3.1 is a proper coloring of $\mathcal{I}(\mathcal{B}, \mathcal{F})$, as required. This starts the induction.

Otherwise, H has $l \geq 3$ vertices. Assume by induction that every intersection hypergraph of pseudo-disks wrt. pseudo-disks better than $\mathcal{I}(\mathcal{B}, \mathcal{F})$ admits a proper 4-coloring.

If H corresponds to some F that contains only 1-deep points, then using Corollary 2.4 and that \mathcal{F} is shrinking-closed we get that H is not containment-minimal, a contradiction. Thus, F contains a point which is at least 2-deep and then using that H is containment-minimal and \mathcal{F} is shrinking-closed there must be a point p (any point inside a pseudo-disk corresponding to H is such) for which actually $H = H_p$. Take two members $B, C \in \mathcal{B}$ for which $v_B, v_C \in H_p$. On $B \cap C$ we do the first two steps of the operation used in the proof of Theorem 1.16 to get a new pseudo-disk family \mathcal{B}' . As we have seen, the intersection hypergraph $\mathcal{I}(\mathcal{B}', \mathcal{F})$ supports $\mathcal{I}(\mathcal{B}, \mathcal{F})$ and thus $\mathcal{I}(\mathcal{B}', \mathcal{F})$ is not worse than $\mathcal{I}(\mathcal{B}, \mathcal{F})$. If $B \cap C$ in \mathcal{B}' contains a 2-deep point, then $\{v_B, v_C\}$ is a hyperedge of $\mathcal{I}(\mathcal{B}', \mathcal{F})$ and the number of size- l hyperedges did actually decrease (as H is not containment minimal anymore) and so $\mathcal{I}(\mathcal{B}', \mathcal{F})$ is better than $\mathcal{I}(\mathcal{B}, \mathcal{F})$. Then by induction $\mathcal{I}(\mathcal{B}', \mathcal{F})$ can be colored with four colors, which by Observation 3.1 is a proper coloring of $\mathcal{I}(\mathcal{B}, \mathcal{F})$ as well.

The only case left is when $B \cap C$ still does not contain a 2-deep point. In this case similarly to Step 3 of the operation in the proof of Theorem 1.16, we redraw ∂C so

that we change a_C to a^* (see again the right side of Fig. 2(a)). In the new family \mathcal{B}' every point in R is still covered at least twice, and nothing changed outside R , thus $\mathcal{I}(\mathcal{B}', \mathcal{F})$ supports $\mathcal{I}(\mathcal{B}', F)$ (and then in turn it supports $\mathcal{I}(\mathcal{B}, \mathcal{F})$). Moreover, the hyperedge corresponding to p is a subset of H_p but it does not contain v_C anymore, so it is a proper subset of H and is of size at least 2. This implies that the number of size- l hyperedges did decrease. Then again $\mathcal{I}(\mathcal{B}', \mathcal{F})$ is better than $\mathcal{I}(\mathcal{B}, \mathcal{F})$ and then by induction can be colored with four colors, which by Observation 3.1 is a proper coloring of $\mathcal{I}(\mathcal{B}, \mathcal{F})$ as well. \square

Observe that already for coloring points wrt. disks we may need four colors (as there are points whose intersection hypergraph wrt. disks is a complete graph on four vertices), so 4 is also a lower bound for coloring pseudo-disks wrt. pseudo-disks.

Without going into details, we note that from the proofs we can create (low-degree) polynomial time coloring algorithms, supposing that initially we are given reasonable amount of information (e.g., the whole structure of the arrangement (all the vertices, edges and faces) of the pseudo-disks of \mathcal{B} and also for every hyperedge H in $\mathcal{I}(\mathcal{B}, \mathcal{F})$ we are given an F corresponding to H).

4 Consequences

First we prove Corollary 1.17 about conflict-free coloring the intersection hypergraph of pseudo-disks wrt. pseudo-disks.

Proof of Corollary 1.17 As mentioned already, the proof uses the by now standard framework of [27]. Given the family \mathcal{B} containing n regions, the first step is to color it properly wrt. \mathcal{F} using c dummy colors using Theorem 1.15 ($c = 4$). Take the biggest color class \mathcal{C}_1 and give it color 1. We continue with $\mathcal{B}_1 = \mathcal{B} \setminus \mathcal{C}_1$. In the s 'th step we start with the subfamily \mathcal{B}_{s-1} and we color it properly wrt. \mathcal{F} using c dummy colors and take the biggest color class \mathcal{C}_s and give it color s and we continue the process with $\mathcal{B}_s = \mathcal{B}_{s-1} \setminus \mathcal{C}_s$. We repeat this until we color all members of \mathcal{B} . In each step the remaining family is reduced by a ratio at least $(c - 1)/c$ and so we finish in $O(\log n)$ steps. It is easy to check that this coloring is indeed conflict-free, as in every hyperedge H_F defined by a region $F \in \mathcal{F}$ take the highest color s appearing on its vertices (corresponding to members of \mathcal{B}). This color must appear only on one vertex otherwise at step s the hyperedge defined by F would have been a monochromatic hyperedge with at least two vertices in the dummy coloring, contradicting the fact that this dummy coloring was proper. \square

Now we continue with the results related to the VC-dimension of the intersection hypergraph of pseudo-disks wrt. pseudo-disks. In order to do this, we assume that the reader is familiar with the definitions of VC-dimension and shattering. As we mentioned, Aronov et al. [3] proved that given a pseudo-disk family \mathcal{F} and a finite subfamily \mathcal{B} of \mathcal{F} , the intersection hypergraph $\mathcal{I}(\mathcal{B}, \mathcal{F})$ has VC-dimension at most 4 and this bound is tight. We now show that the same way Lemma 2.7 implies that $\mathcal{I}(\mathcal{B}, \mathcal{F})$ has VC-dimension at most 4 when \mathcal{B} is any finite pseudo-disk family, not necessarily a subfamily of \mathcal{F} :

Proof of Theorem 1.19 The proof is essentially the same as in [3], except that we use Lemma 2.7. In a shattered subset V_s of vertices in $\mathcal{I}(\mathcal{B}, \mathcal{F})$ by definition every subset of V_s is a hyperedge. Thus, for every vertex $v_B \in V_s$, $\{v_B\}$ is a singleton hyperedge defined by some $F_B \in \mathcal{F}$ and so the points in $B \cap F_B$ must be own-points of B . Also, all pairs of vertices in V_s must define a Delaunay-edge by definition of shattering, while the Delaunay-graph is planar by Lemma 2.7. In fact, Theorem 1.16 also implies that this Delaunay-graph is planar, without using the fact that the singleton hyperedges guarantee the existence of own-points.

As a planar complete graph has at most four points, $|V_s| \leq 4$, as required.

Finally, the construction of [4] showing that the bound was tight already when \mathcal{B} is a subfamily of \mathcal{F} shows that it is still tight in our more general setting. \square

Theorem 1.20 follows directly from the arguments of [3], except that one needs to use Theorem 1.19 instead of the weaker version the authors used (where \mathcal{B} was a subfamily of \mathcal{F}). The details are left to the interested reader.

5 Coloring a Family with Linear Union Complexity wrt. Pseudo-Disks

As in the literature in the definition of union complexity sometimes vertices and other times edges of the arrangement are counted, we first show that this does not affect the property of having linear union complexity (apart from slightly changing the constant c in having c -linear union complexity):

Lemma 5.1 *Given a family \mathcal{B} of n Jordan regions, whose boundary contains $e(\partial\mathcal{B})$ edges and $v(\partial\mathcal{B})$ vertices of the arrangement, $v(\partial\mathcal{B}) \leq e(\partial\mathcal{B}) \leq v(\partial\mathcal{B}) + n$.*

Proof Going clockwise around the closed Jordan curves forming $\partial\mathcal{B}$, the boundary of \mathcal{B} , we see that every vertex of the arrangement is followed by a (different) edge of the arrangement, hence $v(\partial\mathcal{B}) \leq e(\partial\mathcal{B})$. On the other hand every edge is preceded by a (different) vertex except the ones that alone form a closed Jordan curve, that is, that form the boundary of a single region in \mathcal{B} . There are at most n such edges, thus $e(\partial\mathcal{B}) \leq v(\partial\mathcal{B}) + n$. \square

Before proving Theorem 1.25 we make some preparations. First we prove an easy lemma using random sampling, that is, using the Clarkson–Shor method [8]. We say that a face is k -deep if its interior points are k -deep, that is, the face is contained in exactly k members of the family.

Lemma 5.2 *In a family \mathcal{B} of n Jordan regions with linear union complexity, the number of k -deep faces is $O_k(n)$.*

Proof In the arrangement of \mathcal{B} , associate with each k -deep face T one of the edges e_T on its boundary. No edge is associated with two faces this way.

We count the number s of associated edges e_T . Take a random subfamily \mathcal{B}' of \mathcal{B} by taking each $B \in \mathcal{B}$ with probability $1/2$. With probability at least $1/2^{k+3}$ an edge e_T counted in s ends up on $\partial\mathcal{B}'$, the boundary of \mathcal{B}' . Indeed, this happens whenever the at most k regions that contain this edge in their interior are not in \mathcal{B}' but the region

whose boundary contains e_T and two other regions whose boundaries go through the two endvertices of e_T are in \mathcal{B}' . Thus $1/2^{k+3}s \leq E(e(\partial\mathcal{B}')) \leq E(c|\mathcal{B}'|) \leq cn$ implying that $s = O_k(n)$. The second inequality followed from our assumption that \mathcal{B} has c -linear union complexity.

The number of k -deep faces equals the number s of associated edges, which is $O_k(n)$, concluding the proof. \square

It will be comfortable to use the following definition:

Definition 5.3 A graph G is k -degenerate if every subgraph of G contains a point with degree at most k .

If G is k -degenerate for some constant k we also simply say that G is degenerate. For example planar graphs are 5-degenerate. The following is a well-known property of degenerate graphs:

Observation 5.4 If a graph G is k -degenerate, then we can proper color it with $k + 1$ colors.

Proof Take a smallest degree vertex v of G and proper color with $k + 1$ colors by induction the vertices of the subgraph induced by the rest of the vertices. Finally, color v with a color different from the color of its at most k neighbors. \square

Proof of Theorem 1.25 Given the family \mathcal{B} of n regions with bounded union complexity and the family \mathcal{F} of pseudo-disks, we want to color the vertices of $\mathcal{I}(\mathcal{B}, \mathcal{F})$ corresponding to the members of \mathcal{B} so that for every $F \in \mathcal{F}$ the hyperedge H_F in \mathcal{B} containing those v_B for which $B \cap F \neq \emptyset$ is not monochromatic whenever it is of size at least 2.

Using Corollary 2.3 we can assume that \mathcal{F} is shrinking-closed (for \mathcal{B}). In particular, for every point p which is in some $F \in \mathcal{F}$, $H_p = \{v_B : p \in B \in \mathcal{B}\}$ is a hyperedge of $\mathcal{I}(\mathcal{B}, \mathcal{F})$ if it is of size at least two.

First we color \mathcal{B} wrt. points, that is, we color the vertices v_B with c_1 many colors for some constant c_1 so that for an arbitrary point p in the plane the hyperedge $H_p = \{v_B : p \in B \in \mathcal{B}\}$ is not monochromatic whenever it is of size at least 2. Note that this is the same as Theorem 1.24, we include its proof for the sake of completeness.

By Lemma 5.2 there are linearly many different H_p 's with $|H_p| = 2$, in other words the Delaunay-graph of the regions of \mathcal{B} wrt. points has linearly many edges. Thus there is a vertex v_B with degree bounded by a constant. Color by induction the Delaunay-graph of the family we get by deleting B from \mathcal{B} with constantly many colors and then color v_B with a color different from the colors of its neighbors in the Delaunay-graph of \mathcal{B} . We claim that this is a coloring in which each H_p of size at least two is not monochromatic. By induction this is true for every H_p that contains at least two vertices besides v_B . Thus only H_p 's which are of size exactly two and contain v_B may be monochromatic at this point. However, we colored v_B exactly in a way to make these hyperedges also non-monochromatic.⁷

⁷ In order to proper color \mathcal{B} wrt. points, it is tempting to simply proper color the Delaunay-graph by Observation 5.4. This would not be correct, as the Delaunay-graph might not support the intersection hypergraph of \mathcal{B} wrt. points.

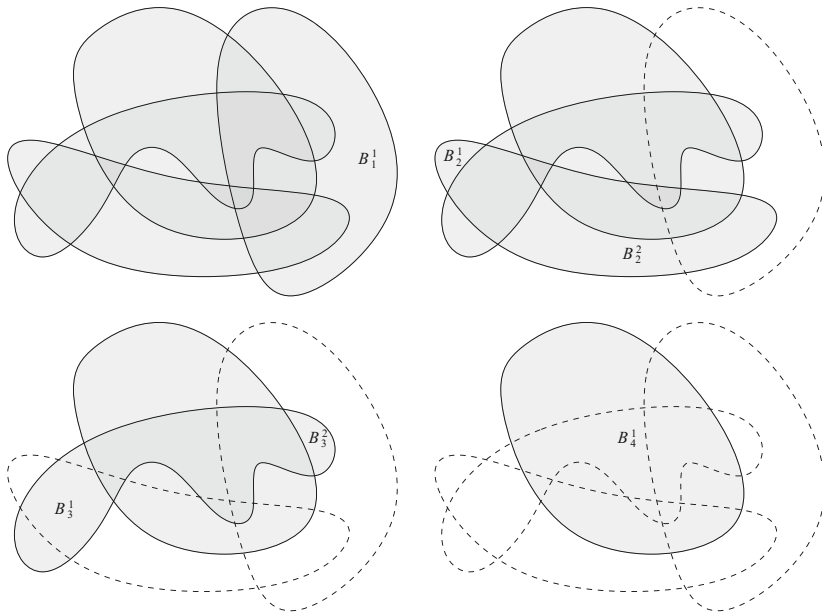


Fig. 3 Steps of the construction of the disjoint subregions

Take an arbitrary $F \in \mathcal{F}$ such that H_F is of size at least 2.

If $F \in \mathcal{F}$ contains a point p which is in at least two members of \mathcal{B} , then H_p is a subset of H_F and is not monochromatic, thus H_F is also not monochromatic.

Thus, we only need to take care of the hyperedges defined by $F \in \mathcal{F}$ that do not contain any point p that is in at least two members of \mathcal{B} . We do this by defining a second coloring using c_2 many colors.

In this case $F \in \mathcal{F}$ intersects the members of \mathcal{B} in disjoint regions. By Corollary 2.4 and as \mathcal{F} is shrinking-closed, there exists an $F' \in \mathcal{F}$, $F' \subseteq F$ such that F' intersects exactly two members of \mathcal{B} . This means exactly that every such \mathcal{F} contains as a subset an edge of the restricted Delaunay-graph G of \mathcal{B} wrt. \mathcal{F} . Next we prove that G can be colored properly with constantly many colors.

We define subregions of the regions of \mathcal{B} in the following way (see Fig. 3 for an example). By Lemma 5.2 there are linearly many 1-deep faces. Thus, there is a $B_1 \in \mathcal{B}$ which contains only constantly many 1-deep faces. Let $B_1^1, B_1^2, \dots, B_1^{z_1}$ be these faces. Now we repeat this process for $\mathcal{B} \setminus \{B_1\}$ (and define in a $B_2 \in \mathcal{B} \setminus \{B_1\}$ constantly many subregions, etc.) until there are no regions left. As by definition linear union complexity is hereditary, we can apply Lemma 5.2 at every step and so at the end in every region we defined constantly many subregions (with the same constant), we collect these into the family \mathcal{B}' .⁸ Observe that the regions in \mathcal{B}' are connected, pairwise disjoint and their union equals the union of the regions in \mathcal{B} (apart from the boundaries). That is, \mathcal{B}' is a partition of the union of the regions in \mathcal{B} (apart from the

⁸ Note that every region in \mathcal{B}' is a face of the arrangement of the subfamily of regions of \mathcal{B} at that moment in the process, but usually a union of more faces of the original arrangement \mathcal{B} .

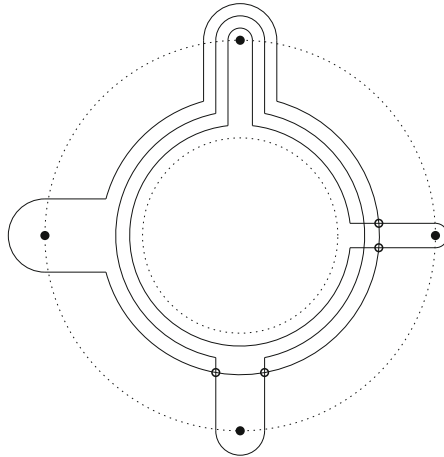


Fig. 4 Construction with $n = 4$ points, three of the $\binom{4}{2}$ regions are shown, their union complexity is 4

boundaries). In particular, every 1-deep face of the arrangement \mathcal{B} is contained in one of the regions in \mathcal{B}' . Notice that \mathcal{B}' has at most constant times n regions.

Now we define a graph G' with the same number of edges as G and with at most constant times n vertices. For every $B_i^j \in \mathcal{B}'$ it has a vertex v_i^j associated with B_i^j . We connect two vertices v_i^j and $v_{i'}^{j'}$ whenever there exists an $F \in \mathcal{F}$ which intersects exactly two members of \mathcal{B} , B_i and $B_{i'}$ and furthermore $F \cap B_i \subseteq B_i^j$ and $F \cap B_{i'} \subseteq B_{i'}^{j'}$. Notice that G' is a subgraph of the restricted Delaunay-graph of \mathcal{B}' wrt. \mathcal{F} and as \mathcal{B}' is a family of disjoint connected regions, by Lemma 2.6 this is a planar graph.

As G' is planar, it has at most three times as many edges as vertices and so it has at most constant times n edges. By definition G and G' have the same number of edges, associated with members $F \in \mathcal{F}$ which intersect exactly two members of \mathcal{B} with disjoint intersections. Thus G also has linearly many edges and so must have a vertex with degree bounded by a constant. As we can apply this argument for any subfamily of \mathcal{B} and any induced subgraph of G is a subgraph of the Delaunay-graph of the family of corresponding members of \mathcal{B} , we get that G is degenerate. Thus by Observation 5.4 it admits a proper coloring with c_2 many colors for some constant c_2 .

The direct product of the above two colorings with c_1 and c_2 many colors is a proper $c_1 c_2$ -coloring of the hypergraph, concluding the proof. \square

Proof of Corollary 1.26 The proof is exactly the same as of Corollary 1.26 just instead of Theorem 1.15 we need to use Theorem 1.25. \square

We note again that the above proof of Theorem 1.25 works without modification also when instead of Jordan-regions (i.e., bounded regions whose boundary is a single closed Jordan curve), \mathcal{B} is a family of bounded regions whose boundary is a union of a finite number of disjoint Jordan curves (in particular the regions are not necessarily simple or connected).

Without going into details, we also note again that from the proofs we can create efficient coloring algorithms.

Now we show that we cannot always color properly points wrt. a family of linear union complexity with constantly many colors and so in particular in Theorem 1.25 \mathcal{F} cannot be changed to be a family of linear union complexity.

Claim 5.5 *For every n there exists a set S of n points and a family \mathcal{F} of $\binom{n}{2}$ simply connected regions with linear union complexity such that $\mathcal{I}(S, \mathcal{F})$ is the complete graph on n vertices. More precisely, for every pair of points in S there is a region in \mathcal{F} that contains exactly these two points, every pair of regions intersects at most four times and for every m the union complexity of any subfamily of size m is at most $4m - 4$.*

Proof The construction is the following: distribute evenly n points $S = \{p_0, \dots, p_{n-1}\}$ on a circle of radius 2 around the origin o . Choose an ϵ small enough. Let $F_{\{i,j\}}$ be the region we get by taking the union of the following three regions: a disk of radius $1 + (in + j)\epsilon$, the Minkowski sum of the segment op_i and a disk of radius $(in + j)\epsilon$ and the Minkowski sum of the segment op_j and a disk of radius $(in + j)\epsilon$. See Fig. 4.

Notice that $F_{\{i,j\}}$ contains p_i and p_j but no other points of S if ϵ is small enough, thus $\mathcal{I}(S, \mathcal{F})$ is the complete graph, as required.

It is easy to see that two regions can intersect at most four times. For an arbitrary subfamily \mathcal{F}' of m regions take the region $F^0 = F_{\{i,j\}} \in \mathcal{F}'$ for which $in + j$ is largest possible. By the construction every intersection point of boundaries that are on the boundary of $\cup \mathcal{F}'$ is on the boundary of F^0 . As F^0 intersects all other regions in at most four points, we get the upper bound $4m - 4$ on the union complexity of \mathcal{F}' .

We remark that it is not much harder to show that $\min(2n - 4, 4m - 4)$ is also an upper bound on the complexity of \mathcal{F}' . \square

6 Discussion

As mentioned, the primary motivation behind considering such problems lies in applications to conflict-free colorings and to cover-decomposability problems.

When the main interest is in conflict-free colorings, the $\log n$ factor makes it less interesting to optimize the bound on the proper coloring result. However, when considering the relation to cover-decomposability problems, finding the exact value is of great interest. This makes it important that in Theorem 1.15 we could prove an exact upper bound for such a wide class of intersection hypergraphs. On the other hand, in Theorem 1.25 we aimed for keeping the proof as simple as possible, and so we did not optimize the number of colors.

In cover-decomposability problems the typical question asks if for some m we can properly 2-color $\mathcal{I}_m(\mathcal{B}, \mathcal{F})$, the subhypergraph of $\mathcal{I}(\mathcal{B}, \mathcal{F})$ containing only hyperedges of size at least m . If we properly color $\mathcal{I}_m(\mathcal{B}, \mathcal{F})$, then we say that we color \mathcal{B} wrt. \mathcal{F} for m .

Usually either \mathcal{B} or \mathcal{F} is the family of points of the plane. Originally researchers concentrated on the problem when the other family, \mathcal{F} or \mathcal{B} , is the family of translates of a convex region (see, e.g., [25] for a history of this research direction).

In the past years researchers started to concentrate more on problems where \mathcal{B} or \mathcal{F} is the family of homothets of a convex region. One of the the most intriguing open problems is whether we can 2-color points wrt. homothets of a given convex polygon for some constant m . The existence of a 2-coloring for some m was proved for triangles [17] and the square [1] and recently for convex polygons the existence of a 3-coloring for some m was proved [20]. On the other hand, for coloring points wrt. disks two colors are not enough (for any m) [24]. As we have seen already in the introduction, a 4-coloring of points wrt. disks always exists (for $m = 2$), however the existence of a 3-coloring (for some m) is an open problem (see also the remark after Problem 6.1).

For the dual problem, about coloring the homothets of a convex polygon wrt. points, while proper 2-coloring exists for the homothets of triangles [17] for some m , for the homothets of any other convex polygon there is not always a 2-coloring [21] (for any m). Also, for coloring disks wrt. points two colors are not enough (for any m) [22].

In [20] it was conjectured that points wrt. pseudo-disks and pseudo-disks wrt. points can be proper 3-colored (for some m). Encouraged by Theorem 1.15 we pose the common generalization of these conjectures:

Problem 6.1 *Does there exist a constant m such that given a family \mathcal{F} of pseudo-disks and a finite family \mathcal{B} of pseudo-disks, $\mathcal{I}_m(\mathcal{B}, \mathcal{F})$ always admits a proper coloring with three colors?*

We note that Tóth (Personal Communication) showed an example that choosing $m = 3$ is not enough already when \mathcal{B} is a set of points and \mathcal{F} is a family of pseudo-disks.

The case when \mathcal{B} is a family of points and \mathcal{F} is the family of disks was asked by the author of this paper more than 10 years ago [15,16] and is still open (in this special case even $m = 3$ might be true).

As we mentioned in the introduction, one can easily change in the definition of the incidence hypergraph the incidence relation to the containment or reverse-containment relation. Thus we can define the following hypergraphs: $\mathcal{H}_m^{\subset}(\mathcal{B}, \mathcal{F})$ contains a hyperedge $H = \{v_B : B \subset F\}$ for every $F \in \mathcal{F}$; similarly, $\mathcal{H}_m^{\supset}(\mathcal{B}, \mathcal{F})$ contains a hyperedge $H = \{v_B : B \supset F\}$ for every $F \in \mathcal{F}$ (if these sets are of size at least two).

Thus we can ask if the respective variants of Theorem 1.15 and Problem 6.1 hold for these hypergraphs:

Problem 6.2 *Given a family \mathcal{F} of pseudo-disks and a finite family \mathcal{B} of pseudo-disks, does $\mathcal{H}_m^{\subset}(\mathcal{B}, \mathcal{F})$ (resp. $\mathcal{H}_m^{\supset}(\mathcal{B}, \mathcal{F})$) always admit a proper coloring with four colors?*

Does there exist a constant m such that $\mathcal{H}_m^{\subset}(\mathcal{B}, \mathcal{F})$ (resp. $\mathcal{H}_m^{\supset}(\mathcal{B}, \mathcal{F})$) always admits a proper coloring with three colors?

The variant of this problem where \mathcal{F} and \mathcal{B} are families of intervals on the line was regarded in [18] where among others it was shown that for intervals these two classes (the class of containment and reverse-containment hypergraphs of intervals) are the same.

Finally we mention that similarly to pseudo-disks, pseudo-halfplanes are natural generalizations of halfplanes. For pseudo-halfplanes it was also possible to reprove the same coloring results that are known for halfplanes [19]. Also, due to the lack of direct

relation to pseudo-disks, we did not mention axis-parallel rectangles, whose intersection hypergraph coloring problems are some of the most interesting open problems of the area.

Acknowledgements Open access funding provided by MTA Alfréd Rényi Institute of Mathematics (MTA RAMKI). This paper is dedicated to my grandfather who died the same day when I found the proof for Theorem 1.15 (at that time with a non-explicit constant). Thanks to Dömötör Pálvölgyi for the inspiring discussions about these and other problems about coloring geometric hypergraphs defined by pseudo-disks. Thanks to Chaya Keller for many useful remarks about the realization of this paper.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Ackerman, E., Keszegh, B., Vizer, M.: Coloring points with respect to squares. *Discrete Comput. Geom.* **58**(4), 757–784 (2017)
2. Agarwal, P.K., Nevo, E., Pach, J., Pinchasi, R., Sharir, M., Smorodinsky, S.: Lenses in arrangements of pseudo-circles and their applications. *J. ACM* **51**(2), 139–186 (2004)
3. Aronov, B., Donakonda, A., Ezra, E., Pinchasi, R.: On pseudo-disk hypergraphs (2018). [arXiv:1802.08799](https://arxiv.org/abs/1802.08799)
4. Buzaglo, S., Pinchasi, R., Rote, G.: Topological hypergraphs. In: Pach, J. (ed.) *Thirty Essays on Geometric Graph Theory*, pp. 71–81. Springer, New York (2012)
5. Cardinal, J., Korman, M.: Coloring planar homothets and three-dimensional hypergraphs. *Comput. Geom.* **46**(9), 1027–1035 (2013)
6. Chan, T.M., Har-Peled, S.: Approximation algorithms for maximum independent set of pseudo-disks. *Discrete Comput. Geom.* **48**(2), 373–392 (2012)
7. Chojnacki, Ch.: Über wesentlich unplättbare Kurven im dreidimensionalen Raume. *Fundamenta Mathematicae* **23**, 135–142 (1934)
8. Clarkson, K.L., Shor, P.W.: Applications of random sampling in computational geometry. II. *Discrete Comput. Geom.* **4**(5), 387–421 (1989)
9. Even, G., Lotker, Z., Ron, D., Smorodinsky, S.: Conflict-free colorings of simple geometric regions with applications to frequency assignment in cellular networks. *SIAM J. Comput.* **33**(1), 94–136 (2003)
10. Fekete, S.P., Keldenich, P.: Conflict-free coloring of intersection graphs. *Internat. J. Comput. Geom. Appl.* **28**(3), 289–307 (2018)
11. Fulek, R., Pelsmajer, M.J., Schaefer, M., Štefankovič, D.: Adjacent crossings do matter. *J. Graph Algorithms Appl.* **16**(3), 759–782 (2012)
12. Holmsen, A.F., Nassajian Mojarrad, H., Pach, J., Tardos, G.: Two extensions of the Erdős–Szekeres problem (2017). [arXiv:1710.11415](https://arxiv.org/abs/1710.11415)
13. Kedem, K., Livné, R., Pach, J., Sharir, M.: On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles. *Discrete Comput. Geom.* **1**(1), 59–71 (1986)
14. Keller, C., Smorodinsky, S.: Conflict-free coloring of intersection graphs of geometric objects. *Discrete Comput. Geom.* <https://doi.org/10.1007/s00454-019-00097-8>
15. Keszegh, B.: Weak conflict-free colorings of point sets and simple regions. In: Bose, P. (ed.) *Proceedings of the 19th Annual Canadian Conference on Computational Geometry (CCCG’07)*, pp. 97–100. Carleton University, Ottawa (2007)
16. Keszegh, B.: Coloring half-planes and bottomless rectangles. *Comput. Geom.* **45**(9), 495–507 (2012)
17. Keszegh, B., Pálvölgyi, D.: Octants are cover-decomposable. *Discrete Comput. Geom.* **47**(3), 598–609 (2012)
18. Keszegh, B., Pálvölgyi, D.: More on decomposing coverings by octants. *J. Comput. Geom.* **6**(1), 300–315 (2015)

19. Keszegh, B., Pálvölgyi, D.: An abstract approach to polychromatic coloring: Shallow hitting sets in ab-free hypergraphs and pseudohalfplanes. In: Mayr, E.W. (ed.) *Graph-Theoretic Concepts in Computer Science*. Lecture Notes in Computer Science, vol. 9224, pp. 266–280. Springer, Berlin (2015)
20. Keszegh, B., Pálvölgyi, D.: Proper coloring of geometric hypergraphs. In: Aronov, B., Katz, M.J. (eds.) *Proceedings of the 33rd International Symposium on Computational Geometry (SoCG'17)*. LIPIcs. Leibniz International Proceedings in Informatics, vol. 77, pp. 47:1–47:15. Schloss Dagstuhl. Leibniz-Zentrum für Informatik, Wadern (2017)
21. Kovács, I.: Indecomposable coverings with homothetic polygons. *Discrete Comput. Geom.* **53**(4), 817–824 (2015)
22. Pach, J., Pálvölgyi, D.: Unsplittable coverings in the plane. In: Mayr, E.W. (ed.) *Graph-Theoretic Concepts in Computer Science*. Lecture Notes in Computer Science, vol. 9224, pp. 281–296. Springer, Berlin (2015)
23. Pach, J., Sharir, M.: On the boundary of the union of planar convex sets. *Discrete Comput. Geom.* **21**(3), 321–328 (1999)
24. Pach, J., Tardos, G., Tóth, G.: Indecomposable coverings. In: Akiyama, J., et al. (eds.) *Discrete Geometry, Combinatorics and Graph Theory*. Lecture Notes in Computer Science, vol. 4381, pp. 135–148. Springer, Berlin (2007)
25. Pach, J., Pálvölgyi, D., Tóth, G.: Survey on decomposition of multiple coverings. In: Bárány, I., et al. (eds.) *Geometry-Intuitive, Discrete, and Convex*. Bolyai Society Mathematical Studies, vol. 24, pp. 219–257. Springer, Berlin (2013)
26. Pinchasi, R.: A finite family of pseudodiscs must include a “small” pseudodisc. *SIAM J. Discrete Math.* **28**(4), 1930–1934 (2014)
27. Smorodinsky, S.: On the chromatic number of geometric hypergraphs. *SIAM J. Discrete Math.* **21**(3), 676–687 (2007)
28. Smorodinsky, S.: Conflict-free coloring and its applications. In: Bárány, I., et al. (eds.) *Geometry-Intuitive, Discrete, and Convex*. Bolyai Society Mathematical Studies, vol. 24, pp. 331–389. Springer, Berlin (2013)
29. Snoeyink, J., Hershberger, J.: Sweeping arrangements of curves. In: *Proceedings of the 5th Annual Symposium on Computational Geometry (SoCG'89)*, pp. 354–363. ACM, New York (1989)
30. Tutte, W.T.: Toward a theory of crossing numbers. *J. Comb. Theory* **8**(1), 45–53 (1970)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.