Geometric hypergraphs and tangency graphs

Balázs Keszegh

Doctoral dissertation submitted to the
Hungarian Academy of Sciences

Budapest, 2024

Contents

1 Introduction

Combinatorial properties of incidence and intersection structures of geometric objects have been central problems of combinatorial geometry since the beginning of this field. The Sylvester-Gallai theorem states that for a finite set of n points in the Euclidean plane, not all of which lie on a line, there exists a line that passes through exactly two of the points. By induction this implies the de Bruijn–Erdős theorem, which states that if not all points are on a line then they determine at least n lines; see the book chapter [107] for the history of these problems. Solving the problem of Erdős, the influential Szemerédi–Trotter theorem [136] shows that the maximal number of incidences among n points and m lines is $O(n^{\frac{2}{3}}m^{\frac{2}{3}}+n+m)$. There are many variants of this result, in higher dimensions, with more general curves instead of lines, etc. The topics of these results are still central areas of research.

There is a lot of research on graphs that can be realized in some way by geometric objects where by a geometric object we mean a subset of some underlying space, usually $\mathbb{R}^d$. First, one can try to characterize graphs that are realizable in a certain way. Second, one can try to prove results about the class of graphs realizable in a certain way.

The most natural way is to realize a graph in the plane such that its vertices are represented by points and its edges by connecting curves, and then we put additional requirements on the possible intersection structures of these curves. Planar graphs are the most notable example, where no intersections among the curves are allowed. Other examples are k-quasi-planar graphs, k-planar graphs. The branch of mathematics that deals with such classes is called graph drawing.

If vertices and edges are represented in other ways, we get further interesting classes. Such graph classes include intersection graphs, tangency graphs and visibility graphs, among others.

Intersection graphs are perhaps the most studied of these. Given a set of n geometric objects in $\mathbb{R}^d$, its intersection graph is the graph with n vertices corresponding to the geometric objects and we connect two vertices by an edge if the corresponding objects intersect. Many well-known classes are defined as or are equivalent to some class of intersection graphs: interval graphs (intersection graphs of intervals), intersection graphs of subtrees of a tree (which are exactly the chordal graphs [54]), circle graphs (intersection graphs of chords of a circle), intersection graphs of interior-disjoint disks (which are exactly the planar graphs by the Koebe–Andreev–Thurston theorem), string graphs (intersection graphs of curves in the plane), segment intersection graphs, L-shape intersection graphs, T-shape intersection graphs, intersection graphs of unit disks, intersection graphs of axis-parallel boxes.[1] Besides the characterization of such a family, its extremal properties (maximum number of edges under certain additional conditions) and its colorings are the most studied topics. In particular, the chromatic number of a graph, denoted by χ, is the least number of colors in a proper coloring of its vertices, i.e., in which no edge is monochromatic.

For example Scheinerman's conjecture, proved by Chalopin and Gonçalves [28] states that every planar graph is an intersection graph of segments in the plane. However, the latter family is significantly broader, as it contains the complete graph of arbitrary size (while the complete graph on 5 vertices is already not planar).

Concentrating on the chromatic number of intersection graphs, it is usually not possible to give a universal bound on the chromatic number of intersection graphs of a given family, instead, the central question became the χ-boundedness of these families. A family of graphs is χ-*bounded* if there exists a function such that the chromatic number of any graph in the family is

[1] In particular the boxicity of a graph is defined as the minimum dimension in which a given graph can be represented as an intersection graph of axis-parallel boxes.

upper bounded by this function of the size of the maximal clique of the graph. It turns out that intersection graphs of a family of geometric objects are often χ-bounded, most notably of axis-parallel rectangles (Asplund and Grünbaum [16], Chalermsook and Walczak [27]), of chords of a circle (Gyárfás [61], Davies [39]), of homothets of a convex shape in the plane (Kim, Kostochka and Nakprasit [126]), of curves lying in a common halfplane and having an endpoint on the boundary of that halfplane (Rok and Walczak [123]). In many cases, although it is much harder to prove, it is also true that the bounding function may even be a polynomial, these are called polynomially χ-bounded families. Esperet [67] therefore asked whether every hereditary class of graphs (intersection graph classes are usually such) that is χ-bounded is also polynomially χ-bounded, but this was recently refuted by Briański, Davies and Walczak [20]. On the other hand, there are natural families whose intersection graphs are not χ-bounded, most notably of boxes in $\mathbb{R}^3$ (Burling [21]), of boundaries of axis-parallel rectangles and of line segments in the plane (Pawlik et al. [113]). For more on this topic see the survey of Scott and Seymour [125].

1.1 Geometric hypergraphs

We can look at intersections in a slightly different way. If we have two sets of objects, we can concentrate on the bipartite intersection graph defined by pairs of objects, one from each set. This can be regarded as the incidence graph of a hypergraph, whose vertices correspond to the first set of objects, the hyperedges to the other set of objects, such that incidences in the hypergraph are the same as for the corresponding objects. Looking at it this way, it is natural to extend our study to settings where the two sets of objects are different. We call hypergraphs defined this way *intersection hypergraphs*. By *geometric hypergraphs* we usually refer to such intersection hypergraphs, unless stated otherwise. Given two sets of geometric objects, $\mathcal{A}$ and $\mathcal{B}$, let the geometric (intersection) hypergraph $\mathcal{I}(\mathcal{A},\mathcal{B})$ be the hypergraph whose vertices correspond to objects of A, and a subset of the vertices is a hyperedge if and only if there is an object of $\mathcal{B}$ that intersects exactly the objects of $\mathcal{A}$ corresponding to this subset of vertices. We omit hyperedges of size smaller than 2. We refer to $\mathcal{I}(\mathcal{A},\mathcal{B})$ as the hypergraph defined by $\mathcal{A}$ with respect to $\mathcal{B}$ (by $\mathcal{A}$ wrt. $\mathcal{B}$, in short):

Definition 1.1. *Given a family $\mathcal{A}$ of objects and a family $\mathcal{B}$ of objects, the* intersection hypergraph *of $\mathcal{A}$* with respect to *(wrt. in short) $\mathcal{B}$ is the hypergraph $\mathcal{I}(\mathcal{A},\mathcal{B})$ which has a vertex v_A corresponding to every $A \in \mathcal{A}$ and a subset H of the vertices is a hyperedge if and only if there exists a $B \in \mathcal{B}$ such that $H = \{v_A : A \cap B \neq \emptyset\}$ (if this set has size at least two), and we say that H corresponds to B.*

The Delaunay graph *of a hypergraph $\mathcal{H}$ is the graph on the same vertex set containing only the hyperedges of $\mathcal{H}$ of size 2. In particular, we refer to the Delaunay graph of $\mathcal{I}(\mathcal{A},\mathcal{B})$ as the Delaunay graph of $\mathcal{A}$ wrt. $\mathcal{B}$.*

Note that $\mathcal{I}(\mathcal{A},\mathcal{B})$ has no repeated hyperedges by definition. Thus, if $\mathcal{A}$ is finite, then even if $\mathcal{B}$ is infinite, $\mathcal{I}(\mathcal{A},\mathcal{B})$ has finitely many hyperedges. In particular, a hyperedge H can correspond to multiple members of $\mathcal{B}$.

The most natural case is when vertices correspond to points (we refer to such a setting as a *primal problem*) or when the hyperedges correspond to points (we refer to such a setting as a *dual problem*). Note that in these two cases the intersection relation is in fact the same as the containment relation. This is the reason why the notation geometric hypergraphs is preferred to intersection hypergraphs, especially in these two settings, as referring to a point intersecting an object sounds unnatural, even if formally correct. See Figure 1.

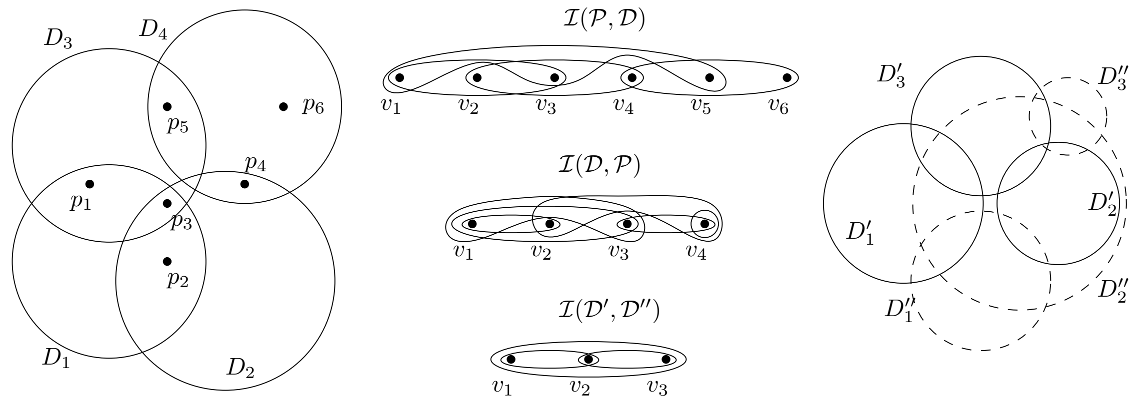

Figure 1: Given a family $\mathcal{P}$ of points and $\mathcal{D}$ of disks, $\mathcal{I}(\mathcal{P},\mathcal{D})$ is the hypergraph in which vertices correspond to points and hyperedges to disks, in $\mathcal{I}(\mathcal{D},\mathcal{P})$ vertices correspond to disks and hyperedges to points. Given two families of disks $\mathcal{D}'$ and $\mathcal{D}''$, in $\mathcal{I}(\mathcal{D}',\mathcal{D}'')$ vertices correspond to disks of $\mathcal{D}'$ and hyperedges correspond to disks of $\mathcal{D}''$.

The primal problem (when vertices correspond to points) is especially natural as in practice hypergraphs are usually represented such that points correspond to vertices and the hyperedges correspond to objects that contain exactly the points corresponding to vertices the hyperedges contain. Given such a representation, the original hypergraph is exactly the primal hypergraph defined by these points and objects. Thus it is very natural to regard hypergraphs that can be represented this way using only simple objects like disks, axis-parallel rectangles, or regions whose boundaries interest only few times, like pseudo-disks, etc. (definitions of these families are postponed).

Yet again, besides extremal questions, most interesting are coloring problems of such hypergraphs. For example, the well-known Delaunay graph of a point set is just the Delaunay graph of these points wrt. all disks in the plane. This is known to be a planar graph, therefore proper 4-colorable. We will see that the Delaunay-graph of a hypergraph in general is in itself an important notion for us, as if the Delaunay graph has only a linear number of edges hereditarily (in particular when it is a planar graph) then this gives a bound on the size and chromatic number of the hypergraph itself.

Given a geometric hypergraph (think about the example where the set of vertices corresponds to a finite set of points and the set of hyperedges to a finite set of disks in the plane), we are mainly interested about its proper and polychromatic colorings.

Definition 1.2. *Given a hypergraph $\mathcal{H}$, a* proper k-coloring *of its vertices is a coloring with k colors such that every hyperedge of $\mathcal{H}$ contains two differently colored vertices. $\chi(\mathcal{H})$ is the smallest possible k for which a proper k-coloring exists. A* polychromatic k-coloring *of the vertices of $\mathcal{H}$ is a coloring in which every hyperedge contains k differently colored vertices.*

Note that for $k = 2$ the notion of a proper 2-coloring and a polychromatic 2-coloring coincides. Further, a proper k-coloring is also a proper $(k+1)$-coloring while merging two colors in a polychromatic $(k+1)$-coloring we get a polychromatic k-coloring. Summarizing, if we increase k, then having a proper k-coloring becomes gradually easier, while having a polychromatic k-coloring becomes harder, so this gives a linear hierarchy of the families that admit such colorings, meeting in the middle at proper and polychromatic 2-colorings:

$\ldots \subseteq$ polychrom. 3-colorable $\subseteq$ polychrom./proper 2-colorable $\subseteq$ proper 3-colorable $\subseteq \ldots$

It is important to note that due to how we defined geometric hypergraphs, it could easily happen that a hyperedge has size one (in our example, a disk may contain a single point), which is always monochromatic, that is why we needed to disregard hyperedges of size smaller than 2. Moreover, for polychromatic colorings there is no hope to polychromatic k-color a hyperedge of size smaller than k, thus we need to disregard hyperedges of size smaller than k when considering polychromatic k-colorings.

It turns out that we get more natural problems if we allow ourselves to disregard hyperedges of size smaller than m for some fixed value m chosen according to the type of geometric objects. In a (primal) *geometric hypergraph polychromatic coloring* problem, we are given a natural number k, a set of points and a collection of objects in $\mathbb{R}^d$, and our goal is to k-color the points such that every object that contains at least $m(k)$ points contains a point of every color, where m is some function that we try to minimize. We call such a coloring a *polychromatic k-coloring*. In a *dual* geometric hypergraph polychromatic coloring problem, our goal is to k-color the objects such that every point which is contained in at least $m(k)$ objects is contained in an object of every color.

Motivated by this, we are ready to define one of the central notions of this dissertation. Given a hypergraph $\mathcal{H}$, let $\mathcal{H}_m$ be the hypergraph on the same set of vertices that contains the hyperedges of H with size at least m. Similarly, for a hypergraph family $\mathcal{F}$ let $\mathcal{F}_m$ be the hypergraph family that contains $\mathcal{H}_m$ for every $\mathcal{H} \in \mathcal{F}$. Finally, we define $\chi_{\mathrm{m}}(\mathcal{F})$ to be the smallest number of colors k with the property that there exists a value m such that every hypergraph in $\mathcal{F}_m$ admits a proper k-coloring[2]. If $\chi_{\mathrm{m}}(\mathcal{F}) = 2$ then the smallest possible such m is denoted by m_2. In this case it makes sense to ask for polychromatic k-colorings for $k > 2$ as well, let m_k be the smallest possible m such that every hypergraph in $\mathcal{F}_m$ admits a polychromatic k-coloring (if no such m exists, we define $m_k = \infty$). Note that m_k is monotone increasing in k.

For infinite families we define the following shorthand notation, used many times during this dissertation. When we talk about infinite sets of geometric objects then (unless stated otherwise) we always mean that for large enough m every finite subfamily has the required property. That is, for $\mathcal{A}$ and $\mathcal{B}$ infinite families of objects we say that $\mathcal{A}$ is proper 2-colorable wrt. $\mathcal{B}$ if there exists a value m such that for any finite subfamily $\mathcal{A}'$ of $\mathcal{A}$ we have that $\mathcal{I}_m(\mathcal{A}', \mathcal{B})$ is proper 2-colorable. We say that for $\mathcal{A}$ wrt. $\mathcal{B}$ we have $m_k \leq f(k)$ if for any finite subfamily $\mathcal{A}'$ of $\mathcal{A}$ we have that $\mathcal{I}_{f(k)}(\mathcal{A}', \mathcal{B})$ admits a polychromatic k-coloring. The smallest possible $f(k)$ is denoted by m_k. For example 'points are 2-colorable wrt. halfplanes' means that there exists an m such that for every finite set of points, the point set is 2-colorable such that every halfplane that contains at least m points is 2-colored. Let us give also a dual example: 'halfplanes are polychromatic k-colorable wrt. points' means that for every k there exists an $f(k)$ such that for every finite set of halfplanes, the halfplanes are k-colorable such that every point that is contained in at least $f(k)$ halfplanes is contained in halfplanes of each of the k colors.

Investigating $\chi_{\mathrm{m}}(\mathcal{F})$ for geometric intersection hypergraph families and the existence of m_k were originally motivated by cover-decomposability problems [99], and the focus was on whether $\chi_{\mathrm{m}}(\mathcal{F}) = 2$, and later also by their connection to so called conflict-free coloring problems [129]. These two problems have real-life motivations, cover-decomposability problems have applications in sensor networks, while conflict-free colorings are used in frequency-assignment problems, among others.

The cover-decomposability problem is the following: given an m-fold covering of (some subset of) the plane (or of $\mathbb{R}^d$) with some geometric objects (e.g., disks), can we split it into 2, or more

[2]Note that the m in χ_{m} is a letter and not a variable, in particular it is not the same m as in the definition.

generally, into k, 1-fold coverings? This is equivalent to k-coloring the objects such that every point of (some subset of) the plane is covered by objects of each color. This is exactly the dual polychromatic k-coloring problem (i.e., when we color objects wrt. points). Moreover, if the objects are translates of some given shape, then it turns out to be also equivalent to the primal polychromatic k-coloring problem (i.e., when we color points wrt. objects).

A conflict-free coloring of a hypergraph is a coloring in which for each hyperedge there is a color which appears exactly once on its vertices. For example if we take n points on a line and the hyperedges are defined by those intervals of the line that contain at least two points, then alternatingly coloring with 2 colors gives a proper 2-coloring of this hypergraph, while a conflict-free coloring needs $\lceil \log n \rceil$ colors and this is enough (give the middle point a unique color and recurse with the same color set on the two sides). There is a general connection between the two notions, and the difference is at most a $\log n$ factor if we have some natural assumptions.

We illustrate the connection between colorings, cover-decomposability problems, conflict-free colorings and real-life applications with an example. We are given a finite set of receivers and a finite set of antennas. The reception range of each antenna is a region. Assume that we want to assign a frequency to each antenna such that every receiver has to be in the range of antennas with at least 3 different frequencies. This is equivalent to polychromatic 3-coloring the ranges wrt. the points where the receivers are. This is also equivalent to splitting the regions into 3 1-fold coverings of this point set. On the other hand, if we require that for every receiver among the antennas whose ranges contain this receiver, there has to be an antenna whose frequency is unique among these, then we get a problem equivalent to conflict-free coloring the ranges wrt. the points.

The main results of the dissertation on this topic are as follows (see later the missing definitions). Theorem 2.12 states that the hypergraph defined by one pseudo-disk family wrt. another pseudo-disk family always admits a proper 4-coloring. This is a common generalization of multiple earlier results. Then Theorem 2.22 extends this result to coloring with constant many colors the hypergraph defined by a family of regions with linear union complexity wrt. a pseudo-disk family. Theorem 2.25 upper bounds the number of hyperedges of size k in the hypergraph defined by a family of lines wrt. a pseudo-disk family. We then define several abstract families of hypergraphs, such as ABA-free, pseudo-halfplane, ABAB-free and prove results about polychromatic coloring pseudo-halfplane hypergraphs (Theorem 2.29 and Theorem 2.30) and proper coloring ABAB-free hypergraphs (Theorem 2.32 and Theorem 2.33), among others. We also show how these abstract families correspond to intersection hypergraphs defined points wrt. pseudo-halfplanes and stabbed pseudo-disks, respectively.

Results in this area will be discussed in Sections 2.1,2.2 and 2.3.

1.2 Tangency graphs

Besides the intersection relation, there are other natural ways to define a graph or hypergraph on geometric objects, based on other relations of the objects, such as visibility graphs. Our main interest is about tangency graphs, whose vertices correspond to geometric objects and we connect two vertices by an edge if and only if the corresponding objects are tangent (i.e., touch each other, for an example see Figure 2). Tangency graphs arise in various seemingly unrelated areas, we mention two highlights. First, the famous Koebe–Andreev–Thurston theorem [86] states that every planar graph is the tangency graph of internally disjoint disks. This result has various surprising applications, for example the author and his coauthors used it to prove bounds on the so called slope number of planar graphs [76]. Another notable example is the unit-

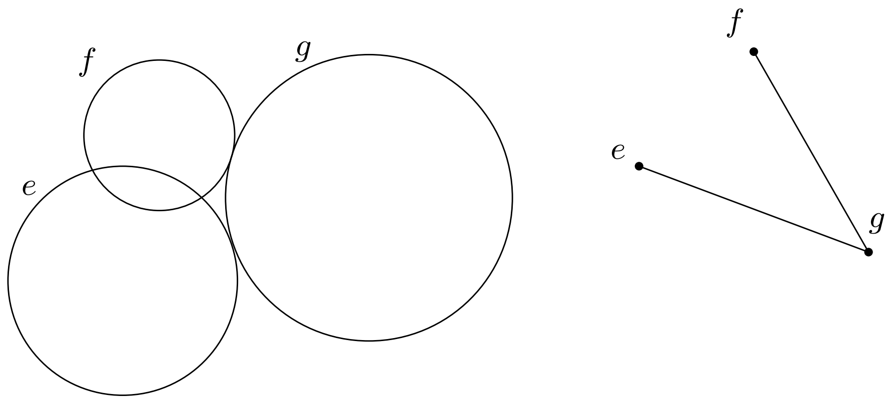

Figure 2: A family of 3 circles and their tangency graph.

distance problem. Erdős asked [42] what is the maximum number of unit distances determined by n points in the plane. Placing a disk of radius 1/2 centered at each point we get that this problem is equivalent to the maximum number of tangencies among congruent (but not necessarily disjoint) disks. While the best known lower bound is only slightly above linear, the best known upper bound is $O(n^{4/3})$. This problem is also related to the famous Hadwiger–Nelson problem concerning the chromatic number of the plane.

We have mentioned earlier that the Delaunay graph of a hypergraph is useful to investigate. Here we give another motivation for this, which in turn connects the area of geometric hypergraphs and tangency graphs in a nice way, which we will indeed apply later in the dissertation. Given a set of curves, whenever two of them are tangent, we can perturb the two curves so that they become locally disjoint and replace the touching point with a small curve (or even a segment) such that this small curve crosses only the two objects that were touching. This way the hypergraph defined by the original curves wrt. these new small curves has hyperedges of size 2 only and the number of hyperedges is equal to the number of tangencies in the original setting. Thus, this hypergraph is equal to its Delaunay graph and bounding the number of edges of this graph bounds the tangencies in the original setting.

The main results of the dissertation on this topic are as follows (see later the missing definitions). Solving a special case of a conjecture of Pach, Theorem 2.39 states that n curves that pairwise intersect exactly once define $O(n)$ tangencies. Theorem 2.40 shows that n curves that pairwise intersect at most t times define $O(n^{2-\frac{1}{t+3}})$ tangencies. Next we consider the case when we are given two families of pairwise disjoint curves and count the number of tangencies among these two families. Theorem 2.43 shows that there may be as much as $\Omega(n^{4/3})$ tangencies in this setting. We prove an optimal upper bound in two special cases. On one hand, if every pair of curves intersect at most once, then there are only $O(n)$ tangencies by Theorem 2.42. On the other hand, if every curve is x-monotone, then there are $O(n \log n)$ tangencies by Theorem 2.41.

Results in this area will be discussed in Section 2.4.

2 Survey of results

In this chapter we give an overview of the most important geometric hypergraphs and the known results about them, concentrating on their colorings. We will also give results about the sizes (maximal number of hyperedges) in most cases. There are other parameters that are usually investigated for hypergraphs, like the size of containment-free hypergraphs, size of ε-nets, VC-dimension, existence of shallow hitting sets, discrepancy results, enumeration results. In this dissertation we concentrate on colorings and the size of (uniform) hypergraphs. The webpage [77] maintained by the author and Pálvölgyi collects these results in an interactive database, where the other above mentioned parameters are also included.

Concentrating on colorings and sizes, there is a nice general connection saying that if the Delaunay graph of every induced subhypergraph has linear many edges then the hypergraph's chromatic number, size and other parameters can be bounded. While this was known for specific hypergraph families, the author and his coauthors has stated this explicitly first, we state it later as Theorem 3.1 (Ackerman-Keszegh-Pálvölgyi [7]).

Still talking about results that apply to a broad set of hypergraphs, there is an intriguing general conjecture, claiming that if m_2 is bounded for a hypergraph and all its induced subhypergraphs, then m_k is bounded as well, in fact it is possible that it is always linear in k (for any fix bound on m_2) [118]. The boundedness of m_k was separately proved in all cases of geometric hypergraphs, as we will see. Yet in general not much is known about it. First, if $m_2 = 2$ for some hypergraph family, then a nice result of Berge [18] shows that $m_k = k$ as well. However, already for $m_2 = 3$ we do not know if an upper bound on m_3 exists. The best general lower bound is given by a construction where $m_k \geq (m_2-1)(k-1)+1$. Note that for $m_2 = 3$ this gives $m_3 \geq 5$. Improving on this slightly, recently Pálvölgyi [118] gave a construction such that $m_2 = 3$ and $m_3 = 6$. It is interesting to compare this problem with the question of Esperet mentioned in the Introduction whether χ-boundedness implies polynomial χ-boundedness, which turned out to have a 'no' answer.

The result mentioned earlier about the connection between the edges of the Delaunay-graph of a hypergraph and its other properties is detailed and proved in Chapter 3, based on [7].

In the remainder of this chapter we first list the geometric hypergraph results. Several of these results are by the author (with different coauthors). In this dissertation only the results about pseudo-disks and about ABA-free hypergraphs and their generalizations will be proved, in Chapters 4-9, based on [74, 69, 81, 4, 75].

After that we turn our attention to results about tangencies of curves. Here some results will be also phrased as a geometric hypergraph result, illuminating the connection between the two areas that was mentioned already in the Introduction. The results of the author from this area will be proved in Chapters 10-12, based on [7, 83, 6].

The figures showing the inclusion hierarchy of geometric hypergraph families are created using the interactive database of the author and Pálvölgyi [77].

The dissertation is based on the following papers of the author, in all of which every coauthor (in case of multiple authors) contributed an equal share: [74, 69, 81, 4, 75, 7, 83, 6].

2.1 A short history: disks, translates, homothets, axis-parallel regions

Usually in a family either the objects are meant to be all open or all closed, and given that we consider finite problems and assuming certain general position assumptions it does not matter which. Thus, in the survey part of the dissertation usually we won't specify this detail.

The area of coloring geometric hypergraphs was first motivated by cover-decomposability problems of translates of a given convex shape in the plane. We first give an overview of the results about translates, for more on these see the survey paper about cover-decomposability problems [99]. Then we will move on to overview the results about homothets, axis-parallel regions and disks. We note that in the context of cover-decomposability, infinite coverings are

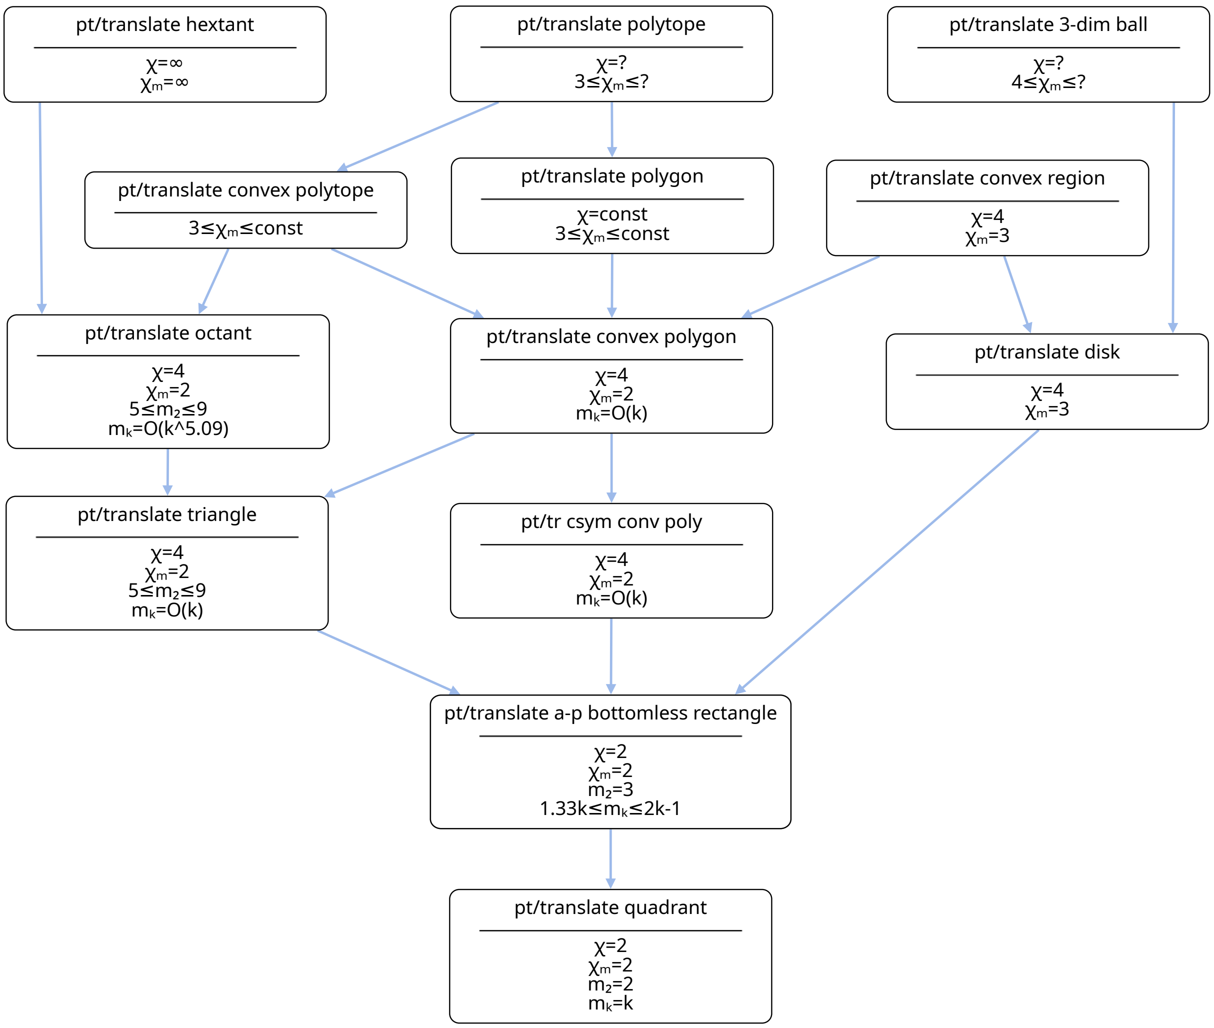

Figure 3: The inclusion hierarchy of hypergraph families defined by translates of geometric objects and the best known results about their respective coloring parameters. 'Tr csym conv poly' refers to the family 'translates of a fixed centrally symmetric convex polygon'. Note that these families defined by points wrt. translates of a given region are self-dual, e.g., the family of hypergraphs defined by points wrt. translates of a disk are the same as the hypergraphs defined by translates of a disk wrt. points, etc.

usually considered too, where compactness arguments and other technicalities arise. Here we only concentrate on the finite versions of the problems.

Pach initiated the area by conjecturing the following (here we use the hypergraph notation, the conjecture and in fact most papers about translates use the cover-decomposability terminology):

Conjecture 2.1 (Pach [96])**.** *Given a convex shape C, points are 2-colorable wrt. the translates of C. In other words, there exists a constant $m(C)$ such that for any finite point set $\mathcal{P}$, we can 2-color the points in $\mathcal{P}$ such that any translate of C that contains at least m points contains points of both colors.*

Note that for a given finite point set a finite number of translates defines the same hypergraph as all translates (this holds in general for any geometric hypergraph).

Further, denoting a fixed point in C as its centerpoint and replacing each translate with its centerpoint, while replacing each point of $\mathcal{P}$ with a reflected copy of C whose centerpoint is this point, the containments get exactly reversed. Thus we get an equivalent dual problem of coloring geometric objects with respect to points. Therefore, for translates, the primal and dual problems are equivalent and so it is enough to concentrate on the primal one, even though formally the dual problem is the one which is equivalent to the cover-decomposability problem that Pach's conjecture was originally about.

Motivated by this conjecture, there came multiple positive results about translates of convex polygons. Pach [97] proved that the conjecture holds for centrally symmetric convex polygons. Tardos and Tóth [137] proved for triangles and then Pálvölgyi and Tóth for every convex polygon [112]. Knowing that m_2 exists for points wrt. translates of a given convex polygon, the interest shifted to the existence and bounds on m_k. Initially the above 3 papers proved exponential bounds depending on k. Then Pach and Tóth [110] and then Aloupis et al. [12] improved this for quadratic and then linear for centrally symmetric convex polygons, respectively. Finally, Gibson and Varadarajan [56] showed that $m_k = O(k)$ for points wrt. translates of a given convex polygon.

On the negative side, for concave quadrilaterals we cannot 2-color points wrt. its translates [105, 111]. Such counterexamples also exist for balls in $\mathbb{R}^3$ [90] and convex polytopes in $\mathbb{R}^3$ [111].

Figure 3 summarizes the known results related to translates.

Points wrt. translates of a halfspace in $\mathbb{R}^3$ are trivial to 2-color. However, there exists a more interesting object in $\mathbb{R}^3$ such that points wrt. its translates are 2-colorable. Namely, the octant at the origin defined by the set of points $\{(x, y, z) : x, y, z > 0\}$ (we call its translates octants). For these the author and Pálvölgyi proved [78] that one can 2-color any point set in $\mathbb{R}^3$ such that any octant that contains at least 12 points contains points of both colors. This has important consequences in the plane. Namely, the same holds for coloring points wrt. homothets of any triangle (a homothet of a shape is a copy that we get by a translation and scaling by a positive factor), generalizing the result of Tardos and Tóth about translates of a triangle. It also implies the same for the dual problem, coloring homothets of a triangle wrt. points. Later the author and Pálvölgyi showed that polychromatic k-coloring is also possible for points wrt. octants [85], that is, m_k exists, then showed that $m_k = O(k^{4.53})$ for points wrt. homothets of a triangle [79]. This was proved by showing a more general result saying that the existence of m_2 implies a polynomial bound on m_k for points wrt. the homothets of any given convex polygon. Then Cardinal et al. showed $m_k = O(k^{5.53})$ for octants (and therefore also for the dual problem about coloring the homothets of a triangle wrt. points). Then the author and Pálvölgyi, by improving

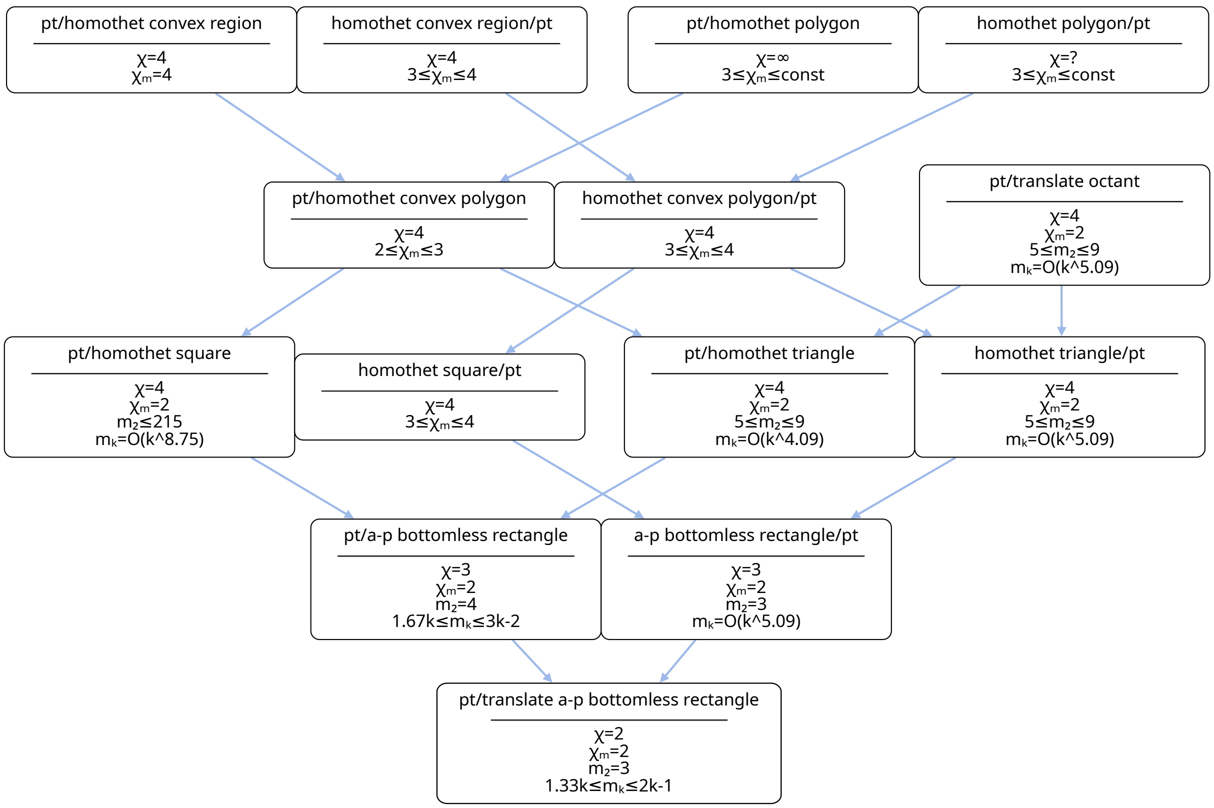

Figure 4: The inclusion hierarchy of hypergraph families defined by homothets of geometric objects and the best known results about their respective coloring parameters.

for octants the bound on m_2 from 12 to 9 [80], improved both of the above bounds to $O(k^{4.09})$ and $O(k^{5.09})$, respectively. It is still not known if an $O(k)$ bound is true or not.

These implications about homothets of a triangle opened the way of regarding homothets of a given convex shape. Note that for these the primal and dual problems are not equivalent anymore. In this area our understanding is still more limited. The author with Ackerman and Vizer proved [5] that m_2 exists for points wrt. the homothets of any fixed parallelogram (in particular, a square). Thus, by our general result also $m_k = O(k^{8.75})$ for these. In general, 2-colorability of points wrt. homothets of a convex polygon may always be true but is not known, already for pentagons. On the positive side, recall that whenever m_2 exists for points wrt. the homothets of some convex polygon, then by our result [79] m_k is bounded by some polynomial of k as well. Further, with Pálvölgyi we showed that proper 3-colorability does always hold in this case [82]. We have noted that the octant result implies the existence of m_2 for the dual problem of coloring homothets of a triangle with respect to points. However, somewhat surprisingly, Kovács proved that this becomes false already for the homothets of any given convex polygon with at least four vertices [88].

Figure 4 summarizes the known results related to homothets.

Further consequence of the results about octants is that it implies the same for 2-coloring points wrt. bottomless rectangles (these are defined as $\{(x, y) : a < x < b, y < c\}$ for any a, b, c) and for the dual problem of coloring bottomless rectangles wrt. points. These were first studied systematically by the author in [73]. In the primal case Asinowski et al. showed directly that $m_k \leq 3k-2$ [15] (the best lower bound for the optimal value is approx. $1.67k$). However, in the dual case the best known result about bottomless rectangles, $m_k = O(k^{5.09})$, is indeed implied by the results about octants. There are some special cases of bottomless rectangle families for which $m_k = O(k)$ is known [25], yet in general this is the best known bound.

The motivation to study coloring points wrt. bottomless rectangles was motivated by its equivalence to a certain dynamic coloring problem of points on the line. Moreover, the dual problem can be used in divide and conquer algorithms when dealing with axis-parallel rectangles in general. Hypergraphs defined by axis-parallel rectangles are the basis of some of the central problems of the area. The flavor of these problems is slightly different from the above mentioned ones, as here the number of colors does depend on the number of points. About the dual problem, $\Theta(\log n)$ colors are needed and enough for proper coloring axis-parallel rectangles wrt. points (for any fixed m) [62, 109]. For the primal problem, if $m \geq 3$ then again $O(\log n)$ colors are enough [8], yet when we color the points such that axis-parallel rectangles containing two points are already properly colored (in other words, we color the respective Delaunay graph), then it is only known that $\Omega(\log n / \log\log n)$ colors are needed [33] and $O(n^{0.368})$ are sufficient [10, 29]. Closing the gap between these bounds is an important open problem of the area.

We have seen that Conjecture 2.1 is verified for convex polygons. What do we know about other shapes? The most natural shape which is not a polygon is certainly the (unit) disk. In an unpublished manuscript Mani-Levitska and Pach [90] already presented a proof of the conjecture for translates of a unit disk. However, this was never verified and later, surprisingly, Pálvölgyi found a counterexample (published together with some of the results from [90] in [106]), thus showing that Pach's conjecture fails for a disk. In the same paper they show that it fails also for convex sets with a smooth boundary with everywhere positive curvature but it holds for unbounded convex sets. Later Pálvölgyi and Tóth [112] gave a complete characterization of the convex shapes for which the conjecture holds.

It was proved even earlier by Pach, Tardos and Tóth [105] that there is no 2-coloring for points wrt. homothets of a disk (that is, all disks).

These raised the question about the exact value of χ_{m} for the family of translates/homothets of a disk. First, the well-known Delaunay graph which coincides with the Delaunay graph of the hypergraph of points wrt. disks is known to be planar and thus 4-colorable, and 4 colors might be needed already for points wrt. unit disks hypergraphs (for example K_4 can be easily realized). Therefore χ is at most 4 even for the homothets of a disk, and so (using that $\chi_{\mathrm{m}} \leq \chi$) χ_{m} is also at most 4 for both families. Smorodinsky [128], using a dualization argument (by mapping points to halfspaces in $\mathbb{R}^3$ and disks to points in $\mathbb{R}^3$) showed that χ is also at most 4 for homothets of a disk wrt. points.

The problem of deciding if for these families χ_{m} is 3 or 4 was posed multiple times, e.g., by the author in [73], by now except for one case the problem is solved. Damásdi and Pálvölgyi [37] showed that χ_{m} is 3 for points wrt. unit disks (and therefore also for the dual setting). On the other hand, the same authors showed [36] that χ_{m} is 4 for points wrt. disks. As here the dual problem is not equivalent anymore, the case of coloring disks wrt. points is left open.

Figure 5 summarizes the known results about hypergraph families related to disks and pseudo-disks (see the forthcoming sections about their definitions).

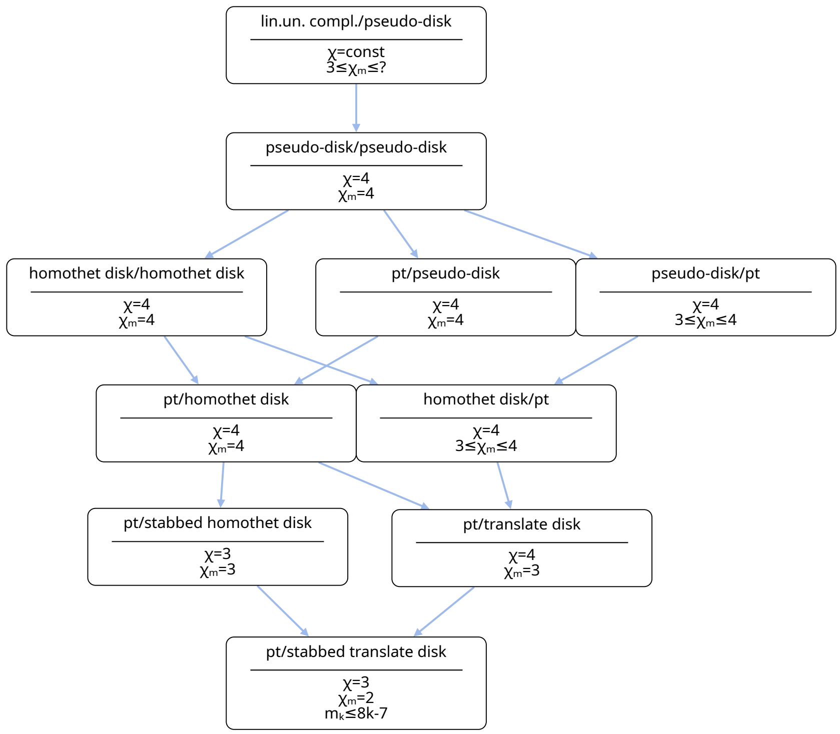

Figure 5: The inclusion hierarchy of hypergraph families defined by disks and related objects and the best known results about their respective coloring parameters.

2.2 Pseudo-disks

The results of the author presented in this section are based on [74, 69] and are proved in Chapters 4,5 and 6.

2.2.1 Pseudo-disks wrt. pseudo-disks

The problem about disks can be extended in two natural ways. First, we can regard families of pseudo-disks. Second, we can consider intersection hypergraphs of some objects wrt. disks or pseudo-disks where the objects are not points anymore. We will show that we can prove coloring results even about the common generalization of these two possible extensions. As we have seen that $\chi_m = \chi = 4$ already for points wrt. disks, in these generalizations it is more interesting to concentrate on χ instead of χ_m, therefore in this section we always talk about χ, that is, we color all hyperedges, not just the large enough hyperedges.

We first define families of pseudo-disks, which extends families of disks.

Definition 2.2. *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.*

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, just like pseudoline arrangements generalize line arrangements (see Section 2.3 for more about them). Problems regarded range from classic algorithmic questions like finding maximum size independent (disjoint) subfamilies [31] to classical combinatorial geometric questions like the Erdős-Szekeres problem [64]. Probably the most important pseudo-disk family is the family of homothets of a convex region. Because of this, results about pseudo-disks generalize to results about homothets of any fixed convex region.

There are only a few earlier results regarding this general setting. In [80] intersection (and also inclusion and reverse-inclusion) hypergraphs of intervals of the line were considered by the author and Pálvölgyi. In [71] and [48] intersection hypergraphs (and graphs) of (unit) disks, pseudo-disks, squares and axis-parallel rectangles were considered.

It is well known that the Delaunay graph of points wrt. disks is planar and thus proper 4-colorable. From the following trivial yet useful observation it then follows that the respective hypergraph is also proper 4-colorable.

Observation 2.3 (Smorodinsky [128])**.** *If $\mathcal{A}$ and $\mathcal{B}$ are families for which every hyperedge of $\mathcal{I}(\mathcal{A},\mathcal{B})$ contains a hyperedge of size 2, then a proper coloring of the Delaunay graph is also a proper coloring of $\mathcal{I}(\mathcal{A},\mathcal{B})$.*

In the literature hypergraphs that have the property assumed in Observation 2.3 are sometimes called *rank two* hypergraphs (e.g., in [128]).

Definition 2.4. *Given a hypergraph $\mathcal{H}$, a* conflict-free k-coloring *of its vertices is a coloring with k colors such that every hyperedge H of $\mathcal{H}$ contains a vertex whose color differs from the color of all the other vertices of H.*

In [128] Smorodinsky developed a general framework (based on the framework presented in [46]): 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. [46] that points wrt. disks admit a conflict-free coloring with $O(\log n)$ colors:

Theorem 2.5 (Even et al.[46])**.** *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 $n = |P|$.*

It is well known that if instead of a family of disks we take a family of pseudo-disks, the bound of Theorem 2.5 still holds:

Theorem 2.6 (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 4 colors and a conflict-free coloring with $O(\log n)$ colors, where $n = |P|$.*

The proof is that the Delaunay graph of points with respect to pseudo-disks is also a planar graph (implied by, e.g., Lemma 4.9, see later) and thus proper 4-colorable and then we can apply Observation 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 2.5 was also proved by Even et al. and was generalized to pseudo-disks by Smorodinsky:

Theorem 2.7 (Even et al. [46]). *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 4 colors and a conflict-free coloring with $O(\log n)$ colors, where $n = |\mathcal{B}|$.*

Theorem 2.8 (Smorodinsky [128]). *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 region, Cardinal and Korman showed that 4 colors are enough, just like for disks:

Theorem 2.9 (Cardinal-Korman [26]). *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 4 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 2.5 and 2.7:

Theorem 2.10 (Keller-Smorodinsky [71]). *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 6 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 2.11 (Keller-Smorodinsky [71]). *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 $n = |\mathcal{B}|$.*

We generalize Theorem 2.10 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 2.5,2.7,2.8,2.9,2.10 and Claim 2.11). Moreover we prove the optimal upper bound of 4 colors, which improves the bound of Theorem 2.8 for coloring pseudo-disks wrt. points from some constant number of colors to 4 colors and improves the bound of Theorem 2.10 for coloring disks wrt. disks from 6 colors to 4 colors. Along the way we prove that the Delaunay graph of the hypergraph defined by one family of pseudo-disks wrt. another family of pseudo-disks is planar. Furthermore, it provides an alternative proof for Theorems 2.7 and 2.9 (both of these were originally proved using dualization and solving equivalent problems about coloring points in the 3 dimensional space):

Theorem 2.12 (Keszegh [74]). *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 4 colors.*

Using the usual framework it easily follows from Theorem 2.12 that:

Corollary 2.13 (Keszegh [74]). *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 2.11 implies in a straightforward way the main result of [71] about conflict-free coloring the (open/closed) neighborhood hypergraphs of intersection graphs of pseudo-disks (for definitions and details see [71]). Thus, Theorem 2.12 also implies their main result.

Note that in Theorem 2.12 $\mathcal{B}$ 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 2.14 (Keszegh [74])**.** *We can proper color with* 4 *colors the family of homothets of a convex region A wrt. the family of homothets of another convex region B.*

Buzaglo et al. [24] 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 t in such a hypergraph is $O(t^2n)$ (and this can be attained already by $\mathcal{F}$ being a specific family of disks). More recently, Aronov et al. [13] proved independently from us with an almost identical proof the case of Lemma 4.10 when $\mathcal{B} \subset \mathcal{F}$ (for details see later, it is a very special case of Theorem 4.3) 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 [24] they show that this implies that the number of hyperedges of size at most t in $\mathcal{I}(\mathcal{B},\mathcal{F})$ in this case is $O(t^3n)$ (they do not give a construction with a matching lower bound, however). Alternatively, it follows from Theorem 3.1 as well. We show that using Lemma 4.10 instead of their weaker variant we get the following more general statements (where $\mathcal{B}$ is not necessarily a subfamily of $\mathcal{F}$):

Theorem 2.15 (Keszegh [74])**.** *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 2.16 (Keszegh [74])**.** *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 t in the intersection hypergraph $\mathcal{I}(\mathcal{B},\mathcal{F})$ is $O(t^3n)$.*

We note that in case of Theorem 2.16 we are again not aware of a matching lower bound (and in fact this upper bound might easily not be tight).

A family $\mathcal{B}$ of Jordan regions is called *non-piercing* if $A \setminus B$ is connected for every pair of sets $A, B \in \mathcal{B}$. Notice that a family of pseudo-disks is always non-piercing. Raman and Ray [120] investigated respective problems about non-piercing families regions. Instead of the Delaunay graph, they were interested in finding planar support graphs, due to their applications for finding PTAS's for packing and covering problems. A *support graph* of a hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ is a graph $\mathcal{G}$ with vertex set $\mathcal{V}$ such that for every $E \in \mathcal{E}$ the subgraph of $\mathcal{G}$ induced by E is connected. Notice that a support graph necessarily contains the Delaunay graph of $\mathcal{H}$ as a subgraph. It is further called a *planar support graph* if G is planar. Their main result is the following:

Theorem 2.17 (Raman-Ray [120])**.** *Given two families $\mathcal{F}$ and $\mathcal{B}$ of non-piercing regions, $\mathcal{I}(\mathcal{B},\mathcal{F})$ admits a planar support graph.*

Notice that while a proper coloring of the Delaunay graph of a hypergraph is not always a proper coloring of the hypergraph itself, proper coloring a support graph is always a proper coloring of the hypergraph as well. Recalling that a pseudo-disk family is also non-piercing, 4-coloring the planar support graph guaranteed by Theorem 2.17 implies Theorem 2.12.

2.2.2 Family with linear union complexity wrt. pseudo-disks

In [128] Theorem 2.8 was shown by proving a more general statement about coloring a family of regions that has linear union complexity wrt. points.

Definition 2.18. *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 2.19. *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 $\cup_{B\in\mathcal{B}}B$.*

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 $c|\mathcal{B}'|$.*[3]

Theorem 2.20 (Kedem et al. [68])**.** *Any finite family of pseudo-disks in the plane has a linear union complexity.*

Theorem 2.20 shows that the following result of Smorodinsky is indeed more general than Theorem 2.8.

Theorem 2.21 (Smorodinsky [128])**.** *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 $n = |\mathcal{B}|$.*[4]

We generalize this in the following way:

Theorem 2.22 (Keszegh [74])**.** *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 2.20 we get that Theorem 2.22 implies Theorem 2.12 with a (non-explicit and worse) upper bound.

Corollary 2.23 (Keszegh [74])**.** *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 $n = |\mathcal{B}|$.*

Corollary 2.24 (Keszegh)**.** *Given a family $\mathcal{F}$ of pseudo-disks and a finite family $\mathcal{B}$ of Jordan regions with linear union complexity, the VC-dimension d of $\mathcal{I}(\mathcal{B}, \mathcal{F})$ is bounded and the number of hyperedges of $\mathcal{I}(\mathcal{B}, \mathcal{F})$ of size at most t is $O(t^{d-1}n)$.*

So far we have considered the hypergraphs defined by pseudo-disks wrt. pseudo-disks and the hypergraphs defined by families with linear union complexity wrt. pseudo-disks. Next we consider hypergraphs defined by lines wrt. pseudo-disks. Note that a family of lines does not have linear union complexity.

---

[3]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.

[4]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.

2.2.3 Lines wrt. pseudo-disks

In this section we consider hypergraphs defined by lines wrt. pseudo-disks, that is, intersection hypergraphs $\mathcal{I}(\mathcal{L},\mathcal{F})$ where $\mathcal{L}$ is a family of lines in the plane, and $\mathcal{F}$ a family of pseudo-disks. We assume that the geometric objects are in general position, in the sense that no 3 lines pass through a common point, no line passes through an intersection point of two boundaries of pseudo-disks.

For geometric hypergraphs defined by lines wrt. pseudo-disks we concentrate only on their size. The reason we don't consider their colorings is that even χ_{m} is infinite for this family [105]. Indeed, it is a more general family then the family of hypergraphs defined by lines wrt. points which by a standard dualization argument is the same as its dual, the family of hypergraphs defined by points wrt. lines. For this latter family for arbitrary c, m we can take the appropriate construction from the Hales–Jewett theorem and project it onto the plane in a general direction to get a construction in which in any c-coloring of the points there will be a monochromatic line with m points.

Unlike for the hypergraphs of points wrt. pseudo-disks, the number of hyperedges in a hypergraph $\mathcal{I}(\mathcal{L},\mathcal{F})$, of lines wrt. pseudo-disks, of any fixed size, may be quadratic in the number of vertices. This holds already when the lines form an $n/2 \times n/2$ grid and we take the lines wrt. disks hypergraph. A better example was shown in a beautiful paper of Aronov et al. [14]. They showed that for any family $\mathcal{L}$ of lines, if $\mathcal{F}$ consists of the inscribed circles of the triangles formed by any triple of lines, then for any $t \geq 3$, the number of t-hyperedges (i.e., hyperedges of size t) in $\mathcal{I}(\mathcal{L},\mathcal{F})$ is exactly $\binom{n-t+2}{2}$.

For any fixed t, there exist hypergraphs $\mathcal{I}(\mathcal{L},\mathcal{F})$ in which the number of t-hyperedges is larger than in the construction of Aronov et al. [14], even when $\mathcal{F}$ is allowed to contain only disks (as some of those disks might not be inscribed in a triangle formed by the lines). We prove that the number of t-hyperedges cannot be significantly larger for any hypergraph $\mathcal{I}(\mathcal{L},\mathcal{F})$ of lines with respect to pseudo-disks.[5] Specifically, we prove ($\mathcal{E}(\mathcal{H})$ denotes the set of hyperedges of $\mathcal{H}$):

Theorem 2.25 (Keller-Keszegh-Pálvölgyi [69])**.** *Let $\mathcal{L}$ be a family of n lines in the plane, let $\mathcal{F}$ be a family of pseudo-disks, and assume both families are in general position. Then*

$$|\{e \in \mathcal{E}(\mathcal{I}(\mathcal{L},\mathcal{F})) : |e| = t\}| = O_t(n^2).$$

Our techniques combine probabilistic and planarity arguments, together with exploiting properties of arrangements of lines, in particular the *zone theorem.*

We note that in Theorem 2.25 the dependence on t is not proved to be optimal. In particular $O(tn^2)$ or even $O(n^2)$ (i.e., a bound independent of t) might be true.

In addition, we show that for any choice of $\mathcal{L}$ and $\mathcal{F}$, the total number of hyperedges in $\mathcal{I}(\mathcal{L},\mathcal{F})$ does not exceed $O(n^3)$. This upper bound is tight, since the total number of hyperedges in the hypergraph presented by Aronov et al. [14] is $\binom{n}{3}$.

Proposition 2.26 (Keller-Keszegh-Pálvölgyi [69])**.** *Let $\mathcal{L}$ be a family of n lines in the plane, let $\mathcal{F}$ be a family of pseudo-disks, and assume both families are in general position. Then $|\mathcal{E}(\mathcal{I}(\mathcal{L},\mathcal{F}))| = O(n^3)$.*

[5] For the difference between hypergraphs induced by pseudo-disks and hypergraphs induced by disks, see [51] and the references therein.

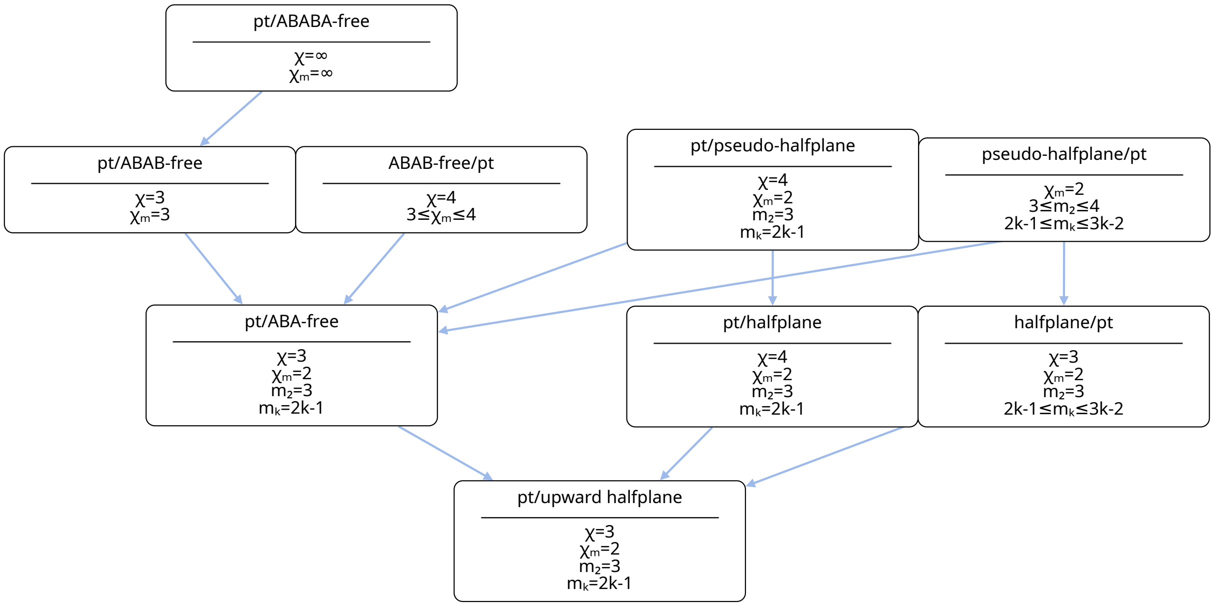

Figure 6: The inclusion hierarchy of hypergraph families defined by halfplanes and related objects and the best known results about their respective coloring parameters.

2.3 ABA-free hypergraphs and their generalizations

The results of the author presented in this section are based on [81, 4, 75] and are proved in Chapters 7 and 8.

2.3.1 ABA-free hypergraphs and pseudo-halfplane hypergraphs

After investigating translates and homothets of convex polygons, disks and then pseudo-disks, the perhaps most natural remaining planar family of shapes is the family of halfplanes. Coloring points wrt. halfplanes and its dual were first investigated by the author [72, 73] followed by Fulek [53], and then polychromatic k-colorings were studied by Smorodinsky and Yuditsky [130]. They proved that for points wrt. halfplanes $m_k = 2k - 1$ while for halfplanes wrt. points $m_k \leq 3k - 2$, this latter bound is not known to be sharp.

Figure 6 summarizes the known results about hypergraph families related to halfplanes (see the forthcoming sections about their definitions).

Similar to how pseudo-disks generalize disks, one can generalize halfplanes to pseudo-half-planes. Informally, a pseudo-line family is a family of bi-infinite curves that pairwise intersect at most once and then a pseudo-halfplane family is a family of regions in the plane whose boundaries form a family of pseudo-lines. For the precise definition see Section 7.6.

We have mentioned earlier that it was proved that points wrt. the translates of any unbounded convex set are 2-colorable. In fact $m_2 = 3$ is sufficient [106]. As translates of an unbounded convex set form a set of pseudo-halfplanes (as their boundaries intersect at most once, and thus form a family of pseudo-lines), our result about pseudo-halfplanes will imply the generalization $m_k = 2k - 1$, which is an optimal function for unbounded convex sets.

We generalize all the (primal and dual) results about halfplanes to pseudo-halfplanes, answering a question left open by the authors of [130].[6] To do this we define families of hypergraphs in a combinatorial way, with no geometry used. We show in Section 7.6 that these are equivalent to points wrt. pseudo-halfplane hypergraphs. Then we are left to deal with these combinatorially defined hypergraphs.

Definition 2.27. *A hypergraph $\mathcal{H}$ with an ordered vertex set is called* ABA-free *if H does not contain two hyperedges A and B for which there are three vertices $x < y < z$ such that $x, z \in A \setminus B$ and $y \in B \setminus A$.*

A hypergraph with an unordered vertex set is ABA-free if its vertices have an ordering with which the hypergraph is ABA-free.[7]

A hypergraph $\mathcal{H}$ on an ordered set of points S is called a pseudo-halfplane hypergraph *if there exists an ABA-free hypergraph $\mathcal{F}$ on S such that $\mathcal{H} \subset \mathcal{F} \cup \bar{\mathcal{F}}$, where $\bar{\mathcal{F}}$ is the family of the complements of the hyperedges of $\mathcal{F}$.*

The following are examples of ABA-free hypergraphs: points wrt. intervals in $\mathbb{R}$, points wrt. translates of an upwards unbounded convex set, points wrt. upwards halfplanes. Additionally, points wrt. all halfplanes is an example of pseudo-halfplane hypergraphs. These examples show how geometric problems follow from abstract problems about ABA-free hypergraphs. Moreover, ABA-free hypergraphs have an equivalent geometric representation with upwards pseudo-halfplanes, which naturally generalize these geometric examples. For more on these geometric connections see Section 7.6. We note that colorings of ABA-free hypergraphs were first defined in [106] under the name *special shift-chains* as a subfamily of so called *shift-chains*.

Our main result about ABA-free hypergraphs is the following.

Theorem 2.28 (Keszegh-Pálvölgyi [81])**.** *Given an ABA-free $\mathcal{H}$ we can color its vertices with k colors such that every $A \in \mathcal{H}$ whose size is at least $2k-1$ contains all k colors.*

This is proved in Section 7.1 (following closely the ideas of Smorodinsky and Yuditsky [130]). We then observe that the dual of this problem is equivalent to the primal, which implies that the hyperedges of every $(2k-1)$-uniform ABA-free hypergraph can be colored with k colors, such that if a vertex v is in a subfamily $\mathcal{H}_v$ of at least $m_k = 2k-1$ of the hyperedges of $\mathcal{H}$, then $\mathcal{H}_v$ contains a hyperedge from each of the k color classes.

About the more general family of pseudo-halfplane hypergraphs we are able to show the same bound, proved in Section 7.2:

Theorem 2.29 (Keszegh-Pálvölgyi [81])**.** *Given a pseudo-halfplane hypergraph $\mathcal{H}$ we can color its vertices with k colors such that every $A \in \mathcal{H}$ whose size is at least $2k-1$ contains all k colors.*

Both results are sharp. Note that, as mentioned before, these results imply the same for hypergraphs defined by points wrt. translates of unbounded convex sets.

About the dual hypergraph we can show the following.

Theorem 2.30 (Keszegh-Pálvölgyi [81])**.** *Given a pseudo-halfplane hypergraph $\mathcal{H}$ we can color its hyperedges with k colors such that every vertex which is contained by at least $3k-2$ hyperedges is contained by hyperedges of all k colors.*

[6]Personal communication, Shakhar Smorodinsky.

[7]While it might seem that using the same notion for ordered and unordered hypergraphs leads to confusion as by forgetting the ordering of an ordered hypergraph it might become ABA-free, from the context it will always be perfectly clear what we mean.

This is proved in Section 7.3. This result might not be sharp, the best known lower bound for $m(k)$ is $2k-1$ [130]. We additionally consider so called pseudo-hemisphere hypergraphs, which are common generalizations of primal and dual pseudo-halfplane hypergraphs.

2.3.2 ABAB-free hypergraphs

Definition 2.27 can be generalized in a straightforward way, similarly to Davenport-Schinzel sequences [38], to more alternations. We show that already one more alternation gives non-two-colorable hypergraphs.

Let $\ell \geq 1$ be a number such that 2ℓ is an integer. We denote by $(AB)^\ell$ the alternating sequence of letters A and B of length 2ℓ. For example, $(AB)^{1.5} = ABA$ and $(AB)^2 = ABAB$.

Definition 2.31 ($(AB)^\ell$-free hypergraphs)**.** *Two subsets A, B of an ordered set of elements form an* $(AB)^\ell$-sequence *if there are 2ℓ elements $a_1 < b_1 < a_2 < b_2 < \ldots < a_\ell < b_\ell$ such that $\{a_1, a_2, \ldots, a_\ell\} \subset A \setminus B$ and $\{b_1, b_2, \ldots, b_\ell\} \subset B \setminus A$.*

A hypergraph with an ordered vertex set is $(AB)^\ell$-free *if it does not contain two hyperedges A and B that form an $(AB)^\ell$-sequence.*

A hypergraph with an unordered vertex set is $(AB)^\ell$-free *if there is an order of its vertices such that the hypergraph with this ordered vertex set is $(AB)^\ell$-free.*

Our results about ABAB-free hypergraphs are the following:

Theorem 2.32 (Keszegh-Pálvölgyi [81])**.** *For every $m \geq 2$ there exists an ABAB-free m-uniform hypergraph which is not 2-colorable.*

Theorem 2.33 (Ackerman-Keszegh-Pálvölgyi [4])**.** *Every ABAB-free hypergraph is proper 3-colorable.*

Similar to how ABA-free hypergraphs can be realized by points wrt. upwards pseudo-halfplanes, ABAB-free hypergraphs can be realized by points wrt. upwards pseudo-parabolas and more generally $(AB)^\ell$-free hypergraphs by points wrt. the upper sides of curves in an arrangement of bi-infinite x-monotone curves that intersect at most $2\ell - 2$ times. While this is easy to prove, we also show a non-trivial realization of ABAB-free hypergraphs, by points wrt. a family of stabbed pseudo-disks (stabbed means that their intersection is non-empty):

Theorem 2.34 (Ackerman-Keszegh-Pálvölgyi [4])**.** *A hypergraph is ABAB-free if and only if it can be realized by points wrt. a family of stabbed pseudo-disks.*

This theorem implies the following:

Theorem 2.35 (Ackerman-Keszegh-Pálvölgyi [4])**.** *Let $\mathcal{F}$ be a family of pseudo-disks whose intersection is non-empty and let S be a finite set of points. Then it is possible to color the points in S with 3 colors such that any pseudo-disk in $\mathcal{F}$ that contains at least two points from S contains two points of different colors. Moreover, for every integer m there is a set of points S and a family of pseudo-disks $\mathcal{F}$ with a non-empty intersection, such that for every 2-coloring of the points there is a pseudo-disk containing at least m points, all of the same color.*

From Theorem 2.35 it is easy to conclude the following.

Corollary 2.36 (Ackerman-Keszegh-Pálvölgyi [4])**.** *Given a finite set of points S it is possible to color the points of S with three colors such that any disk that contains the origin and at least two points from S contains two points with different colors.*

Recently Damásdi and Pálvölgyi proved that already for points wrt. stabbed disks 2 colors are not enough [36] (for any fixed m) and so this bound of 3 is best possible. They also show that for points wrt. stabbed unit disks 2 colors are enough. Same goes for points wrt. stabbed translates and homothets of a convex region (2 and 3 colors, respectively, are enough and needed). However, for points wrt. stabbed homothets of a convex polygon they prove that 2 colors are enough. Note that this is yet another special case solved of the open problem of 2-coloring points wrt. homothets of a convex polygon.

While we omit the proof of Theorem 2.34 from this dissertation (the details can be found in [4]), we provide a direct proof of Theorem 2.35 (in Section 8.5) which is shorter, yet less self-contained, as it uses some previous results about the so-called "shrinkability" of a family of pseudo-disks [24, 115] that rely on a highly nontrivial sweeping machinery from [131].

2.3.3 Further results

We have dealt with ABA-free hypergraphs, their duals and their other generalizations. We also dealt with ABAB-free hypergraphs. We also presented several consequences about hypergraphs realizable by geometric objects.

Similar questions can be studied about dual-ABAB-free hypergraphs as well, which is equivalent to the so-called cover-decomposition problem for stabbed pseudo-disks. In this respect with Pálvölgyi we have proved (work in progress) that they are proper 4-colorable, while for every $m \geq 2$ there exists a dual-ABAB-free m-uniform hypergraph which is not 2-colorable. This leaves open the problem to decide if 3 or 4 colors are needed for proper coloring m-uniform dual-ABAB-free hypergraphs.

Another version is to forbid $ABABA$-sequences cyclically (instead of linearly); such 3-uniform hypergraphs have a nice geometric representation, as *convex geometric 3-hypergraphs* without *strongly crossing* edges, see Suk [134]. It is also a natural question to ask whether strongly crossing convex geometric (non-uniform) hypergraphs can be always 3-colored.

Going further, we consider $(AB)^\ell$-free hypergraphs for $\ell > 2$ and show that $\chi_{\mathrm{m}} = \infty$ already for ABABA-free hypergraphs, using a construction similar to the one in the proof of Theorem 2.32, which in turn uses the hypergraph from [105].

Theorem 2.37 (Ackerman-Keszegh-Pálvölgyi [4])**.** *For every $c \geq 2$ and $m \geq 2$ there exists an $ABABA$-free m-uniform hypergraph which is not c-colorable.*

This is proved in Section 8.4.

What can we say about the size of $(AB)^\ell$-free hypergraphs on n vertices? First, trivially an $(AB)^\ell$-free hypergraph has VC-dimension at most $2\ell - 1$, and therefore by the Sauer-Shelah lemma has at most $O(n^{2\ell-1})$ hyperedges. On the other hand take an ordered set of n vertices and split it into $2\ell-1$ intervals of almost the same size. Take all hyperedges that are the unions of prefixes of these intervals. We get a hypergraph with $\Theta(n^{2\ell-1})$ hyperedges and it is easy to see that it is $(AB)^\ell$-free. Thus this is the maximum size of an $(AB)^\ell$-free hypergraph for fixed ℓ.

We remark that having VC-dimension at most $2\ell - 1$ is a weaker assumption than being $(AB)^\ell$-free. For any c and m there are m-uniform hypergraphs of VC-dimension 2 that are not c-colorable; the main construction from both [105] and [111] can be generalized from 2-colors to c-colors as m-uniform hypergraphs of VC-dimension 2.

In [23], working in the geometric setting, they consider the problem of bounding the size-k hyperedges. They determine that the size of k-uniform ($k \leq n/2$) $(AB)^\ell$ hypergraphs on n vertices is $O(k^{\ell-1}n^{\ell-1})$ and show that this bound is best possible.

Finally, let us mention another interesting direction. In our proofs of pseudo-halfplane hypergraph results, we will need the following Helly-type result (Lemma 7.36): If any three hyperedges of a pseudo-halfplane hypergraph intersect, then we can extend the hypergraph with a vertex contained in every hyperedge so that it is still a pseudo-halfplane hypergraph. Recently a discrete Helly-type result was shown for halfplanes by Jensen et al. [66]. It turns out that this also generalizes to pseudo-halfplanes: among others, in [75] we have proved that given a pseudo-halfplane hypergraph, if every triple of hyperedges has a common vertex then there exists a set of at most two vertices that hits every hyperedge. Summarizing, we can hit every hyperedge either with one vertex 'outside' the original vertex set, or with two vertices from the original vertex set. Similar results about dual pseudo-halfplane hypergraphs are shown as well in [75]. This thread of research is continued in [84] by the author of the dissertation in the following way. One way to define convex sets in the plane is that they are intersections of halfplanes. Analogously, we can consider so called pseudoconvex sets defined by intersections of hyperedges in a pseudo-halfplane hypergraph. In [84] several classical results of convexity are generalized to this notion by the author, including Helly theorem, Carathéodory theorem, Kirchberger theorem, Separation theorem, Radon theorem and the Cup-Cap theorem. The details of these results are out of scope of this dissertation, for more on them and their connections to several other areas of convexity theory, see [84].

2.4 Tangencies

The results of the author presented in this section are based on [7, 83, 6] and are proved in Sections 11,10 and 9.

In this section, we study the maximum number of tangencies between curves, where two curves are called tangent if they intersect exactly once, and in that intersection point they do not cross, but touch each other.

A planar curve is a *Jordan arc*, that is, the image of an injective continuous function from a closed interval into $\mathbb{R}^2$. If no two points on a curve have the same x-coordinate, then the curve is *x-monotone*. We consider families of curves such that every pair of curves intersect at a finite number of points. Such a family is called *t-intersecting* (sometimes we refer to these as *at most t-intersecting*) if every pair of curves intersects at at most t points. An intersection point p of two curves is a *crossing point* if there is a small disk D centered at p which contains no other intersection point of these curves, each curve intersects the boundary of D at exactly two points and in the cyclic order of these four points no two consecutive points belong to the same curve. If two curves intersect at exactly one point which is not a crossing point, then we say that they are *touching* or *tangent* at that point.

The number of *tangencies* is the number of tangent pairs of curves. If more than two curves are allowed to intersect at a common point, then every pair of curves might be tangent, e.g., for the graphs of the functions x^{2i}, $i = 1, 2, \ldots, n$, in the interval $[-1, 1]$. Therefore, we restrict our attention to families of curves in which no three curves intersect at a common point.

In many cases one can locally perturb curves such that tangencies become digons (faces of size two in the arrangement of the curves) and vice versa, although one has to be careful since this might lead to more intersection points or a different notion of tangency (e.g., two curves may form many digons but can they touch in more than one point?). There are quite a few problems in combinatorial geometry that can be phrased in terms of bounding the number of tangencies (or digons) among certain curves, see, e.g., [9]. The most famous of which is the unit distance problem of Erdős [42] mentioned already at the very beginning of the Introduction.

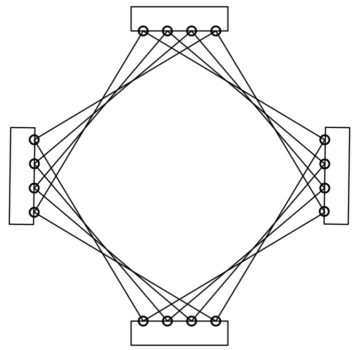

Figure 7: n convex polygons can have $\Omega(n^2)$ tangencies.

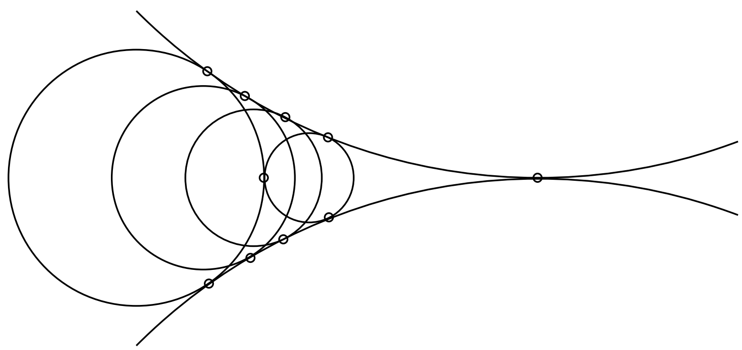

Figure 8: n circles can have $2n-2$ tangencies.

Given a family of shapes, their *tangency graph* has vertices corresponding to the shapes and edges corresponding to tangent pairs. With this notation, our primary aim is to bound the number of edges of the tangency graphs of certain families of shapes by a function of the number of shapes.

If the curves in the family are not allowed to cross, i.e., they are disjoint apart from tangencies, then their tangency graph is a planar graph, and it follows from Euler's formula that the number of tangencies is at most $3n-6$ among $n \geq 3$ curves [44].

On the other hand, if there is no restriction on the number of intersections, we can have $\lfloor n^2/4 \rfloor$ tangencies, as we can easily draw two families of $n/2$ curves such that every pair of curves, with one from each family, is tangent. We can even do this easily such that the curves are boundaries of convex shapes [47], see Figure 7 . However, in such a drawing we necessarily have a lot of intersections between pairs of curves from the same family.

Therefore, we are interested in bounds assuming various restrictions on the possible intersections.

First we concentrate on the case when we have one family of curves, either required to be pairwise intersecting or not, and then we consider the case when we have two families, each of which consists of pairwise disjoint curves.

2.4.1 Tangencies among a single family of pairwise intersecting curves

A natural condition which gives rise to many interesting problems is if we require that the curves are pairwise intersecting.

We first consider circles and pseudo-circles. Alon et al. [11] proved that every arrangement of n pairwise intersecting circles contains $O(n)$ digons (see also [9]). A lower bound for this case is a simple construction of circles with $2n - 2$ tangencies, see Figure 8. A famous conjecture of Grünbaum [59] states that any arrangement of n pairwise intersecting pseudo-circles in the plane can in fact have at most $2n-2$ digons (for $n \geq 4$). Agarwal et al. [9] proved this conjecture for arrangements in which there is a common point surrounded by all pseudo-circles. Recently, Felsner, Roch and Scheucher [50] showed that Grünbaum's conjecture is true also for arrangements of pseudo-circles in which there are three pseudo-circles, every pair of which creates a digon. Most recently the author with several coauthors showed this 50 year old conjecture first for circles [2] and then also the full conjecture for pseudo-circles [3]. Note that for pseudo-circles counting digons is equivalent to counting tangencies.

Pairwise intersecting pseudo-circles are closed curves that need to intersect pairwise at least once and at most twice. If instead we have not-necessarily closed curves, then with the same assumption of having to intersect pairwise at least once and at most twice, we can already have $\Omega(n^{4/3})$ tangencies,[8] while the best upper bound is $O(n^{9/5})$ that uses only the at most 2-intersecting property (follows from Theorem 2.40). However, Pach [95] conjectured that for not necessarily closed curves, requiring every pair of curves to intersect exactly once (either at a crossing or a tangency point) leads again to a linear upper bound on the number of tangencies:

Conjecture 2.38 (Pach [95]). *Let $\mathcal{C}$ be a set of n curves such that no three curves in $\mathcal{C}$ intersect at a single point and every pair of curves in $\mathcal{C}$ intersect at exactly one point which is either a crossing or a tangency point. Then the number of tangencies among the curves in $\mathcal{C}$ is $O(n)$.*

Györgyi, Hujter and Kisfaludi-Bak [60] proved Conjecture 2.38 for the special case where there are constantly many faces in the arrangement of $\mathcal{C}$ that together contain all the endpoints of the curves. We show that Conjecture 2.38 also holds for x-monotone curves.

Theorem 2.39 (Ackerman-Keszegh [6]). *Let $\mathcal{C}$ be a set of n x-monotone curves such that no three curves in $\mathcal{C}$ intersect at a single point and every pair of curves in $\mathcal{C}$ intersect at exactly one point which is either a crossing or a tangency point. Then the number of tangencies among the curves in $\mathcal{C}$ is at most $1160n$.*

2.4.2 Tangencies among a single family of curves

Next we consider the case when the curves are not necessarily pairwise intersecting. The number of tangencies within an at most 1-intersecting family of n curves can be $\Omega(n^{4/3})$, even if all curves are required to be x-monotone; this again follows from the construction of Erdős and Purdy [43] of n points and n lines that determine $\Omega(n^{4/3})$ point-line incidences by replacing each point with a small curve, and slightly perturbing the lines; see [98] (for an illustration consider Figure 42 and trim every curve beyond its rightmost touching point). For x-monotone at most 1-intersecting curves, an almost matching upper bound of $O(n^{4/3}\log^{2/3} n)$ follows from a result of Pach and Sharir [102] which can be improved with a more careful analysis to $O(n^{4/3}\log^{1/3} n)$.[9] See also [41, 44]. It also follows from [102] that for bi-infinite x-monotone 1-intersecting curves the maximum number of tangencies is $\Theta(n \log n)$.

[8] This can be seen by modifying the famous construction of Erdős and Purdy [43] of n points and n lines that determine $\Omega(n^{4/3})$ point-line incidences by replacing each point with an appropriate curve, and slightly perturbing the lines. See later Figure 42, which can be easily perturbed to be 2-intersecting. Variations of this construction will appear repeatedly.

[9] Personal communication with Eyal Ackerman, who observed this with Rom Pinchasi.

Concerning arrangements of (pseudo-)circles that are not necessarily pairwise intersecting, one can construct such arrangements of n circles with $\Omega(n^{4/3})$ digons based yet again on the construction of Erdős and Purdy of n lines and n points admitting that many point-line incidences (see [98]). The best known upper bound for pseudo-circles is $O(n^{3/2}\log n)$ by Marcus and Tardos [91]. A slightly better upper bound of $O(n^{3/2})$ for the number of touching points among n circles follows from a result of Ellenberg, Solymosi and Zahl [41]. For unit circles counting the number of tangencies is equivalent to the famous unit distance problem of Erdős for which the best known lower and upper bounds are $\Omega(n^{1+c/\log\log n})$ [42] and $O(n^{4/3})$ [108, 132, 135], respectively.

Pach, Rubin and Tardos [100, 101] settled a long-standing conjecture of Richter and Thomassen [121], showing that the number of crossing points (i.e., intersections that are *not* tangencies) determined by pairwise intersecting curves is $\Omega(1-o(1)n^2$. In particular, they showed that in any set of (not necessarily pairwise crossing) curves having at least linearly many tangencies, the number of crossing points is superlinear with respect to the number of tangencies. This implies that for any fixed t every set of n t-intersecting curves admits $o(n^2)$ tangencies.

Also motivated by the conjecture of Richter and Thomassen, Salazar [124] already pointed this out for families which are also pairwise intersecting. More specifically, Salazar showed that if the curves are pairwise intersecting, then for some large enough $c = c(t)$ the tangency graph of any family of at most t-intersecting curves avoids $K_{2t,c}$ as a subgraph. Then one can apply the Kővári–Sós–Turán theorem to conclude that the number of edges in the tangency graph is $o(n^2)$. More recently, Bechler-Speicher [17] adapted this idea to the case when the curves are not necessarily pairwise intersecting, showing that any family of n at most t-intersecting curves determines at most $O(n^{2-\frac{1}{3t+15}})$ tangencies. She proved that if there are many edges, then there is a doubly-grounded subfamily (we postpone its definition) of the curves which induces a big enough proportion of the tangencies, and that the tangency graph of such a family avoids $K_{t+5,c}$ as a subgraph for some large enough $c = c(t)$. Then again the Kővári–Sós–Turán theorem can be applied. We improve further this bound by showing that the tangency graph avoids a $K_{t+3,c}$ for some large enough $c = c(t)$. Note that we do not need to assume pairwise intersection or doubly-groundedness, also our proof seems to be simpler. As before, the next theorem then directly follows using the Kővári–Sós–Turán theorem.

Theorem 2.40 (Keszegh-Pálvölgyi [83])**.** *Any family of n at most t-intersecting curves determines at most $O(n^{2-\frac{1}{t+3}})$ tangencies.*

Specifically, it follows that n 1-intersecting curves determine $O(n^{7/4})$ tangencies while n 2-intersecting curves determine $O(n^{9/5})$ tangencies. We are not aware of any better lower bound (even for arbitrary t) than the aforementioned construction of a family of at most 1-intersecting curves with $\Omega(n^{4/3})$ tangencies.

Recently in [140] the author and its coauthors proved a generalization of Theorem 2.40 from curves to topological trees, using the same methods.

2.4.3 Tangencies among red and blue curves

Let us now turn our attention to the maximum number of tangencies between the members of two families (together having n curves), each of which consists of pairwise disjoint curves, where two curves are called tangent if they intersect exactly once, and in that intersection point they do not cross, but touch each other. Note that in this case the tangency graph is bipartite.

Pach, Suk, and Treml [104] attribute to Pinchasi and Ben-Dan (personal communication) the first result about this problem. They proved that the maximum number of tangencies among such a family of n curves is $O(n^{3/2}\log n)$. Their proof is based on a theorem of Marcus and Tardos [91], and of Pinchasi and Radoičić [116]; see also the discussion after Claim 11.3. Moreover, Pinchasi and Ben-Dan suggested that the correct order of magnitude may be linear in n.

Pach, Suk, and Treml [104] proved this conjecture in the special case where both families consist of closed convex regions instead of arbitrary curves. Ackerman [1] have improved the multiplicative constant from 8 to 6 for this case, which is asymptotically optimal.

On the other hand, Pach, Suk, and Treml [104] refuted the conjecture in general by showing an example that the number of tangencies may be more than linear. More precisely, they constructed two families of n pairwise disjoint x-monotone curves with $\Omega(n\log n)$ tangencies between the two families. They have also shown an upper bound of $O(n\log^2 n)$ for x-monotone curves. We determine the exact order of magnitude for x-monotone curves, by giving the following improved upper bound, matching the previous lower bound.

Theorem 2.41 (Keszegh-Pálvölgyi [83])**.** *Given a family of n red and blue x-monotone curves such that no two curves of the same color intersect, the number of tangencies among the curves is $O(n\log n)$.*

Having a closer look at the lower bound construction of Pach et al. [104] with $\Omega(n\log n)$ curves, one can notice that the two families together form an at most 2-intersecting family, i.e., any pair of red and blue curves intersects at most twice. This raises the question if assuming that the family is at most 1-intersecting leads to a better upper bound. We prove that this is indeed the case, i.e., given an at most 1-intersecting family of n red and blue curves such that no two curves of the same color intersect, the number of tangencies between the curves is $O(n)$:

Theorem 2.42 (Ackerman-Keszegh-Pálvölgyi [7])**.** *Given a 1-intersecting family of n red and blue curves such that no two curves of the same color intersect, the number of tangencies among the curves is $O(n)$.*

Note that it is trivial to construct an example with $\Omega(n)$ tangencies and we have seen that Theorem 2.42 does not hold if a pair of curves in $\mathcal{S}$ may intersect twice.

For general (not necessarily x-monotone) curves our main result is a construction which improves the previous $\Omega(n\log n)$ lower bound considerably. Our lower bound construction has the special property that it contains a red curve and a blue curve that each touch all curves of the other color. Alternately, the construction can be realized such that all curves lie within a vertical strip, every red curve touches the left boundary of the strip and every blue curve touches the right boundary of the strip. We call such a family of curves a *doubly-grounded family.*[10]

Theorem 2.43 (Keszegh-Pálvölgyi [83])**.** *There exists a family of n red and blue curves such that no two curves of the same color intersect and the number of tangencies between the curves is $\Omega(n^{4/3})$. Moreover, the family can be doubly-grounded.*

There still remains a polynomial gap between this lower bound and the best known upper bound of $O(n^{3/2}\log n)$. Even getting rid of the $\log n$ from the upper bound would be an interesting improvement.

To complement this lower bound, we show that for doubly-grounded families this is best possible, provided a conjecture about forbidden 0-1 matrices of Pach and Tardos [108] holds:

[10] Note that a slightly different, but combinatorially equivalent definition was used in [52].

Theorem 2.44 (Keszegh-Pálvölgyi [83]). *Assume that the following holds: If a* 0-1 *matrix avoids all positive orthogonal cycles, then it has* $O(n^{4/3})$ 1*-entries.*

Then the following is also true: The number of tangencies is $O(n^{4/3})$ *between* n *red and blue curves that form a doubly-grounded family, i.e., that lie within a vertical strip such that no two curves of the same color intersect, every red curve touches the left boundary of the strip and every blue curve touches the right boundary of the strip.*

The missing definitions related to 0-1 matrices are postponed to Section 11.3. As an easy first step of the proof we get a slight improvement over the best known general upper bound $O(n^{3/2}\log n)$ for doubly-grounded families (which we think was folklore already):

Claim 2.45 (Keszegh-Pálvölgyi [83]). *Given a doubly-grounded family of* n *red and blue curves, the number of tangencies between the curves is* $O(n^{3/2})$.

The known results about the number of tangencies among two families of pairwise disjoint curves are summarized in Table 1.

Curve type	**# of tangencies**	**Lower bd. ref**	**Upper bd. ref**
general	$\Omega(n^{4/3})$, $\tilde{O}(n^{3/2})$	Thm. 2.43	Pinchasi & Ben-Dan [104]
doubly-grounded	$\Omega(n^{4/3})$, $O(n^{3/2})$, Conj. 11.5 $\Rightarrow\Theta(n^{4/3})$	Thm. 2.43	Claim 2.45 Thm. 2.44
x-monotone	$\Theta(n\log n)$	Pach et al. [104]	Thm. 2.41
1-intersecting	$\Theta(n)$	trivial	Thm.2.42
convex regions	$\Theta(n)$	trivial	Pach et al. [104]

Table 1: Summary of results for two families of disjoint curves.

Connecting curves. Instead of considering touching points among curves in a family of curves $\mathcal{S}$, we may consider pairs of disjoint curves that are intersected by a curve c from a different family of curves $\mathcal{C}$, such that c does not intersect any other curve. Indeed, each touching point of two curves from $\mathcal{S}$ can be replaced by a new, short curve that connects the two previously touching curves that become disjoint by redrawing one of them near the touching point. Conversely, if $\mathcal{C}$ consists of curves that connect *disjoint* curves from $\mathcal{S}$, then each connecting curve in $\mathcal{C}$ can be replaced by a touching point between the corresponding curves by redrawing one of these two curves. Therefore, studying touching points of such connecting curves are equivalent problems. This gives the following reformulation of Theorem 2.42.

Theorem 2.46 (Ackerman-Keszegh-Pálvölgyi [7]). *Let* $\mathcal{S}$ *be a set of* 1*-intersecting* n *red and blue curves such that no two curves of the same color intersect. Suppose that* $\mathcal{C}$ *is a set of pairwise disjoint curves such that each of them intersects exactly a distinct pair of disjoint curves from* $\mathcal{S}$*. Then* $|\mathcal{C}| = O(n)$.

If $\mathcal{S}$ and $\mathcal{C}$ are two families of curves, then we say that $\mathcal{C}$ is *grounded* with respect to $\mathcal{S}$ if there is a connected region of $\mathbb{R}^2\setminus\mathcal{S}$ that contains at least one point of every curve in $\mathcal{C}$. If $\mathcal{C}$ is grounded with respect to $\mathcal{S}$, then we can drop the assumption that $\mathcal{S}$ is 1-intersecting and prove the following variant.

Theorem 2.47 (Ackerman-Keszegh-Pálvölgyi [7]). *Let $\mathcal{S}$ be a set of n red and blue curves such that no two curves of the same color intersect. Suppose that $\mathcal{C}$ is a set of pairwise disjoint curves grounded with respect to $\mathcal{S}$, such that each of them intersects exactly a distinct pair of curves from $\mathcal{S}$. Then $|\mathcal{C}| = O(n)$.*

Note that we have also dropped the assumption that a red curve and a blue curve which are connected by a curve in $\mathcal{C}$ are disjoint. Therefore it is essential that $\mathcal{C}$ is grounded with respect to $\mathcal{S}$, for otherwise we might have $|\mathcal{C}| = \Omega(n^2)$. Indeed, let $\mathcal{S}$ be a (1-intersecting) set of $n/2$ horizontal segments and $n/2$ vertical segments such that every horizontal segment and every vertical segment intersect. Then each such pair can be connected by a curve in $\mathcal{C}$ very close to their intersection point. Hence $|\mathcal{C}| = n^2/4$.

Clearly, instead of *curves*, Theorem 2.47 could also be stated with a red and a blue family of disjoint *shapes*, with no requirement at all about the shapes, except that each of them is connected.

We can use Theorem 2.47 to improve a result of Keller, Rok and Smorodinsky [70] about conflict-free colorings of L-shapes. We prove in particular that given a family $\mathcal{L}$ of n L-shapes, all touching the x-axis from below, it is possible to color every L-shape in $\mathcal{L}$ with one of $O(\log^2 n)$ colors such that for each $\ell \in \mathcal{L}$ there is an L-shape with a unique color among the L-shapes whose intersection with ℓ is non-empty. The exact definitions, proofs and consequences can be found in Section 10.4.

3 On hypergraphs whose Delaunay graphs have linear size

The results of this chapter appeared in [7], a joint work with Ackerman and Pálvölgyi.

The upper bound on the number of edges in the Delaunay graph implies upper bounds for the number of hyperedges of size at most k, the chromatic number of the hypergraph and its VC-dimension.[11] This was already shown, e.g., for *pseudo-disks* [13, 24], however, the same arguments apply in general. We summarize these facts in the following statement.

Theorem 3.1. *Let $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ be an n-vertex hypergraph. Suppose that there exist absolute constants $c, c' \geq 0$ such that for every $\mathcal{V}' \subseteq \mathcal{V}$ the Delaunay graph of the sub-hypergraph[12] induced by $\mathcal{V}'$ has at most $c|\mathcal{V}'| - c'$ edges, then:*

(i) The chromatic number of $\mathcal{H}$ is at most $2c + 1$ (at most $2c$ if $c' > 0$);

(ii) The VC-dimension of $\mathcal{H}$ is at most $2c + 1$ (at most $2c$ if $c' > 0$); and

(iii) $\mathcal{H}$ has $O(k^{d-1}n)$ hyperedges of size at most k where d is the VC-dimension of $\mathcal{H}$.

Note that the well-known Sauer-Shelah lemma implies in addition that the hypergraph has $O(n^{d-1})$ hyperedges altogether, where d is the VC-dimension of $\mathcal{H}$.

Using a result of Chan et al. [30] this has another consequence about finding hitting sets. We follow the statement and the usage of this theorem as in [13] (see therein the definition of the minimum weight hitting set problem).

[11] Recall that the VC-dimension of a hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ is the size of its largest subset of vertices $\mathcal{V}' \subseteq \mathcal{V}$ that can be *shattered*, that is, for every subset $\mathcal{V}'' \subseteq \mathcal{V}'$ there exists a hyperedge $h \in \mathcal{E}$ such that $h \cap \mathcal{V}' = \mathcal{V}''$.

[12] The hyperedges of this sub-hypergraph are the non-empty subsets in $\{h \cap \mathcal{V}' \mid h \in \mathcal{E}\}$; this is sometimes called the *trace*, or the restriction to $\mathcal{V}'$.

Theorem 3.2. *[30] Let $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ be a hypergraph, where the number of edges of cardinality k for any restriction of $\mathcal{H}$ to a subset $\mathcal{V}' \subset \mathcal{V}$ is at most $O(|\mathcal{V}'|k^c)$, where $c > 0$ is an absolute constant and $k \leq |\mathcal{V}'|$ is an integer parameter. Then there exists a randomized polynomial-time $O(1)$-approximation algorithm for the minimum weight hitting set problem for $\mathcal{H}$.*

Theorem 3.1 and Theorem 3.2 together imply:

Theorem 3.3. *Let $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ be an n-vertex hypergraph. Suppose that there exist absolute constants $c, c' \geq 0$ such that for every $\mathcal{V}' \subseteq \mathcal{V}$ the Delaunay graph of the sub-hypergraph induced by $\mathcal{V}'$ has at most $c|\mathcal{V}'| - c'$ edges, then there exists a randomized polynomial-time $O(1)$-approximation algorithm for the minimum weight hitting set problem for $\mathcal{H}$.*

Proof of Theorem 3.1 (i). Let $\mathcal{H}$ be a hypergraph as assumed. Then the average degree of the Delaunay graph of every induced sub-hypergraph of $\mathcal{H}$ is at most $2c_1$ (and strictly less than $2c_1$ if $c_2 > 0$) and thus it has a vertex of that degree or smaller. Now we can easily get a proper $(2c_1 + 1)$-coloring (even a $2c_1$-coloring if $c_2 > 0$) by induction. Remove a vertex v with the smallest degree in the Delaunay graph of $\mathcal{H}$ and let $\mathcal{H}'$ be the hypergraph which is induced by the remaining vertices. As $\mathcal{H}'$ still has the above property, by induction it has a proper $(2c_1+1)$-coloring (even a $2c_1$-coloring if $c_2 > 0$). Now color v with a color that is different from the colors of all of its neighbors in the Delaunay graph of $\mathcal{H}$. We claim that this is a proper coloring of $\mathcal{H}$. Indeed, if a hyperedge contains at least two vertices other than v, then it is non-monochromatic by induction, otherwise it contains exactly two vertices, one of them being v and then it is non-monochromatic by the choice of color for v. □

Proof of Theorem 3.1 (ii). Suppose that that the VC-dimension of $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ is d. Then there exists a subset $\mathcal{V}' \subseteq \mathcal{V}$ of size d such that it induces a hypergraph containing all the subsets of $\mathcal{V}'$. In particular, it contains the $\binom{d}{2}$ hyperedges of size two, and thus $\binom{d}{2} \leq c_1 d - c_2$, by the properties of $\mathcal{H}$. Thus $d \leq \frac{(2c_1+1)+\sqrt{(2c_1+1)^2-8c_2}}{2}$. Since d is an integer, we have $d \leq 2c_1 + 1$ and $d \leq 2c_1$ if $c_2 > 0$. □

The proof of part (iii) is more involved, but it does not require new ideas; we follow standard techniques and the proofs in [13] and [24] (almost verbatim).

For a hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ and an integer $k \geq 2$ we denote by $\mathcal{E}_k$ (resp., $\mathcal{E}_{\leq k}$) the set of hyperedges whose size is exactly (resp., at most) k. We say that $\mathcal{H}$ is *c-linear* for some constant c, if $|\mathcal{E}'_2| \leq c|\mathcal{V}'|$ for every induced sub-hypergraph $(\mathcal{V}', \mathcal{E}')$. We first consider an upper bound for $|\mathcal{E}_{\leq k}|$, following the methods in [13, 24].

For $k \geq 2$, a pair of vertices of a hypergraph $\mathcal{H}$ is *k-good* if there exists a hyperedge of size at most k in $\mathcal{H}$ which contains both vertices.

Lemma 3.4. *Let $\mathcal{H}$ be an n-vertex c-linear hypergraph for some constant c. Then for every $k \geq 2$ there are at most $neck$ k-good pairs in $\mathcal{H}$.*[13]

Proof. Set $q = 1/k$ and remove every vertex of $\mathcal{H}$ independently with probability $1 - q$ to get a vertex set $\mathcal{V}'$ which induces a sub-hypergraph $\mathcal{H}' = (\mathcal{V}', \mathcal{E}')$. A k-good pair of $\mathcal{H}$ becomes a size-two hyperedge of $\mathcal{H}'$ with probability at least $q^2(1 - q)^{k-2}$. On the other hand $|\mathcal{E}'_2| \leq c|\mathcal{V}'|$ since $\mathcal{H}$ is c-linear. Let g denote the number of k-good pairs of $\mathcal{H}$. Then $gq^2(1-q)^{k-2} \leq \mathbb{E}[|\mathcal{E}'_2|] \leq c\mathbb{E}[|\mathcal{V}'|] = cqn$. Thus, we get $g \leq \frac{cn}{\frac{1}{k}(1-\frac{1}{k})^{k-2}} \leq neck$, as required. □

[13] The e in $neck$ stands for Euler's number.

Lemma 3.5 ([24]). *Let G be an n-vertex c-linear graph for some constant c. Then, for every $h \geq 2$, the number of copies of K_h (the complete graph on h vertices) in G is at most $t_{c,h}\, n$, where $t_{c,h} = \frac{(2c)^{h-1}}{h!}$.*

Corollary 3.6. *If $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ is a c-linear hypergraph, then $|\mathcal{E}_{\leq k}| \leq b_{c,k}\, n$, where $b_{c,k}$ depends only on c and k.*

Proof. Define a graph G whose vertex set is $\mathcal{V}$ such that there is an edge between each pair of k-good vertices. It follows from Lemma 3.4 that G is (eck)-linear. By Lemma 3.5 for every $2 \leq h \leq k$ the number of copies of K_h in G is at most $\frac{(2eck)^{h-1}}{h!}\, n$. Therefore $|\mathcal{E}_{\leq k}| \leq b_{c,k}\, n$ where $b_{c,k} = 1 + \sum_{h=2}^{k} \frac{(2eck)^{h-1}}{h!}$. □

Note that the constant $b_{c,k}$ in Corollary 3.6 is huge. Next, we will use this bound to obtain an upper bound of the form $|\mathcal{E}_{\leq k}| \leq O_{c,d}(k^{d-1}n)$, where d is the VC-dimension of $\mathcal{H}$ (often quite small).

The following is a well-known property of hypergraphs of bounded VC-dimensions, a stronger form of the Sauer-Shelah-lemma, used also in Buzaglo et al. [24].

Theorem 3.7. *Let $\mathcal{H}$ be a hypergraph with VC-dimension d. Then it is possible to assign to each hyperedge a subset of at most d of its vertices, such that distinct hyperedges are assigned distinct subsets.*

Proof of Theorem 3.1 (iii). Let $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ be an n-vertex c-linear hypergraph with VC-dimension at most d and let $k \geq 2$ be an integer. We wish to show that $|\mathcal{E}_{\leq k}| \leq O_c(k^{d-1}n)$ using an argument similar to the proof of Lemma 3.4.

By Theorem 3.7 we can assign to every hyperedge h of $\mathcal{H}$ a signature $h' \subseteq h$ of size at most d. Set $q = 1/k$ and remove every vertex of $\mathcal{H}$ independently with probability $1-q$ to get a vertex set $\mathcal{V}'$ that induces a sub-hypergraph $\mathcal{H}' = (\mathcal{V}', \mathcal{E}')$. We say that a hyperedge $h \in \mathcal{E}_{\leq k}$ *survives* if $h \cap \mathcal{V}' = h'$, where h' is the signature of h. Observe that a hyperedge h (of size at most k) survives with probability

$$q^{|h'|}(1-q)^{|h|-|h'|} \geq q^{|h'|}(1-q)^{k-|h'|} \geq q^d(1-q)^{k-d},$$

where the first inequality holds since $|h| \leq k$ and the second inequality holds since $q \leq 1-q$ and $|h'| \leq d$.

By Corollary 3.6 we have $|\mathcal{E}'_{\leq d}| \leq b_{c,d}|\mathcal{V}'|$, where $b_{c,d}$ is the constant from the corollary. Therefore,

$$q^d(1-q)^{k-d}|\mathcal{E}_{\leq k}| \leq \mathbb{E}[|\mathcal{E}'_{\leq d}|] \leq b_{c,d}\, \mathbb{E}[|\mathcal{V}'|] = b_{c,d}qn.$$

This implies that $|\mathcal{E}_{\leq k}| \leq b_{c,d}(1-q)^{d-k}q^{1-d}n \leq b_{c,d}\, k^{d-1}\, n$. Since $d \leq 2c+1$ by Theorem 3.1 (ii) we have that the number of hyperedges of size at most k is $O(k^{d-1}n)$ where the constant hiding in the big-O notation depends only on c. □

4 Pseudo-disks wrt. pseudo-disks

The results of this chapter appeared in [74], written by the author.

The primary aim of this chapter is to prove Theorem 2.12. We start with some definitions:

Definition 4.1. *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 $B \in \mathcal{B} \setminus \{B_1, B_2\}$.*

Definition 4.2. *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.*

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

Theorem 4.3. *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 2.22) this already implies Theorem 2.12 with 6 colors instead of 4. 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 hypergraph[14] and all planar hypergraphs can be generated this way. From this perspective Theorem 2.12 says that even with a much more relaxed definition of planarity of a hypergraph it remains 4-colorable.

The forthcoming sections are structured as follows. In Section 4.1 we prove Theorem 4.3. In Section 4.2 we prove Theorem 2.12. Implications about conflict-free colorings and VC-dimensions and the number of bounded size hyperedges in intersection hypergraphs is detailed in Section 4.3.

4.1 The Delaunay graph of pseudo-disks wrt. pseudo-disks

In this section we prove Theorem 4.3. First we list some tools we need from the papers of Pinchasi [115] and Buzaglo et al. [24] about pseudo-disks:

Lemma 4.4 ([115])**.** *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 4.5. *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 2.16):

Corollary 4.6. *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.*

[14] The most common definition of a planar hypergraph is that its bipartite incidence graph is planar.

The second part of Corollary 4.6 holds as for a point p contained in some $F \in \mathcal{F}$ we can apply Lemma 4.4.

By Corollary 4.6 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 4.4 will be useful when we deal with a shrinking-closed pseudo-disk family:

Corollary 4.7 ([115]). *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 $2 \leq \ell \leq k$ there exists a set $D' \subset D$ such that D' intersects B and exactly $\ell - 1$ other regions from $\mathcal{B}$, and $\mathcal{F} \cup \{D'\}$ is again a family of pseudo-disks.*

Lemma 4.8 ([24]). *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 4.9 ([115]). *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 4.9 becomes a special case of Theorem 4.3. Before proving Theorem 4.3 we prove another special case which for Jordan regions strengthens Lemma 4.9, while its proof remains relatively simple:

Lemma 4.10. *Given a family $\mathcal{F}$ of pseudo-disks and a finite family $\mathcal{B}$ of pseudo-disks such that each member $B \in B$ 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 [34, 138] 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., [119]).

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 Figure 9.

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

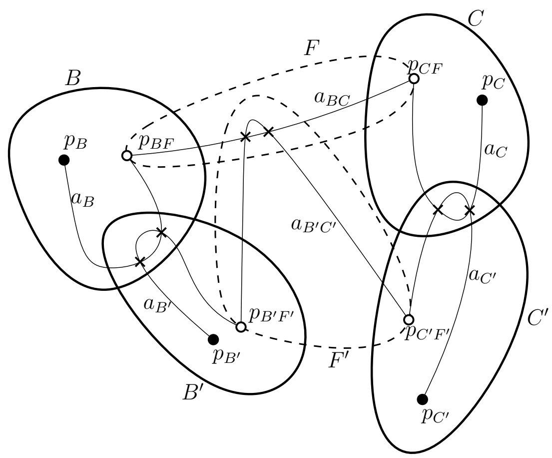

Figure 9: The edges e_f and $e_{f'}$ intersect an even number of times.

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 4.8 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 4.8 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. □

We note that Pach and Sharir [103] proved that (among others) pseudo-disks have linear union complexity using a similar approach as the proof of Lemma 4.10, connecting own-points of intersecting regions along their boundaries.

Corollary 4.11. *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 4.10 to conclude that the Delaunay graph (and so the restricted Delaunay graph as well) is planar. □

Observation 4.12. *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 2 members of $\mathcal{B}$.*

Corollary 4.11 shows that Lemma 4.10 takes care of the planarity of the restricted Delaunay graph. From Observation 4.12 we may think that we are left to take care of hyperedges that contain a point that is contained in at least 2 members of $\mathcal{B}$. This intuition turns out to be good, in all of our main results Lemma 4.10 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 4.3 we make some more preparations:

Definition 4.13. *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 4.14. *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 4.15. *If $\mathcal{H}''$ supports $\mathcal{H}'$ and $\mathcal{H}'$ supports $\mathcal{H}$, then $\mathcal{H}''$ supports $\mathcal{H}$.*

Observation 4.16. *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 the Sweeping theorem):

Theorem 4.17 ([131])**.** *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 such that every member of the sweeping of d intersects at most twice every other curve of Γ.*

A sweeping of d in Theorem 4.17 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 4.17 is stated in [131] for smooth curves but considering instead Jordan curves is not a significant difference.

Proof of Theorem 4.3. Using Corollary 4.6 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 $\mathcal{B}$ such that for the new family $\hat{\mathcal{B}}$ of pseudo-disks every $B \in \hat{\mathcal{B}}$ contains an own-point and furthermore we do this in a way that $\mathcal{I}(\hat{\mathcal{B}}, \mathcal{F})$ supports $\mathcal{I}(\mathcal{B}, \mathcal{F})$. Using Lemma 4.10 we get that the Delaunay graph of $\hat{\mathcal{B}}$ wrt. $\mathcal{F}$ is planar and then by Observation 4.16 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.

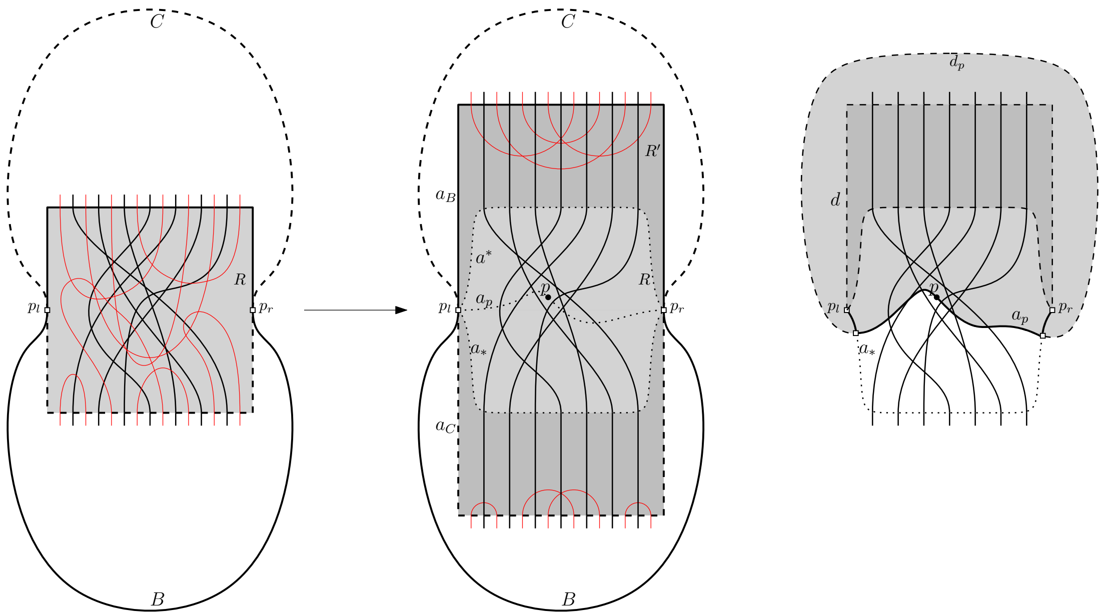

(a) (b)

Figure 10: Major steps of the operation used in the proof of Theorem 4.3: (a) pulling apart $B \cap C$, (b) drawing a_p.

- 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 such 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 Figure 10a 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 such 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 Figure 10a, 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 a $H' \in \mathcal{I}(\mathcal{B}', \mathcal{F})$ 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 Figure 10a 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 such 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 Figure 10b. 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 4.17. 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 4.17 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 Figure 10b. 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.[15]

Thus, with the above 3-step operation we changed $\mathcal{B}$ such 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 $\hat{\mathcal{B}}$ in which every member has an own-point and whose intersection hypergraph $\mathcal{I}(\hat{\mathcal{B}}, \mathcal{F})$ supports the intersection hypergraph $\mathcal{I}(\mathcal{B}, \mathcal{F})$ we started with, concluding the proof. □

4.2 Proper coloring pseudo-disks wrt. pseudo-disks

Note that Observation 2.3 can be rephrased equivalently such 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 4.18. *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 2.12. 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 $\mathcal{B}$ such 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 4.6 and Observation 4.18 we can assume that $\mathcal{F}$ is shrinking-closed (for $\mathcal{B}$).

Taking the direct product of a constant coloring provided by Theorem 2.8 and a 4-coloring of the restricted Delaunay graph which is planar by Corollary 4.11, it is not hard to see that we get a required coloring with a constant number of colors. Furthermore, even the existence of a proper 6 coloring follows immediately from Theorem 4.3 using known frameworks (e.g., [71] or Theorem 3.1).However, with the operation introduced in the proof of Theorem 4.3 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

[15] 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 et al. [9] 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.

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 4.3 we can color it with 4 colors, which by Observation 4.18 is a proper coloring of $\mathcal{I}(\mathcal{B},\mathcal{F})$, as required. This starts the induction.

Otherwise, H has $\ell \geq 3$ 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 4.7 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 $B, C \in B$ 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 4.3 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 4 colors, which by Observation 4.18 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 4.3, we redraw ∂C such that we change a_C to a^* (see again the right side of Figure 10a). 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 4 colors, which by Observation 4.18 is a proper coloring of $\mathcal{I}(\mathcal{B},\mathcal{F})$ as well. □

Observe that already for coloring points wrt. disks we may need 4 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).

4.3 Further consequences

First we prove Corollary 2.13 about conflict-free coloring the intersection hypergraph of pseudo-disks wrt. pseudo-disks.

Proof of Corollary 2.13. As mentioned already, the proof uses the by now standard framework of [128]. 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 2.12 ($c = 4$). 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 $\mathcal{B}_s = \mathcal{B}_{s-1} \setminus C_s$. 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. □

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. [13] 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 4.10 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 2.15. The proof is essentially the same as in [13], except that we use Lemma 4.10. 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 4.10. In fact, Theorem 4.3 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 4 points, $|V_s| \leq 4$, as required.

Finally, the construction of [24] 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. □

Theorem 2.16 follows from Theorem 2.15 and Theorem 3.1.

5 Family with linear union complexity wrt. pseudo-disks

The results of this chapter appeared in [74], written by the author.

Now we turn our attention to results about families with bounded linear union complexity.

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 $e(\partial\mathcal{B}) \leq v(\partial\mathcal{B}) + n$. □

Before proving Theorem 2.22 we make some preparations. First we prove an easy lemma using random sampling, that is, using the Clarkson-Shor method [35]. 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 $1/2^{k+3}s \le E(e(\partial\mathcal{B}')) \le E(c|\mathcal{B}'|) \le cn$ 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. □

Next we prove Theorem 2.22.We note that the proof we present here is much simplified compared to our original proof in [74].

Proof of Theorem 2.22. 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 $\mathcal{B}$ such 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.

We will prove that the Delaunay-graph of $\mathcal{B}$ with respect to $\mathcal{F}$ is c'-linear for some constant c', which then implies this also for every subhypergraph (as linear union complexity was defined hereditarily), and then from Theorem 3.1 it follows that there is a coloring with constant many colors, as required.

First, consider the edges of the Delaunay-graph that are already defined by $\mathcal{B}$ with respect to points. This has at most as many edges as the number of 2-deep faces in the arrangement of $\mathcal{B}$, which is linear by Lemma 5.2.

Second, consider the edges of the Delaunay-graph that are defined by a pseudo-disk from $\mathcal{F}$ that intersects exactly two 1-deep faces in the arrangement of $\mathcal{B}$ (that lie in different members of $\mathcal{B}$). The number of such edges is at most as much as the number of edges in the Delaunay of $\mathcal{B}'$ with respect to $\mathcal{F}$, where $\mathcal{B}'$ is the family of 1-deep faces in the arrangement of $\mathcal{B}$. By Lemma 5.2 $\mathcal{B}'$ has $O(n)$ members, and by definition these are pairwise disjoint regions. Now we can apply for example Lemma 4.9 to conclude that the number of such edges is also $O(n)$.

Note that every edge of the Delaunay graph is covered by one of these two cases, and therefore the total number of its edges is also $O(n)$, as needed. □

Proof of Corollary 2.23. The proof is exactly the same as of Corollary 2.23 just instead of Theorem 2.12 we need to use Theorem 2.22. □

Proof of Corollary 2.24. Note that in the proof of Theorem 2.22 we showed that the Delaunay-graph has $O(n)$ edges (for every subhypergraph). Thus we can apply Theorem 3.1 not just to

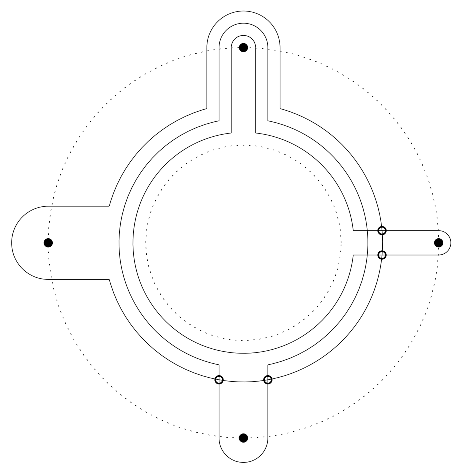

Figure 11: Construction with $n = 4$ points, 3 of the $\binom{4}{2}$ regions are shown, their union complexity is 4.

bound the chromatic number but also to bound the VC-dimension and the number of hyperedges of size t, and we indeed get the required bounds. □

We note that the above proof of Theorem 2.22 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 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 2.22 $\mathcal{F}$ cannot be changed to be a family of linear union complexity.

Claim 5.3. *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 Figure 11.

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 4 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 4 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 $\mathcal{F}'$. □

6 Lines wrt. pseudo-disks

The results of this chapter appeared in [69], a joint work with Keller and Pálvölgyi.

Our primary aim is to prove Theorem 2.25. First we present earlier results and simple lemmata that will be used in our proofs. We will also make use of Theorem 3.1.

The two following lemmata are standard useful tools when handling families of pseudo-disks:

Lemma 6.1 (Lemma 1 in [115], based on [131]). *Let $\mathcal{F}$ be a family of pseudo-disks, $D \in \mathcal{F}, x \in D$. Then D can be continuously shrunk to the point x, such that at each moment during the shrinking process, the family obtained from $\mathcal{F}$ remains a family of pseudo-disks.*

Lemma 6.2 (Lemma 2 in [115]). *Let $\mathcal{B}$ be a family of pairwise disjoint closed connected sets in $\Re^2$. Let $\mathcal{F}$ be a family of pseudo-disks. Define a graph G whose vertices correspond to the sets in B and connect two sets $B, B' \in \mathcal{B}$ if there is a set $D \in \mathcal{F}$ such that D intersects B and B' but not any other set from B. Then G is planar, hence $|E(G)| < 3|V(G)|$.*

A finite set $\mathcal{L}$ of lines in $\Re^2$ determines an *arrangement* $\mathcal{A}$. The 0-dimensional faces of $\mathcal{A}$ (namely, the intersections of two distinct lines from $\mathcal{L}$), are called *the vertices of* $\mathcal{A}$, the 1-dimensional faces are called *the edges of* $\mathcal{A}$, and the 2-dimensional faces are *the cells of* $\mathcal{A}$. Clearly, all cells are convex. The *cell complexity* of a cell f in $\mathcal{A}$, denoted by $comp(f)$, is the number of lines incident with the cell. The *zone* of an additional line ℓ, is the set of faces of $\mathcal{A}$ intersected by ℓ. The *complexity of a zone* is the sum of the cell complexities of the faces in the zone of ℓ, i.e., total number of edges of these faces, counted with multiplicities.

Theorem 6.3 (Zone Theorem [32]). *In an arrangement of n lines, the complexity of the zone of a line is $O(n)$.*

The best possible upper bound in the theorem is $\lfloor 9.5(n-1) \rfloor - 3$, obtained by Pinchasi [114].

We shall use a generalization of the theorem, for which an extra definition is needed. Given an arrangement $\mathcal{A}$ and a line ℓ, the 1-zone of ℓ is defined as the zone of ℓ, and for $t > 1$ the t-zone of ℓ is defined as the set of all faces adjacent to the $(t-1)$-zone, that do not belong to any i-zone for $i < t$. The $(\leq t)$*-zone* of ℓ is the union of the i-zones of ℓ for all $1 \leq i \leq t$.

The following generalization of the zone theorem was given as Exercise 6.4.2 in [92]. Its proof can be found in [133, Prop. 1].

Lemma 6.4 ([133]). *Let $\mathcal{A}$ be an arrangement of n lines. Then for any t, the $\leq t$-zone of any additional line ℓ contains at most $O(tn)$ vertices.*

By planarity, this implies:

Corollary 6.5. *Let $\mathcal{A}$ be an arrangement of n lines. Then for any t, the $\leq t$-zone of any additional line ℓ has complexity $C_{\leq t}(\ell) = O(tn)$.*

6.1 The number of t-hyperedges

In this section we prove Theorem 2.25. We prove the following stronger statement:

Proposition 6.6. *Let $\mathcal{L}$ be a family of n lines in the plane, let $\mathcal{F}$ be a family of pseudo-disks, and assume both families are in general position. Then for each $\ell \in \mathcal{L}$,*

$$|\{e \in \mathcal{E}(\mathcal{I}(\mathcal{L},\mathcal{F})) : |e| = t, \ell \in e\}| = O_t(n).$$

Consequently, $|\{e \in \mathcal{E}(\mathcal{I}(\mathcal{L},\mathcal{F})) : |e| = t\}| = O_t(n^2)$.

Proof of Proposition 6.6. First we prove the statement for hyperedges of size 3, and then we leverage the result to general hyperedges.

3-hyperedges. Fix a line ℓ. We observe that for a pseudo-disk c that defines a 3-hyperedge $\{\ell, \ell', \ell''\}$ there exists a cell of $\mathcal{A}(\mathcal{L} \setminus \{\ell\})$ which is in the ≤ 2-zone of ℓ in $\mathcal{A}(\mathcal{L} \setminus \{\ell\})$ such that c intersects two edges of this cell where one of these edges is on ℓ' and the second is on ℓ''. With every such pseudo-disk c we associate one such cell f_c and one such pair of edges of this cell, and denote this pair by e_c.

Define a graph $G = (V, E)$ whose vertices are all edges in the (≤ 2)-zone of ℓ in $\mathcal{A}(\mathcal{L} \setminus \{\ell\})$, and whose edges are the pairs e_c associated with the pseudo-disks that define a 3-hyperedge. Note that for any hyperedge $e = \{\ell, \ell', \ell''\}$ we choose exactly one pair of edges of $\mathcal{A}(\mathcal{L} \setminus \{\ell\})$ - one is on ℓ' and one is on ℓ'' - that form a corresponding edge of G. Thus by construction, $|E|$ is equal to the number of 3-hyperedges containing ℓ, and so, we want to prove that $|E| = O(n)$.

Consider a single cell f of $\mathcal{A}(\mathcal{L} \setminus \{\ell\})$. For each pseudo-disk c that defines a 3-hyperedge containing ℓ and has $f_c = f$, c does not intersect any other edge of f besides the two edges in e_c (as otherwise, c would intersect at least 4 lines of $\mathcal{L}$). Hence, the restriction of G to the edges of the cell f (after removing their endpoints), satisfies the assumptions of Lemma 6.2. Thus, by Lemma 6.2, the subgraph of G induced by the edges of f is planar, and hence, its number of edges is at most 3 times the complexity of f. Summing over all cells in the (≤ 2)-zone of ℓ, we obtain $|E| \leq 3 \sum_f comp(f) = O(n)$ by Corollary 6.5, and therefore, $|E| = O(n)$, as asserted.

t-hyperedges. Fix a line ℓ, and consider the hypergraph H' whose vertex set is $\mathcal{L} \setminus \{\ell\}$ and whose edge set is $\{e \setminus \{\ell\} : e \in \mathcal{E}(H), \ell \in e\}$. The 2-hyperedges of H' correspond to 3-hyperedges of H containing ℓ, and thus, by the first step, their number is $O(n)$. Furthermore, for any $\mathcal{L}' \subset \mathcal{L} \setminus \{\ell\}$, the number of 2-hyperedges in the restriction of H' to $\mathcal{L}'$ is $O(|\mathcal{L}'|)$, by the same argument. Therefore, H' satisfies the assumptions of Theorem 3.1, which implies that the VC-dimension d of H' is constant, and that the number C_{t-1} of $(t-1)$-hyperedges of H' is $O(t^{d-1}n)$.

Finally, the number of t-hyperedges of H that contain ℓ is equal to C_{t-1}. This completes the proof. □

6.2 The total number of hyperedges

In this section we prove Proposition 2.26.

Proof of Proposition 2.26. By Lemma 6.1 we can shrink the pseudo-disks one by one, such that the shrinking of each pseudo-disk $c \in \mathcal{F}$ is stopped when it becomes tangent to two lines. (Formally, first c is shrunk until the first time it is tangent to some line in $\mathcal{L}$, and then it is shrunk towards the tangency point until the next time it is tangent to some line in $\mathcal{L}$.) By the general position assumption, we can perform the shrinking process in such a way that the obtained geometric objects (i.e., lines and shrinked pseudo-disks) are also in general position. We replace each $c \in \mathcal{F}$ by its shrunk copy. Let $\mathcal{F}'$ be the obtained family. Then $\mathcal{I}(\mathcal{L}, \mathcal{F}) = \mathcal{I}(\mathcal{L}, \mathcal{F}')$, and by a tiny perturbation we can assume that all tangencies are in a point.

For any two lines $\ell_1, \ell_2 \in \mathcal{L}$, denote by $\mathcal{F}'(\ell_1, \ell_2)$ the set of all pseudo-disks in $\mathcal{F}'$ that are tangent to both ℓ_1 and ℓ_2. We claim that for any $\ell_1, \ell_2 \in \mathcal{L}$, $|\mathcal{E}(H(\mathcal{L}, \mathcal{F}'(\ell_1, \ell_2)))| = O(n)$, and this implies $|\mathcal{E}(H)| = O(n^3)$, the assertion of Proposition 2.26.

To show this, for any $c \in \mathcal{F}'(\ell_1, \ell_2)$, we define $x_{\ell_1,\ell_2}(c) = c \cap \ell_1 \in \Re^2$ and $y_{\ell_1,\ell_2}(c) = c \cap \ell_2 \in \Re^2$ (see Figure 12). In each of the four wedges that ℓ_1, ℓ_2 form, we define a linear order relation on the elements of $\mathcal{F}'(\ell_1, \ell_2)$: $c \prec c'$ if the segment $[x_{\ell_1,\ell_2}(c), y_{\ell_1,\ell_2}(c)]$ is completely above the segment $[x_{\ell_1,\ell_2}(c'), y_{\ell_1,\ell_2}(c')]$ (that is, if the points $x_{\ell_1,\ell_2}(c), y_{\ell_1,\ell_2}(c)$ are closer to the intersection point within the wedge than the points $x_{\ell_1,\ell_2}(c'), y_{\ell_1,\ell_2}(c')$, respectively).

First, we claim that this relation is well defined, since for $c \neq c'$ two such segments never intersect. Indeed, assume to the contrary they intersect, so that $y_{\ell_1,\ell_2}(c')$ is above $y_{\ell_1,\ell_2}(c)$, while $x_{\ell_1,\ell_2}(c')$ is below $x_{\ell_1,\ell_2}(c)$. The pseudo-disk c divides the remainder of the wedge into two connected components – the part 'above' it and the part 'below' it. Now, consider the points $x_{\ell_1,\ell_2}(c'), y_{\ell_1,\ell_2}(c')$. In the boundary of c', these points are connected by two curves. As these points are in different connected components wrt. c, each of these curves intersects c at least twice, which means that c, c' intersect at least 4 times, a contradiction.

Second, we claim that in each wedge, every line in $\mathcal{L}$ intersects a subset of consecutive elements of $\mathcal{F}'(\ell_1, \ell_2)$ under the order $\prec$. Indeed, assume that some line ℓ intersects two pseudo-disks c_1, c_3, as depicted in Figure 12. We want to show it must intersect c_2 as well. Like above, c_2 divides the wedge (without it) into two connected components. By the same argument as above, c_1 cannot intersect the component below c_2 (as otherwise, it would cross c_2 four times). Similarly, c_3 cannot intersect the component above c_2. Thus, either ℓ intersects at least one of c_1, c_3 inside c_2, or ℓ contains a point above c_2 and a point below c_2. In both cases, ℓ must intersect c_2.

Finally, by passing over all elements of $\mathcal{F}'(\ell_1, \ell_2)$ in each wedge, from the smallest to the largest, according to the order $\prec$, the number of times that the hyperedge defined by the current pseudo-disk is changed is linear in $|\mathcal{L}|$. Indeed, any such change is caused by appearance or disappearance of some line, and each line in $\mathcal{L}$ appears at most once and disappears at most once, along the process. Therefore, in each wedge, $|\mathcal{E}(H(\mathcal{L}, \mathcal{F}'(\ell_1, \ell_2)))| = O(n)$, and summing over all pairs $\{\ell_1, \ell_2\} \in \mathcal{L}$, we get $|\mathcal{E}(H)| = O(n^3)$. □

6.3 Open problems

Here we present a few open problems related to lines wrt. pseudo-disks.

Hypergraph of lines and inscribed pseudo-disks. A natural question is whether the arguments of Aronov et al. [14] can be extended from disks to pseudo-disks. We have found that all their arguments would go through if we knew that every triangle has an inscribed pseudo-disk. More precisely, we would need that for any triangle formed by three sides a, b, c, there is a pseudo-disk $d \in \mathcal{F}$, contained in the closed triangle, that intersects every side in exactly one point, or if there is no such $d \in \mathcal{F}$, then we can add such a new pseudo-disk d to $\mathcal{F}$ such that

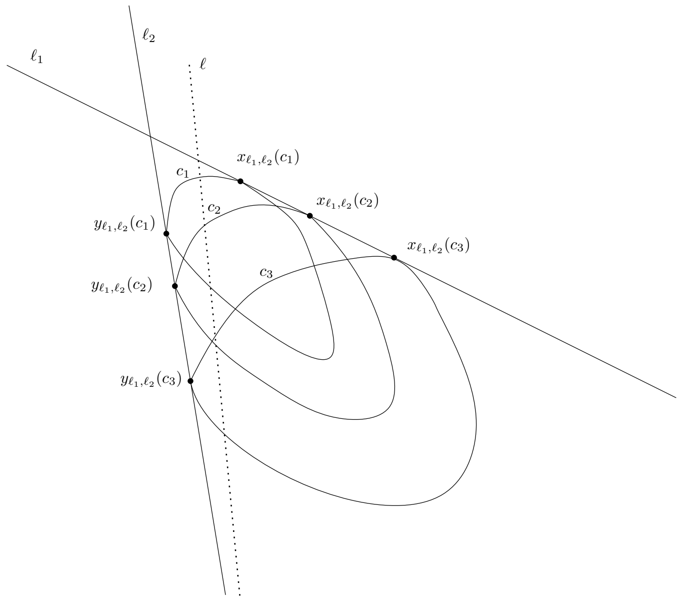

Figure 12: Illustration for the proof of Proposition 2.26 - c_1, c_2, c_3 are tangent to the lines ℓ_1, ℓ_2, and $c_1 \prec c_2 \prec c_3$.

$\mathcal{F} \cup \{d\}$ still forms a pseudo-disk family. Unfortunately, it seems that such a theory has not been developed yet, not even for $\mathcal{F}$ all whose elements are convex.

We note that for the related problem regarding circumscribed pseudo-disks, even a stronger result is known. Specifically, it was shown in [131, Thm. 5.1] that for any three points a, b, c, there is a pseudo-disk $d \in \mathcal{F}$ such that $a, b, c \in \partial d$, or if there is no such $d \in \mathcal{F}$, then we can add such a new pseudo-disk d to $\mathcal{F}$ such that $\mathcal{F} \cup \{d\}$ still forms a pseudo-disk family.

Dependence on t in Theorem 2.25. While we showed the quadratic dependence on n in Theorem 2.25 to be tight, the dependence on t is not clear. It seems plausible that

$$|\{e \in \mathcal{E}(\mathcal{I}(\mathcal{L}, \mathcal{F})) : |e| = t\}| = O(tn^2),$$

but we have not been able to prove this. On the other hand, even the stronger upper bound $O(n^2)$ for any fixed t, that would immediately imply Proposition 2.26 might hold.

Analogue of Theorem 3.1 for 3-sized hyperedges. It seems plausible that one can prove the following analogue of Theorem 3.1 for 3-sized hyperedges: If in some hypergraph on n vertices, for any induced hypergraph, the number of 3-sized hyperedges is quadratic in the number of vertices, then for any fixed t, the number of t-sized hyperedges is $O_t(n^2)$. Such a strong leveraging lemma would allow an easier proof of Theorem 2.25.

7 ABA-free hypergraphs and pseudo-halfplane hypergraphs

The results of this chapter appeared mostly in [81], a joint work Pálvölgyi, with the exception of some results (Section 7.4 and Theorem 7.50 from Section 7.6) that appeared in [75], written by the author.

In this chapter we prove our results related to ABA-free hypergraphs. We consider ABA-free hypergraphs in Section 7.1, pseudo-halfplane hypergraphs in Section 7.2, dual pseudo-halfplane hypergraphs in 7.3. As these sections consider mostly only polychromatic colorings, we consider the chromatic number of these hypergraphs in Section 7.4. Some consequences about ε-nets are mentioned in Section 7.5 and finally in Section 7.6 we detail the geometric consequences of these combinatorial results.

7.1 ABA-free hypergraphs and the general coloring algorithm

Suppose we are given an ABA-free hypergraph $\mathcal{H}$ on n vertices. As the hypergraph is ABA-free, for any pair of sets $A, B \in \mathcal{H}$ either there are $a < b$ such that $a \in A \setminus B$ and $b \in B \setminus A$, or there are $b < a$ such that $a \in A \setminus B$ and $b \in B \setminus A$, or none of them, but not both as that would contradict ABA-freeness.

Define $A < B$ if and only if there are $a < b$ such that $a \in A \setminus B$ and $b \in B \setminus A$, and $A \leq B$ if and only if either $A = B$ (as sets) or $A < B$. By the above, this is well-defined, and below we show that it gives a partial ordering of the sets.

Observation 7.1. *If $A < B$ and $a \in A \setminus B$, then there is a $b \in B \setminus A$ such that $b > a$.*

Proposition 7.2. *If $A < B$ and $B < C$, then $A < C$.*

Proof. Take an $a \in A \setminus B$. If $a \notin C$, then take a $b \in B \setminus A$. If $b \in C$, then $A < C$ and we are done. Otherwise, there has to be a $c > b$ such that $c \in C \setminus B$. If $c \in A$, then $a < b < c$ forms a forbidden sequence for A and B, thus $c \notin A$. Then by definition a and c show that $A < C$.

If $a \in C$, then also $a \in C \setminus B$, thus there has to be a $b_1 < a$ such that $b_1 \in B \setminus C$. As $a \in A \setminus B$ and $A < B$, we also have $b_1 \in A$ and so $b_1 \in A \setminus C$. There also has to be a $b_2 > a$ such that $b_2 \in B \setminus A$. If $b_2 \notin C$, then $b_1 < a < b_2$ forms a forbidden sequence for B and C. Thus $b_2 \in C \setminus A$, and by definition b_1 and b_2 show that $A < C$. $\square$

We proceed with another definition. Denote the smallest (resp. largest) element of an ordered set H by $\min(H)$ (resp. $\max(H)$).

Definition 7.3. *A vertex a is* skippable *if there exists an $A \in \mathcal{H}$ such that $\min(A) < a < \max(A)$ and $a \notin A$. In this case we say that A* skips *a. A vertex a is* unskippable *if there is no such A.*

Observation 7.4. *If a vertex a is unskippable in some ABA-free hypergraph $\mathcal{H}$, then after adding the one-element hyperedge $\{a\}$ to $\mathcal{H}$, it remains ABA-free.*

Note that the following two lemmas show that the unskippable vertices of an ABA-free hypergraph behave with respect to hyperedges similarly to how the vertices on the convex hull of a point set behave with respect to halfplanes. These two lemmas make it possible to use the framework of [130] on ABA-free hypergraphs.

Lemma 7.5. *If $\mathcal{H}$ is ABA-free, then every $A \in \mathcal{H}$ contains an unskippable vertex.*

Remark 7.6. Note that finiteness (recall that we have supposed that all our hypergraphs are finite) is needed, as the hypergraph whose vertex set is $\mathbb{Z}$ and hyperedge set is $\{\mathbb{Z} \setminus \{n\} \mid n \in \mathbb{Z}\}$ is ABA-free without unskippable vertices.

Proof of Lemma 7.5. Take an arbitrary set $A \in \mathcal{H}$, suppose that it does not contain an unskippable vertex, we will reach a contradiction. Call $a \in A$ *rightskippable* if there is a $B \in \mathcal{H}$ rightskipping a, that is for which $a \in A \setminus B$ and there are $b_1, b_2 \in B$ such that $b_1 < a < b_2$ where $b_2 \in B \setminus A$.

If A contains no unskippable vertex, $\max(A)$ must be rightskippable (any set skipping $\max(A)$ must also rightskip $\max(A)$). Also, $\min(A)$ cannot be rightskippable, as otherwise A and the set B rightskipping $\min(A)$ would violate ABA-freeness (we would get $b_1 < \min(A) < b_2$ where $b_1, b_2 \in B \setminus A, \min(A) \in A \setminus B$). Therefore we can take the largest $a \in A$ that is not rightskippable. By the assumption, it is skipped by a set, call it B, i.e., $b_1 < a < b_2$ where $b_1, b_2 \in B \not\ni a$. Moreover, suppose without loss of generality that b_2 is the smallest element of B which is bigger than a. Since a is not rightskippable, $b_2 \in A$ must also hold. As $b_2 \in A$ is rightskippable, there is a C such that $c_1 < b_2 < c_2$ where $c_1, c_2 \in C$ and $b_2 \notin C, c_2 \notin A$. Without loss of generality, suppose that c_1 is the largest element of C which is smaller than b_2. If $c_1 < a$, then C would rightskip a, a contradiction. Thus, $b_1 < a \leq c_1$, and from the choice of b_2 we conclude that $c_1 \notin B$. As $c_2 \notin A$, also $c_2 \notin B$, otherwise B would rightskip a. Putting all together, we get $c_1 < b_2 < c_2$, thus B and C contradict ABA-freeness. □

Lemma 7.7. *If $\mathcal{F}$ is an ABA-free hypergraph on vertex set V and $v \in V$ is skippable, then every hyperedge H which contains v must contain at least one of the two unskippable vertices before and after v that are closest to v.*

Proof. Assume on the contrary. Let ℓ (resp. r) be the closest unskippable vertex to v left to v (resp. right to v). By Lemma 7.5 H contains some unskippable vertex w different from ℓ and r. If w is left to ℓ (resp. right to r) then H skips ℓ, contradicting that ℓ (resp. r) is unskippable. Thus $\ell < w < r$ in the vertex order, contradicting that ℓ and r were the closest unskippable vertices to v. □

Definition 7.8. *A hypergraph is called* containment-free *if none of its hyperedges contains another hyperedge.*[16] *A hypergraph $\mathcal{H}'$ is a* subhypergraph *of a hypergraph $\mathcal{H}$ on vertex set S if we can get $\mathcal{H}'$ by taking a subset $S' \subset S$ as its vertex set and the family of the hyperedges of $\mathcal{H}'$ is a subfamily of the hyperedges of $\mathcal{H}$ restricted to S'. We call a hypergraph property $\mathcal{P}$* hereditary *if for every hypergraph $\mathcal{H}$ that has property $\mathcal{P}$, all of its subhypergraphs also have property $\mathcal{P}$.*

Observation 7.9. *ABA-freeness is a hereditary property.*

We further assume in the rest that our hypergraphs are nonempty in the sense that they contain at least one hyperedge which is not the empty set. Notice that for an ABA-free containment-free hypergraph the ordering $<$ of its sets is a total order, i.e., any two hyperedges are comparable.

To study polychromatic coloring problems, we also introduce the following definition, which is implicitly used in [130], but deserves to be defined explicitly as it seems to be important in the study of polychromatic colorings. Actually, almost all our results about ABA-free hypergraphs are based on shallow hitting sets.

Definition 7.10. *A set R is a c-*shallow hitting set *of the hypergraph $\mathcal{H}$ if for every $H \in \mathcal{H}$ we have $1 \leq |R \cap H| \leq c$.*

Lemma 7.11. *If $\mathcal{H}$ is ABA-free and containment-free, then any minimal hitting set of $\mathcal{H}$ that contains only unskippable vertices is 2-shallow.*

[16]Equivalently, the hyperedges form an *antichain*. This property is also called *Sperner*.

Proof. Let R be a minimal (for containment) hitting set of unskippable vertices. Assume to the contrary that there exists a set A such that $|A \cap R| \geq 3$. Let $\ell = \min(A \cap R)$ and $r = \max(A \cap R)$. There exists a third vertex $\ell < a < r$ in $A \cap R$. We claim that $R' = R \setminus \{a\}$ hits all sets of $\mathcal{H}$, contradicting the minimality of R. Assume on the contrary that R' is disjoint from some $B \in \mathcal{H}$. As R must hit B, we have $R \cap B = \{a\}$. If there is a $b \in B \setminus A$ such that $\ell < b < r$, that would contradict the ABA-free property. If there is a $b \in B$ such that $b < \ell < a$ or $a < r < b$, that would contradict that ℓ and r are unskippable. Thus $B \subset A$, contradicting that $\mathcal{H}$ is containment-free. □

Lemma 7.12. *Every containment-free ABA-free hypergraph has a 2-shallow hitting set.*

Proof. Given a containment-free ABA-free hypergraph, take the set of all unskippable vertices, it is a hitting set by Lemma 7.5. Then we can delete vertices from this set until it becomes a minimal hitting set, which is 2-shallow by Lemma 7.11. □

Now we present an abstract and generalized version of the framework of [130] to give polychromatic k-colorings of hypergraphs.

Theorem 7.13. *Assume that $\mathcal{P}$ is a hereditary hypergraph property such that every containment-free hypergraph with property $\mathcal{P}$ has a c-shallow hitting set. Then every hypergraph $\mathcal{H}$ with hyperedges of size at least $ck - (c-1)$ that has property $\mathcal{P}$ admits a polychromatic k-coloring, i.e., a coloring of its vertices with k colors such that every hyperedge of $\mathcal{H}$ contains vertices of all k colors.*

Proof. We present an algorithm that gives a polychromatic k-coloring. First, we repeat $k-1$ times ($i = 1, \ldots, k-1$) the *general step* of the algorithm:

At the beginning of step i we have a hypergraph $\mathcal{H}$ with hyperedges of size at least $ck - ci + 1$ that has property $\mathcal{P}$. If any hyperedge contains another, then delete the bigger hyperedge. Repeat this until no hyperedge contains another, thus making our hypergraph containment-free. Next, take a c-shallow hitting set (using our assumptions), and color its vertices with the i-th color. Delete these vertices from $\mathcal{H}$ (the hyperedges of the new hypergraph are the ones induced by the remaining vertices). As $\mathcal{P}$ is hereditary, the new hypergraph also has property $\mathcal{P}$ and we can proceed to the next step.

After $k-1$ iterations of the above, we are left with a 1-uniform hypergraph whose vertices we can color with the k-th color. □

First, we use this algorithm to give a polychromatic k-coloring of the vertices of an ABA-free hypergraph with hyperedges of size at least $2k - 1$.

Theorem 7.14. *Given an ABA-free $\mathcal{H}$ we can color its vertices with k colors such that every $A \in \mathcal{H}$ whose size is at least $2k - 1$ contains all k colors.*

Proof. By Observation 7.9 ABA-freeness is a hereditary property. Together with Lemma 7.12 we get that all the assumptions of Theorem 7.13 with $c = 2$ hold for ABA-free hypergraphs with hyperedges of size at least $2k - 1$ and thus we get a required k-coloring. □

Notice that the above theorem is sharp, as taking $\mathcal{H}$ to be all subsets of size $2k - 2$ from $2k - 1$ vertices, in any coloring of the vertices, one color must occur at most once and is thus missed by some hyperedge.

We state another corollary of Lemma 7.5 that we need later. Before that, we need another simple claim.

Proposition 7.15. *If we insert a new vertex, v, somewhere into the (ordered) vertex set of an ABA-free hypergraph, $\mathcal{H}$, and add v to every hyperedge that contains a vertex before and another vertex after v, then we get an ABA-free hypergraph.*

Proof. We show that if in the new hypergraph, $\mathcal{H}'$, two hyperedges A' and B' violate ABA-freeness, then we can find two hyperedges A and B in the original hypergraph, $\mathcal{H}$, that also violate ABA-freeness, which would be a contradiction. We define $A = A' \setminus \{v\}$ and $B = B' \setminus \{v\}$. If both A' and B' contain or do not contain v, then by definition A and B also violate the condition. If, say, $v \in A'$ and $v \notin B'$, then without loss of generality we can suppose that all the vertices of $B = B'$ are before v. This means that if there are $x < y < z$ such that $x, z \in A' \setminus B'$ and $y \in B' \setminus A'$, then necessarily $v = z$. But as A' has an element z' that is bigger than v, we have $x, z' \in A \setminus B$ and $y \in B \setminus A$, a contradiction. □

Lemma 7.16. *If $\mathcal{H}$ is ABA-free, $A \in \mathcal{H}$, then there is a vertex $a \in A$ such that $\mathcal{H} \cup \{A \setminus \{a\}\}$ is also ABA-free.*

Proof. If $|A| = 1$, then trivially $\mathcal{H}$ can be extended with $\emptyset$. If $|A| > 1$, then we proceed by induction on the size of A. Using Lemma 7.5, there is an unskippable vertex $v \in A$. Delete this vertex from $\mathcal{H}$ to obtain some ABA-free $\mathcal{H}_v$ and let $A_v = A \setminus \{v\}$. Using induction on A_v, there is an $A'_v = A_v \setminus \{a\}$ such that $\mathcal{H}_v \cup \{A'_v\}$ is also ABA-free. We claim that with $A' = A'_v \cup \{v\} = A \setminus \{a\}$, the family $\mathcal{H} \cup \{A'\}$ is also ABA-free.

Notice that adding back v to $\mathcal{H}_v$ is very similar to the operation of Proposition 7.15, as v is unskippable in $\mathcal{H}$. The only difference is that we might also have to add it to some further hyperedges, ending in or starting at v. But a hyperedge that contains v cannot violate the ABA-free condition with A', since it also contains v, so the corresponding hyperedges in $\mathcal{H}_v$ would also violate the ABA-free condition. □

Notice that with the repeated application of Lemma 7.16 we can extend any ABA-free hypergraph, such that in any set A there is a vertex a for which $\{a\}$ is a singleton hyperedge, implying that a is unskippable in A. Thus in fact Lemma 7.16 is equivalent to Lemma 7.5. Moreover, in Section 7.2, in the more general context of pseudo-halfplanes, it will be the abstract equivalent of a known and important property of pseudo-halfplanes.

We prove another interesting property of ABA-free hypergraphs before which we need the following definition.

Definition 7.17. *The dual of a hypergraph $\mathcal{H}$, denoted by $\mathcal{H}^*$, is such that its vertices are the hyperedges of $\mathcal{H}$ and its hyperedges are the vertices of $\mathcal{H}$ with the same incidences as in $\mathcal{H}$.*

Proposition 7.18. *If $\mathcal{H}$ is ABA-free, then its dual $\mathcal{H}^*$ is also ABA-free (with respect to some ordering of its vertices).*

Proof. Take the partial order "$<$" of the hyperedges of $\mathcal{H}$ and extend this arbitrarily to a total order $<^*$. We claim that $\mathcal{H}^*$ is ABA-free if its vertices are ordered with respect to $<^*$. To check the condition, suppose for a contradiction that $H_x <^* H_y <^* H_z$ and $a \in (H_x \cap H_z) \setminus H_y$ and $b \in H_y \setminus (H_x \cup H_z)$. Without loss of generality, suppose that $a < b$. But in this case $H_z < H_y$ holds, contradicting $H_y <^* H_z$. □

Corollary 7.19. *The hyperedges of every ABA-free hypergraph can be colored with k colors, such that if a vertex v is in a subfamily $\mathcal{H}_v$ of at least $m(k) = 2k - 1$ of the hyperedges of $\mathcal{H}$, then $\mathcal{H}_v$ contains a hyperedge from each of the k color classes.*

Corollary 7.20. *Any $(2k-1)$-fold covering of a finite point set with the translates of an unbounded convex planar set is decomposable into k coverings.*

In fact, there is a slightly different proof for Proposition 7.18. For that we give an equivalent definition of ABA-free hypergraphs in relation to their incidence matrices, which will be useful also for other purposes later. In an *incidence matrix* of a hypergraph $\mathcal{H}$, rows correspond to the vertices of $\mathcal{H}$, columns correspond to the hyperedges of $\mathcal{H}$. An entry is 1 if the hyperedge corresponding to the column contains the vertex corresponding to the row, and 0 otherwise. Note that this is not unique as we can order the rows and columns arbitrarily. We say that a matrix M *contains* another matrix P if P is a submatrix of M. If M does not contain P, then it is called *P-free*.

Theorem 7.21. *Given a hypergraph $\mathcal{H}$, the following are equivalent:*

(a) *$\mathcal{H}$ is an ABA-free hypergraph,*

(b) *there is a permutation of the rows of the incidence matrix of $\mathcal{H}$ such that the matrix becomes $\begin{bmatrix} 0 & 1 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}$-free and $\begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 0 \end{bmatrix}$-free,*

(c) *there is a permutation of the rows and columns of the incidence matrix of $\mathcal{H}$ such that the matrix becomes $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$-free.*

Proof. First, ordering the vertices of the hypergraph corresponds to permuting the rows of its incidence matrix. Thus, the equivalence of (a) and (b) follows from the definition of ABA-free hypergraphs.

To prove $(c) \to (b)$, suppose (b) is false, i.e., that in any permutation of the rows of the incidence matrix of $\mathcal{H}$ there is an occurrence of one of the two matrices forbidden in (b). In any permutation of the two columns of these two matrices forbidden in (b), we get back one of these two matrices, both of which contains a copy of $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$. Thus by any permutation of the rows and columns of the incidence matrix of $\mathcal{H}$, we get a matrix that contains $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$. Thus we can conclude that $\neg(b) \to \neg(c)$, which is the contrapositive of $(c) \to (b)$.

Finally, extending to a complete order the partial ordering "$<$" defined on the hyperedges at the beginning of this section, Proposition 7.2 implies that by permuting the columns according to any extension of this order "$<$" of the hyperedges we get a $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$-free matrix, and thus $(b) \to (c)$. □

Now observe that in Theorem 7.21 the property in (c) holds for an incidence matrix if and only if it holds for its transpose (as the forbidden matrix in (c) is its own transpose). Taking the transpose of an incidence matrix in terms of the hypergraph means taking the dual of the hypergraph, thus Proposition 7.18 follows.

7.2 Pseudo-halfplane hypergraphs

Definition 7.22. *A hypergraph $\mathcal{H}$ on an ordered set of points S is called a* pseudo-halfplane hypergraph *if there exists an ABA-free hypergraph $\mathcal{F}$ on S such that $\mathcal{H} \subset \mathcal{F} \cup \bar{\mathcal{F}}$.*

Note that $\bar{\mathcal{F}}$ is also ABA-free with the same ordering of the points. We refer to the hyperedges of a pseudo-halfplane hypergraph also as pseudo-halfplanes.

Using Lemma 7.16 on a hyperedge of a pseudo-halfplane hypergraph, we get the following.

Proposition 7.23. *Given a pseudo-halfplane hypergraph $\mathcal{H}$, and a hyperedge A of $\mathcal{H}$, we can add a new hyperedge A' contained completely in A that contains all but one of the points of A, such that $\mathcal{H}$ remains a pseudo-halfplane hypergraph.*

In the geometric setting this corresponds to the known and useful fact that given a pseudo-halfplane arrangement and a finite set of points A contained in the pseudo-halfplane H, we can add a new pseudo-halfplane H' contained completely in H that contains all but one of the points of A.

Now we show how to extend Theorem 7.14 to pseudo-halfplane arrangements, i.e., to the case when the points of S *below* a line also define a hyperedge.

Theorem 7.24. *Given a finite set of points S and a pseudo-halfplane arrangement $\mathcal{H}$, we can color S with k colors such that any pseudo-halfplane in $\mathcal{H}$ that contains at least $2k-1$ points of S contains all k colors. Equivalently, the vertices S of a pseudo-halfplane hypergraph can be colored with k colors such that any hyperedge containing at least $2k-1$ vertices contains all k colors.*

Remark 7.25. The similar statement is not true for the union of two arbitrary ABA-free hypergraphs (instead of an ABA-free hypergraph and its complement), as the union of two arbitrary ABA-free hypergraphs might not be 2-colorable, see [106] for such a construction.

Proof of Theorem 7.24. Our proof is completely about the abstract setting, yet it translates naturally to the geometric setting, also the figures illustrate the geometric interpretations.

By definition there exists an ABA-free $\mathcal{F}$ such that $\mathcal{H} \subset \mathcal{F} \cup \bar{\mathcal{F}}$. Call $\mathcal{U} = \mathcal{H} \cap \mathcal{F}$ the upsets and $\mathcal{D} = \mathcal{H} \cap \bar{\mathcal{F}}$ the downsets, observe that both $\mathcal{U}$ and $\mathcal{D}$ are ABA-free.

Further, the unskippable vertices of $\mathcal{U}$ (resp. $\mathcal{D}$) are called top (resp. bottom) vertices. The top and bottom vertices are called the extremal vertices of $\mathcal{H}$. Recall that by adding these extremal vertices as one-element hyperedges to $\mathcal{H}$, $\mathcal{H}$ remains to be a pseudo-halfplane hypergraph, as we can extend $\mathcal{F}$ and $\bar{\mathcal{F}}$ with the appropriate hyperedge (this is a convenient way of thinking about top/bottom vertices in the geometric setting, as seen later in the figures).

Observation 7.26. *If x is top and X is a downset and $x \in X$, then X contains all vertices that are bigger or all vertices that are smaller than x. The same holds if x is bottom, X is an upset and $x \in X$.*

Lemma 7.27. *If $\mathcal{H}$ is a containment-free pseudo-halfplane hypergraph, then any minimal hitting set of $\mathcal{H}$ that contains only extremal vertices is 2-shallow.*

Proof. Let R be a minimal hitting set of extremal vertices. Suppose for a contradiction that $\{a, b, c\} \subset R \cap X$ and $a < b < c$ for some $X \in \mathcal{H}$. Without loss of generality, suppose that b is top. As R is minimal, let B be a set for which $B \cap R = \{b\}$. From Observation 7.26 it follows that B is an upset.

First suppose that X is an upset. As $B \not\subset X$, take a $b_2 \in B \setminus X$. As B and X are both upsets and thus have the ABA-free property, we have $b_2 < a$ or $c < b_2$. Without loss of generality, we can suppose $c < b_2$. If c is top, $\{c\}$ and B violate ABA-freeness. See Figure 13a. If c is bottom, then using Observation 7.26, X contains all the vertices that are smaller than c. Take a set

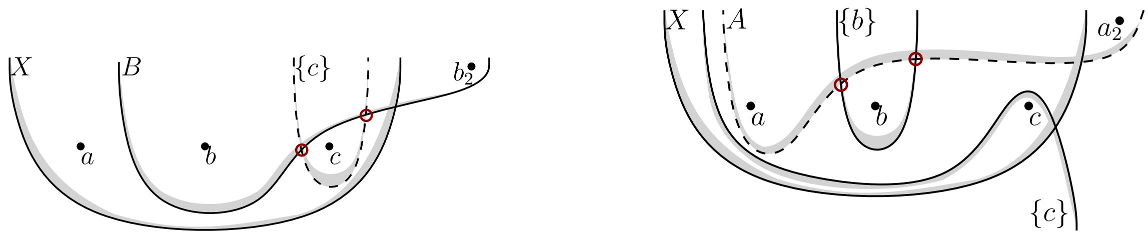

(a) c cannot be top

(b) A cannot be an upset

Figure 13: Proof of Lemma 7.27

$A \not\subset X$ for which $A \cap R = \{a\}$. This set must contain an $a_2 \in A \setminus X$ and so we must have $c < a_2$. If A is an upset, as it does not contain b and recall $a < b < a_2$, A and $\{b\}$ violate ABA-freeness. See Figure 13b. If A is a downset, as it does not contain c and recall $a < c < a_2$, A and $\{c\}$ violate ABA-freeness, both cases lead to a contradiction.

The case when X is a downset is similar. Using Observation 7.26 for X and $\{b\}$ we can suppose without loss of generality that X contains all vertices that are smaller than b. Take a set $A \not\subset X$ for which $A \cap R = \{a\}$ and an $a_2 \in A \setminus X$. As X contains all vertices smaller than b, we have $b < a_2$. A cannot be an upset, as then it would contain b, so it is a downset. If $b < a_2 < c$, then A and X would violate ABA-freeness, thus we must have $c < a_2$. This means c cannot be bottom, so it is top. Using Observation 7.26, X contains all the vertices that are smaller than c. But then $B \setminus X$ must have an element that is bigger than c, contradicting the ABA-freeness of B and $\{c\}$. □

It is easy to see that being a pseudo-halfplane hypergraph is a hereditary property. Thus, Lemma 7.27 implies that all the assumptions of Theorem 7.13 hold with $c = 2$ to get a polychromatic k-coloring as required. This finishes the proof of Theorem 7.24. □

7.3 Dual problem and pseudo-hemisphere hypergraphs

We are also interested in coloring pseudo-halfplanes with k colors such that all points that are covered many times will be contained in a pseudo-halfplane of each k colors. For example, we can also generalize the dual result about coloring halfplanes of [130] to pseudo-halfplanes.

Theorem 7.28. *Given a pseudo-halfplane arrangement $\mathcal{H}$, we can color $\mathcal{H}$ with k colors such that if a point p belongs to a subset $\mathcal{H}_p$ of at least $3k-2$ of the pseudo-halfplanes of $\mathcal{H}$, then $\mathcal{H}_p$ contains a pseudo-halfplane of every color.*

Theorem 7.28 follows from Theorem 7.33, that we will state and prove later.

However, instead of coloring pseudo-halfplanes, we stick to coloring points with respect to pseudo-halfplanes and work with *dual hypergraphs*, where the vertex-hyperedge incidences are preserved, but vertices become hyperedges and hyperedges become vertices.

We denote the symmetric difference of two sets, A and B, by $A\Delta B$, the complement of a hyperedge F by $\bar{F}$ and for a family $\mathcal{F}$ we use $\bar{\mathcal{F}} = \{\bar{F} \mid F \in \mathcal{F}\}$.

Proposition 7.29. *A hypergraph $\mathcal{H}$ on an ordered set of vertices S is a* dual pseudo-halfplane hypergraph *if and only if there exists a set $X \subset S$ and an ABA-free hypergraph $\mathcal{F}$ on S such that the hyperedges of $\mathcal{H}$ are the hyperedges $F\Delta X$ for every $F \in \mathcal{F}$ (where Δ denotes the symmetric difference of two sets).*

Proof. Recall that pseudo-halfplane hypergraphs are hypergraphs that we can get by taking the complement of some hyperedges in an ABA-free hypergraph.[17] In relation to their incidence matrix, using Theorem 7.21, this means that a hypergraph is a pseudo-halfplane hypergraph if and only if there is a permutation of the rows and columns of its incidence matrix such that inverting some of the columns (i.e., exchanging 0's and 1's in these columns) we get a matrix which is $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$-free. Taking the dual of such a hypergraph means taking the transpose of such an incidence matrix.

Thus a hypergraph $\mathcal{H}$ is a dual pseudo-halfplane hypergraph if and only if there is a permutation of the rows and columns of its incidence matrix such that inverting some of the rows we get a matrix which is $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$-free. Using again Theorem 7.21 we get that this is equivalent to the fact that the incidence matrix of $\mathcal{H}$ is the incidence matrix of an ABA-free hypergraph with some of the rows inverted. Finally, this is equivalent to the statement of the proposition with X being the subset of vertices corresponding to the inverted rows. □

Now we define a common generalization of the primal and dual definitions.

Definition 7.30. *A* pseudo-hemisphere hypergraph *is a hypergraph $\mathcal{H}$ on an ordered set of vertices S such that there exists a set $X \subset S$ and an ABA-free hypergraph $\mathcal{F}$ on S such that the hyperedges of $\mathcal{H}$ are some subset of $\{F\Delta X, \bar{F}\Delta X \mid F \in \mathcal{F}\}$.*

Proposition 7.31. *The dual of a pseudo-hemisphere hypergraph is also a pseudo-hemisphere.*

Proof. Notice that by definition a hypergraph $\mathcal{H}$ is a pseudo-hemisphere hypergraph if and only if some rows and columns of its incidence matrix can be inverted such that it becomes the incidence matrix of an ABA-free hypergraph. Using Theorem 7.21 we get that this is equivalent to the fact that we can permute the rows and columns of the incidence matrix of $\mathcal{H}$ and invert some of the rows and columns to get a $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$-free matrix. This property obviously holds for a matrix if and only if it holds for its transpose and thus, similarly to Proposition 7.29, we can conclude that the dual of $\mathcal{H}$ is also a pseudo-hemisphere hypergraph. □

Furthermore, there is a nice geometric representation of such hypergraphs using *pseudo-hemisphere arrangements*, a generalization of hemisphere arrangements on a sphere.

In a pseudo-hemisphere arrangement the pseudo-hemispheres are regions whose boundaries are centrally symmetric simple curves such that any two intersect exactly twice. (For more on pseudo-hemisphere arrangements, see, e.g., [19].) Without changing the combinatorial properties of the arrangement, we can suppose that the boundary of one of the pseudo-hemispheres is the equator. Using a stereographic projection from the center of the sphere such that this pseudo-hemisphere is mapped to a whole plane, the other pseudo-hemispheres are mapped to pseudo-halfplanes. Thus, we can conclude that $\mathcal{H}$ is a pseudo-hemisphere hypergraph if and only if there is a set of points, S, on the surface of a sphere and a *pseudo-hemisphere arrangement* $\mathcal{F}$ on the sphere such that the incidences among S and $\mathcal{F}$ give $\mathcal{H}$. (Here X corresponds to the points on the southern hemisphere and $S \setminus X$ to the points on the northern hemisphere.)

Another popular geometric representation on the plane, adding *signs* to lines and points, is the following. The vertices correspond to a set of points in the plane together with a direction

[17] According to the definition we may need to duplicate some of the hyperedges so that we have both the original and its complement, but by duplicating some hyperedges the hypergraph remains ABA-free.

(up or down), and the hyperedges correspond to a set of (x-monotone) pseudo-lines with a sign (+ or −). The hyperedge corresponding to a positive pseudo-line is the set of points that point *towards* the pseudo-line, while the hyperedge corresponding to a negative pseudo-line is the set of points that point *away* from the pseudo-line. Positive pseudo-lines correspond to $\mathcal{F}$, negative pseudo-lines to $\bar{\mathcal{F}}$, up points correspond to X and down points correspond to $\bar{X}$. With this interpretation, ABA-free hypergraphs have only + and up signs, pseudo-halfplane hypergraphs have $\pm$ and up signs, dual pseudo-halfplane hypergraphs have + and up/down signs.

In the next table we summarize the best known results about these hypergraphs, with respect to how many points each hyperedge has to contain to have a polychromatic k-coloring and the values of the smallest c for which there exists a c-shallow hitting set for containment-free families.

	Polychromatic k-coloring	Shallow hitting set
ABA-free hypergraphs	$2k-1$ (Theorem 7.14)	2 (Lemma 7.11)
pseudo-halfplane hypergraphs	$2k-1$ (Theorem 7.24)	2 (Lemma 7.27)
Dual pseudo-halfplane hypergraphs	$\leq 3k-2$ (Theorem 7.28)	≤ 3 (Theorem 7.33)
pseudo-hemisphere hypergraphs	$\leq 4k-3$ (Corollary 7.32)	≤ 4 (Theorem 7.33)

We conjecture that even containment-free pseudo-hemisphere arrangements have a 2-shallow hitting set, which would also imply, using Theorem 7.13, that any family whose sets have size at least $2k-1$ admits a polychromatic k-coloring. Towards this conjecture, the only result not in the table is about the special case of dual (ordinary) halfplanes, for which Fulek [53] showed that in the $k=2$ case $2k-1=3$ is the right answer. That is, he showed that we can 2-color any family of halfplanes such that every point of the plane which belongs to at least 3 halfplanes is covered by halfplanes of both colors.

As we can find a polychromatic k-coloring of the points of X and $\bar{X}$ independently with respect to the sets of $\mathcal{F}$ and $\bar{\mathcal{F}}$, respectively, of size at least $2k-1$ using Theorem 7.24, the following is true.

Corollary 7.32. *Given a finite set of points S on the sphere and a pseudo-hemisphere arrangement $\mathcal{H}$, we can color S with k colors such that any pseudo-hemisphere in $\mathcal{H}$ that contains at least $4k-3$ points of S contains all k colors. Equivalently, the vertices S of a pseudo-hemisphere hypergraph can be colored with k colors such that any hyperedge containing at least $4k-3$ vertices contains all k colors.*

To finish, we first prove the following theorem, which, using Theorem 7.13, will imply Theorem 7.28, and also provides another proof for Corollary 7.32.

Theorem 7.33. *Every containment-free dual pseudo-halfplane hypergraph has a 3-shallow hitting set and every containment-free pseudo-hemisphere hypergraph has a 4-shallow hitting set.*

The proof of this result follows again closely the argument of [130]. We note that the next few statements can also be proved using the geometric representation, but here we develop further our completely abstract approach. The reason for this is to demonstrate the power of our method, hoping that in the future it enables attacking completely different problems as well. For an ordered set of vertices $S = Y \cup^* Z$, write $S = (Y, Z)$ if the vertices in Y precede the ones in Z.

Lemma 7.34. *Suppose $\mathcal{F}$ is an ABA-free hypergraph on an ordered vertex set $S = (Y, Z)$. Then $\mathcal{F}' = \mathcal{F}\Delta Y = \{F\Delta Y \mid F \in \mathcal{F}\}$ is an ABA-free hypergraph on the vertices ordered as $S' = (Z, Y)$, i.e., Z precedes Y but otherwise the order inside Y and Z is unchanged.*

Moreover, if $\mathcal{F}$ and $X \subset S$ define a pseudo-hemisphere hypergraph $\mathcal{H}$, i.e., the hyperedges of $\mathcal{H}$ are $\{F\Delta X \mid F \in \mathcal{F}\}$ and $\{\bar{F}\Delta X \mid F \in \mathcal{F}\}$, then $\mathcal{F}'$ and $X' = X\Delta Y \subset S'$ also define the same (if unordered) pseudo-hemisphere hypergraph $\mathcal{H}'$.

Proof. It is enough to show the statement if $|Y| = 1$, as then by induction we can proceed with the vertices of $|Y| > 1$ one by one. Let us denote the original order by $<$ and the new one by $\prec$. It is enough to show that for any $A, B \in \mathcal{F}$ we have no ABA-sequence in $A' = A\Delta Y, B' = B\Delta Y \in \mathcal{F}'$ according to the order $\prec$. We will only use that there is no ABA-sequence in A, B according to $<$. Denote the only element of Y by y. If $y \notin A\Delta B$, then $A\Delta B$ is unchanged by the transformation, thus an ABA-sequence in A', B' according to $\prec$ would also be an ABA-sequence in A, B according to $<$, a contradiction. Thus, without loss of generality, $y \in B \setminus A$ and so $y \in A' \setminus B'$. An ABA-sequence in A', B' according to $\prec$ not containing y would be an ABA-sequence also in A, B according to $<$. Otherwise, if three vertices $a \prec b \prec y$ form an ABA-sequence in A', B', then the three vertices $y < a < b$ form an ABA-sequence in B, A, a contradiction.

For the moreover part, notice that as $F\Delta Y\Delta X' = F\Delta Y\Delta X\Delta Y = F\Delta X$, the hyperedges of $\mathcal{H}$ and $\mathcal{H}'$ are indeed the same. □

Remark 7.35. Lemma 7.34 suggests that instead of our linear ordering of the vertices, we could consider them in circular order. Indeed, let the vertices be points in a circle, where for every vertex the point opposite to it on the circle is also a vertex, called its negated pair. Now take a hypergraph on such a circular point set which contains exactly one point from each opposite pair and is *circular ABAB-free*, that is, it does not contain two sets, A and B, and four points, a, b, c, d, that are in this order around the circle for which $a, c \in A \setminus B$ and $b, d \in B \setminus A$. It is easy to see that such a hypergraph is also circular ABABAB-free, and restricting it to any consecutive subset of half of the vertices is an ABA-free hypergraph with the same (non-circular) order. For example if the original base set is $S = (a, b, c)$ in this order and a set in the family is $F = \{a, c\}$, then in the circular order the base set is $(a, b, c, \bar{a}, \bar{b}, \bar{c})$ and $\{a, c, \bar{b}\}$ is F. After we apply Lemma 7.34 with $Y = \{a\}$, we essentially rotate the non-circular base set by one in the circular order and the "new" base set becomes $S' = \{b, c, \bar{a}\}$. In the circular order F is still $\{a, c, \bar{b}\} = \{\bar{b}, c, \bar{\bar{a}}\}$ which is $\{c\}$ over S' (as only c is non-negated compared to S').

Our earlier results could be translated to this abstraction as well, which models the above rotational symmetry of pseudo-hemispheres in a more natural way. However, further statements we prove are still non-trivial even in this model, so we will stick with our original linear ordering of the vertices.

Lemma 7.36. *[Helly's theorem for pseudo-halfplanes] If any three hyperedges of a pseudo-halfplane hypergraph intersect, then we can add a vertex contained in all pseudo-halfplanes of the hypergraph.*

Proof. We prove the dual statement, as it will be more convenient. That is, suppose that we are given a pseudo-hemisphere hypergraph $\mathcal{H}$, such that all its hyperedges are derived from $\mathcal{F}$, i.e., $\mathcal{H}$ has a representing ABA-free $\mathcal{F}$ and vertex set $X \subset S$ such that for every $H \in \mathcal{H}$ there is an $F \in \mathcal{F}$ such that $H = F\Delta X$. We need to show that if for any three vertices there exists a hyperedge that contains all three of them, then we can add the hyperedge $\bar{X}$ to $\mathcal{F}$ such that it stays ABA-free. This is indeed the dual equivalent of the statement, as $\bar{X}\Delta X \in \mathcal{H}$ contains all the vertices.

For a contradiction, suppose that $\bar{X}$ and some $F \in \mathcal{F}$ violate ABA-freeness because of some vertices x, y, z. By our assumption, there exists another hyperedge $G\Delta X$ which contains all of

x, y, z, thus G and $\bar{X}$ contain the same subset of x, y, z. Thus $F, G \in \mathcal{F}$ contain an ABA-sequence on the vertices x, y, z as $F, \bar{X}$ contain an ABA-sequence on x, y, z, a contradiction. □

Applying this to the complements of the pseudo-halfplanes we get the following.

Corollary 7.37. *Given a pseudo-halfplane hypergraph, either there are already three hyperedges that cover all the vertices, or we can add a vertex which is in none of the hyperedges.*

Now we show that reordering the vertices in an appropriate way keeps the ordered hypergraph ABA-free.

Lemma 7.38. *Suppose $\mathcal{F}$ is an ordered ABA-free hypergraph on vertex set S. Let $F \in \mathcal{F}$ be a smallest hyperedge in the partial ordering of the hyperedges of $\mathcal{F}$. If we reorder S as $(F, \bar{F})$, i.e., the vertices of F go to the front but otherwise the order inside F and $\bar{F}$ is unchanged, then the ordered hypergraph remains ABA-free.*

Proof. Let us denote the original order by $<$ and the new one by $\prec$. Suppose on the contrary, that for some $A, B \in \mathcal{F}$ we have some $a, c \in A \setminus B$ and $b \in B \setminus A$ that satisfy $a \prec b \prec c$. The proof is a simple case analysis of how this could happen. Notice that $c \in F$ implies $b \in F$ and $b \in F$ implies $a \in F$, so there are four cases. If $a, b, c \in F$ or $a, b, c \notin F$, then $a < b < c$. In this case A and B contradict that $\mathcal{F}$ is ABA-free. If $a \in F$ and $b, c \notin F$, then we must have $b < a$. In this case $B < F$, contradicting that F is smallest. If $a, b \in F$ and $c \notin F$, then we must have $c < b$. In this case $A < F$, contradicting that F is smallest. □

Remark 7.39. If $S = \{a < b < c\}$ and $\mathcal{F} = \{\{a\}, \{c\}, \{a, b\}, \{a, c\}, \{b, c\}\}$, then in any reordering of S where the elements of the hyperedge $\{a, c\}$ go to front (i.e., in $\{a < c < b\}$ and $\{c < a < b\}$) ABA-freeness is violated. This shows that in the above Lemma 7.38 the assumption that F is a smallest hyperedge cannot be removed. We might hope that the lemma can be modified to remain true for all hyperedges by first applying Lemma 7.34 for an appropriate prefix set of the points, however this is also not possible. Consider the ABA-free hypergraph $\mathcal{F} = \{\emptyset, \{a\}, \{c\}, \{a, b\}, \{a, c\}, \{b, c\}, \{a, b, c\}\}$ and define $\mathcal{F}_X = \{F \Delta X \mid F \in \mathcal{F}\}$ for any $X \subset S = \{a, b, c\}$. In this case there is no X for which there is a reordering of S that starts with the elements of $\{a, c\} \Delta X$ and for which $\mathcal{F}_X$ is ABA-free with this new order.

Lemma 7.40. *If all the hyperedges of a pseudo-hemisphere hypergraph $\mathcal{H}$ avoid some vertex p in S, then $\hat{\mathcal{H}}$, the dual hypergraph of $\mathcal{H}$, is a pseudo-halfplane hypergraph.*

Proof. Start with a representation of $\mathcal{H}$: an ABA-free hypergraph $\mathcal{F}$ and a point set X such that $\mathcal{H} \subset \{F \Delta X, \bar{F} \Delta X \mid F \in \mathcal{F}\}$. Apply Lemma 7.34 with Y being the vertices before p, this way we get a representation of $\mathcal{H}$ in which p is the first point. Take $\hat{\mathcal{H}}$, the dual of $\mathcal{H}$, with representation $\hat{\mathcal{F}}$ and $\hat{X}$. In $\hat{\mathcal{H}}$, the set corresponding to p is $H_p = F_p \Delta \hat{X}$ for some $F_p \in \hat{\mathcal{F}}$, where we can choose the representation such that F_p is the smallest set of $\hat{\mathcal{F}}$ (because of the ordering used in Proposition 7.18, as p was the smallest point of $\mathcal{F}$). Now apply Lemma 7.38 to get another representation of $\hat{\mathcal{H}}$ in which the points of F_p are at the beginning in the order. As p was a point that was in none of the hyperedges of $\mathcal{H}$, in the dual H_p contains no points and so $F_p = H_p \Delta \hat{X} = \emptyset \Delta \hat{X} = \hat{X}$. Now apply again Lemma 7.34 to $\hat{\mathcal{F}}$ with $Y = \hat{X}$. We get a representation $(\hat{\mathcal{F}}', \hat{X}')$ of $\hat{\mathcal{H}}$ in which $\hat{X}' = \hat{X} \Delta \hat{X} = \emptyset$, that is, $\hat{\mathcal{H}}$ is a pseudo-halfplane hypergraph. □

Applying Lemma 7.40 to the dual of a pseudo-hemisphere hypergraph we get the following dual statement:

Corollary 7.41. *If the empty set is (or can be added as) a hyperedge of a pseudo-hemisphere hypergraph $\mathcal{H}$, then $\mathcal{H}$ is a pseudo-halfplane hypergraph.*

Lemma 7.42. *[Helly's theorem for pseudo-hemispheres] If any four hyperedges of a pseudo-hemisphere hypergraph intersect, then we can add a vertex contained in all pseudo-hemispheres of the hypergraph.*

Proof. Let $\mathcal{H}$ be defined by $\mathcal{F}$ and $X \subset S$. We prove the following stronger statement. If there is a pseudo-hemisphere $F_0 \Delta X = H_0 \in \mathcal{H}$ that has a non-empty intersection with any three other pseudo-hemispheres, then we can add a vertex contained in all the pseudo-hemispheres of the hypergraph. Let $X' = \bar{F_0} = S \setminus F_0$ and denote by $\mathcal{H}'$ the pseudo-hemisphere hypergraph defined on S by $\mathcal{F}$ and X'. As $H'_0 = F_0 \Delta X' = F_0 \Delta (S \setminus F_0) = S$ contains all the points, we can apply Corollary 7.41 to $\mathcal{H}'$ and the complement of H'_0 to conclude that $\mathcal{H}'$ is a pseudo-halfplane hypergraph. It follows from our definitions that the hyperedges in $\mathcal{H}$ and $\mathcal{H}'$ are in bijection such that for every $H \in \mathcal{H}$ there is an $G \in \mathcal{H}'$ (and vice versa) such that $H = G \Delta X' \Delta X$.

Next, we prove that in $\mathcal{H}'$ any three pseudo-hemispheres intersect. Suppose that the original intersection point of these pseudo-hemispheres with H_0 in $\mathcal{H}$ was some $p \in H_0 \cap H_1 \cap H_2 \cap H_3$, where $H_i = F_i \Delta X$ for $1 \leq i \leq 3$. This implies $p \in (F_0 \setminus X) \subset (S \setminus X')$ or $p \in (X \setminus F_0) \subset X'$. In the first case, $p \in F_i$ and $p \in F_i \Delta X' = H'_i$. In the second case, $p \notin F_i$ and $p \in F_i \Delta X' = H'_i$.

Therefore, any three pseudo-halfplanes of $\mathcal{H}'$ intersect. Using Lemma 7.36 for the pseudo-halfplane hypergraph representation of $\mathcal{H}'$, we can add a new point q to all the hyperedges of $\mathcal{H}'$. Denote this new pseudo-halfplane hypergraph by $\mathcal{H}^+$, and let $S^+ = S \cup \{q\}$ and $X^+ = X' \Delta X$ (note that $q \notin X^+$). The hypergraph $\mathcal{H}'^+ = \{G^+ \Delta X^+ \mid G^+ \in \mathcal{H}^+\}$ on the base set S^+ is also a pseudo-hemisphere hypergraph. Moreover, we claim that it is the same as $\{H \cup \{q\} \mid H \in \mathcal{H}\}$, which proves the lemma. Indeed, recall that each hyperedge $H \in \mathcal{H}$ is in bijection with a hyperedge $G \in \mathcal{H}'$ with $H = G \Delta X' \Delta X = G \Delta X^+$. Thus, each hyperedge $H^+ = G \cup \{q\} \in \mathcal{H}^+$ is in bijection with the corresponding $H \cup \{q\} = G \Delta X^+$. This implies that $\mathcal{H}'^+ = \{G \cup \{q\} \Delta X^+ : G \in \mathcal{H}'\}$=$\{H \cup \{q\} : H \in \mathcal{H}\}$. ☐

Applying this to the complements of the pseudo-hemispheres we get the following.

Corollary 7.43. *Given a pseudo-hemisphere hypergraph, either there are four hyperedges that cover all the vertices, or we can add a vertex which is in none of the hyperedges.*

Now we are ready to prove Theorem 7.33.

Proof of Theorem 7.33. First we prove that every containment-free dual pseudo-halfplane hypergraph $\mathcal{H}$ has a 3-shallow hitting set. Consider the dual of $\mathcal{H}$, the pseudo-halfplane hypergraph $\hat{\mathcal{H}}$.

If in $\hat{\mathcal{H}}$ there is a set of at most 3 hyperedges covering every point, then in $\mathcal{H}$ the corresponding 3 vertices form a 3-shallow hitting set. Otherwise, by Corollary 7.37 we could add a point to $\hat{\mathcal{H}}$ that is in none of the pseudo-halfplanes. In this case, by Lemma 7.40 the dual of $\hat{\mathcal{H}}$, which is actually $\mathcal{H}$ itself, is a pseudo-halfplane hypergraph (note that we do not include the empty hyperedge that would be the dual of the newly added point). By Lemma 7.27 it has a 2-shallow hitting set, which is also a 3-shallow hitting set. This finishes the proof of the first statement of Theorem 7.33.

Now we can similarly prove that every containment-free pseudo-hemisphere-hypergraph has a 4-shallow hitting set. Let $\mathcal{H}$ be this hypergraph and take again its dual, $\hat{\mathcal{H}}$. If there is a set of at most 4 hyperedges covering every point in $\hat{\mathcal{H}}$, then in $\mathcal{H}$ the corresponding 4 vertices form a 4-shallow hitting set. Otherwise, by Corollary 7.43 we could again add a point to $\hat{\mathcal{H}}$ that is in none of the pseudo-halfplanes. As before this and Lemma 7.40 imply that $\mathcal{H}$ is a pseudo-halfplane hypergraph and thus by Lemma 7.27 it has a 2-shallow hitting set. □

Proof of Theorem 7.28 and of Corollary 7.32. Being a dual pseudo-halfplane hypergraph and being a pseudo-hemisphere hypergraph are hereditary properties. Thus, Theorem 7.33 implies that all the assumptions of Theorem 7.13 hold with $c = 3$ and $c = 4$, respectively, to get the polychromatic colorings required. □

7.4 Chromatic number of pseudo-halfplane hypergraphs

One may notice that our lower bounds for polychromatic colorings primal and dual pseudo-halfplane hypergraphs look the same. A reason for this could be that the duals of pseudo-halfplane hypergraphs are the same as pseudo-halfplane hypergraphs. However, this is not the case, as shown already by a small example.

Claim 7.44. *The family of pseudo-halfplane hypergraphs and dual pseudo-halfplane hypergraphs is not equal nor does contain one another.*

Proof. It is easy to see that K_4 (the hypergraph on 4 vertices containing all 6 pairs as hyperedges) can be realized as a pseudo-halfplane hypergraph yet one can check (e.g., by a computer program) that it cannot be realized as a dual pseudo-halfplane hypergraph. Thus the dual of K_4 can be realized as a dual pseudo-halfplane hypergraph but not as a pseudo-halfplane hypergraph, proving both containments. □

Before we can determine the chromatic number for most of the hypergraphs studied in the previous sections, we need some preparations. Recall that given a pseudo-halfplane hypergraph $\mathcal{H}$, by definition there exists an ABA-free $\mathcal{F}$ such that $\mathcal{H} \subset \mathcal{F} \cup \bar{\mathcal{F}}$. We called $\mathcal{U} = \mathcal{H} \cap \mathcal{F}$ the upsets and $\mathcal{D} = \mathcal{H} \cap \bar{\mathcal{F}}$ the downsets, both ABA-free families. Further, the unskippable vertices of $\mathcal{U}$ (resp. $\mathcal{D}$) are called top (resp. bottom) vertices. The top and bottom vertices are called the extremal vertices of $\mathcal{H}$. The next statements are easy consequences of the definition of extremal vertices:

Observation 7.45 (Top vertices in an upset)**.** *Given a pseudo-halfplane hypergraph, if X is an upset and $x, y \in X$, then X contains all top vertices that are between x and y. The same holds with bottom vertices if X is a downset.*

Observation 7.46 (Bottom vertex in an upset)**.** *Given a pseudo-halfplane hypergraph, if X is an upset and $x \in X$ is a bottom vertex, then X contains all vertices that are bigger or all vertices that are smaller than x. The same holds if X is a downset and $x \in X$ is a top vertex.*

Let $T = (t_1 = v_1, t_2, \dots t_k = v_n)$ and $B = (b_1 = v_1, b_2, \dots b_\ell = v_n)$ be the sets of top and bottom vertices ordered according to the ordering on P. Call T to be the upper hull and B the lower hull. Note that a vertex may appear in both sets. Let us give the following circular order on C, the set of unskippable vertices: $C = (v_1, t_2, \dots, t_{k-1}, v_n, b_{\ell-1}, \dots, b_2)$

Lemma 7.47. *Every pseudo-halfplane intersects the set of its extremal vertices in an interval of the circular order defined on them.*

Proof. By symmetry it is enough to prove the statement when the pseudo-halfplane H is a topset.

If the pseudo-halfplane hyperedge is empty then the claim trivially holds. Otherwise, by Observation 7.46 H intersects B in an interval that has v_1, v_n or both as an endvertex. By Observation 7.45 H intersects T in an interval. As v_1 and v_n are also endvertices of T, v_1, v_n or both (whichever was in $H \cap B$) must be an endpoint of $H \cap T$. Together the two intervals $H \cap B$ and $H \cap T$ form $H \cap C$ which is thus an interval. □

We are ready to state our results about proper colorings. Note that it also gives us another proof of why K_4 (a 4-chromatic hypergraph) cannot be realized as a dual pseudo-halfplane hypergraph, which was used in Claim 7.44.

Theorem 7.48. *The chromatic number of every ABA-free hypergraph is at most* 3*, the chromatic number of every pseudo-halfplane hypergraph is at most* 4*, the chromatic number of every dual pseudo-halfplane hypergraph is at most* 3 *and these bounds are best possible.*

Proof of Theorem 7.48. The lower bounds are trivial as K_3 can be realized easily as an ABA-free hypergraph which is also a dual pseudo-halfplane hypergraph by definition, also K_4 can be realized easily as a pseudo-halfplane hypergraph as we can realize it already in the plane with 4 points whose convex hull is a triangle and with appropriate 6 halfplanes.

We proceed with the upper bounds. For ABA-free hypergraphs we can alternately color the unskippable vertices with 2 colors and use a third color for the skippable vertices. As every hyperedge intersects the unskippable vertices in an interval, it must be properly colored. In fact Theorem 2.33 states that ABAB-free hypergraphs have chromatic number at most 3, which also implies proper 3-colorability of ABA-free hypergraphs as every ABA-free hypergraph is also ABAB-free.

For pseudo-halfplane hypergraphs the upper bound follows from the more general Theorem 2.12 about the 4-colorability of pseudo-disk wrt. pseudo-disk intersection hypergraphs. However, in our case there is a much simpler proof. Just take the extremal vertices in the order of the vertices, and for each of them if it is a top vertex (resp. bottom vertex) then we give a color different from the previous top vertex (resp. bottom vertex). With 3 colors this can be done even when a vertex is both a top and a bottom vertex. Skippable vertices get the fourth color. Then by Lemma 7.47 every hyperedge with at least two vertices either contains both an unskippable and a skippable vertex or at least two vertices that are consecutive among the top vertices or bottom vertices, in every case the hypergraph is non-monochromatic.

It remains to properly 3-color dual pseudo-halfplane hypergraphs. This is the most complicated part, the proof follows the idea of the respective result of 3-coloring dual halfplane hypergraphs of the author from [73]. Recall that by Proposition 7.29 a hypergraph $\mathcal{H}$ on an ordered set of vertices S is a dual pseudo-halfplane hypergraph (that is, there exists a pseudo-halfplane hypergraph whose dual is $\mathcal{H}$) if and only if there exists a set $X \subset S$ and an ABA-free hypergraph $\mathcal{F}$ on S such that the hyperedges of $\mathcal{H}$ are the hyperedges $F \Delta X$ for every $F \in \mathcal{F}$ (where Δ denotes the symmetric difference of two sets). In the rest we assume this setup. First let $\mathcal{F}_1$ (resp. $\mathcal{F}_2$) be the subhypergraph of $\mathcal{F}$ induced by $S \setminus X$ (resp. X), note that the vertex set of $\mathcal{F}_1$ (resp. $\mathcal{F}_2$ is $S \setminus X$ (resp. X). Every hyperedge $H = F \Delta X$ with at least 2 vertices intersects $S \setminus X$ or X in at least 2 vertices or both of them in exactly one vertex.

We define a graph G on S. The vertex set of G is the set U of the unskippable vertices of $\mathcal{F}_1$ and $\mathcal{F}_2$. First we connect to vertices $v \in S \setminus X$ and $w \in X$ by an edge if there exists a hyperedge $H = \{v, w\}$ in $\mathcal{H}$. We connect to vertices $v, w \in S \setminus X$ if they are consecutive (in the order of unskippable vertices) unskippable vertices of $\mathcal{F}_1$. Finally, we connect to vertices $v, w \in X$ if they are consecutive unskippable vertices of $\mathcal{F}_2$. We claim that this graph is outerplanar and thus 3-colorable. The second and third set of edges form two paths that follow the order of the vertices. It is enough prove that in the first set of edges the connected vertices are in reversed order along the paths, i.e., there are no two edges v_1w_1 and v_2w_2 with $v_1, v_2 \in S \setminus X$ and $w_1, w_2 \in X$ such that $v_1 < v_2$ and $w_1 < w_2$. Showing this it follows that if we reverse one of the paths, the first set of edges forms a caterpillar between the two paths, thus the three parts together obviously form an outerplanar graph. Thus assume on the contrary that there are such two edges. This implies that there is a hyperedge $F_1 = \{v_1\} \cup X \setminus \{w_1\}$ and $F_2 = \{v_2\} \cup X \setminus \{w_2\}$ in $\mathcal{F}$ (these are the hyperedges of $\mathcal{F}$ with $\{v_1, w_1\} = F_1 \Delta X$ and $\{v_2, w_2\} = F_2 \Delta X$). If $w_1 < v_1$ then these two hyperedges form an ABA-sequence on w_1, v_1, w_2, a contradiction. Otherwise, if $v_1 < w_1$ then these two hyperedges form an ABA-sequence on v_1, w_1, w_2, again a contradiction.

Finally, we color the remaining vertices $S \setminus U$. By Lemma 7.7 for each skippable vertex v of $\mathcal{F}_1$ there are 2 unskippable vertices of $\mathcal{F}_1$ such that every hyperedge that contains v contains at least one of these. Thus color v with a color different from the colors of these two unskippable vertices. We color similarly the skippable vertices of $\mathcal{F}_2$.

We claim that the color we get is a proper 3-coloring of $\mathcal{H}$. First, if $H \in \mathcal{H}$ contains a vertex from $S \setminus U$ then it is good by the last step of our coloring process. Otherwise, if H contains at least two unskippable vertices of $\mathcal{F}_1$ or of $\mathcal{F}_2$ then we are done as then it contains two consecutive unskippable vertices in one of them, which get different colors as they are consecutive on one of the two paths we added to G. Finally, if H contains exactly one unskippable vertex of $\mathcal{F}_1$ and one of $\mathcal{F}_2$ then these get different color as they were connected in G in the first set of edges we added to G. □

Determining the maximal chromatic number of pseudo-hemisphere hypergraphs is an interesting open problem we leave open.

7.5 Small epsilon-nets for pseudo-halfplane hypergraphs

Here we briefly mention the consequences of our results to ϵ-nets of hypergraphs defined by pseudo-halfplanes. We omit proofs as they are not hard and can be obtained exactly as the corresponding results in [130].

Let $\mathcal{H} = (V, E)$ be a hypergraph where V is a finite set. Let $\epsilon \in (0, 1]$ be a real number. A subset $N \subseteq V$ is called an ϵ-net if for every hyperedge $S \in E$ such that $|S| \geq \epsilon|V|$, we also have $S \cap N \neq \emptyset$, i.e., N is a hitting set for all "large" hyperedges. It is known that hypergraphs with VC-dimension d have small ϵ-nets (of size $O(d/\epsilon \log(1/\epsilon))$ [63] and in general this is best possible [87]. However, for geometric hypergraphs this is usually not optimal, in particular for halfplanes the following is true. Consider a hypergraph $\mathcal{H} = (P, E)$ where P is a finite set of points in the plane and $E = \{P \cap H \mid H \textit{ is a halfplane}\}$. For this hypergraph there is an ϵ-net of size $2/\epsilon - 1$ for every ϵ [139, 130]. Theorem 7.24 implies that the same bound holds if the hypergraph is defined by pseudo-halfplanes instead of halfplanes. Also, for the dual hypergraph $\bar{\mathcal{H}}$, Theorem 7.28 implies that there exists an ϵ-net of size $3/\epsilon$. Note that our results are in fact stronger as in the appropriate polychromatic coloring each color class intersects all large enough hyperedges, thus we get a partition of the vertices into ϵ-nets (and at least one of them is a small ϵ-net by the pigeonhole principle).

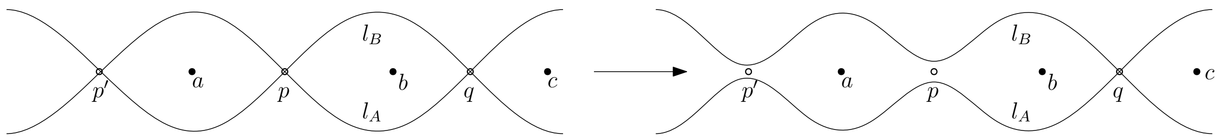

Figure 14: Redrawing a lens to decrease the number of intersections

7.6 Pseudo-halfplanes

First we list some well-known facts about pseudo-line arrangements.

A *pseudo-line arrangement* is a finite collection of simple curves in the plane such that each curve cuts the plane into two components (i.e., both endpoints of each curve are at infinity) and any two of the curves are either disjoint or intersect once, and in the intersection point they *cross*, meaning that any finite perturbation of the curves contains an intersection point. We also suppose that the curves are in general position, i.e., no three curves have a common point. A curve is *graphic* if it is the graph of a function, i.e., an x-monotone infinite curve that intersects every vertical line of the plane. A *graphic* pseudo-line arrangement is such that every curve is graphic. We say that two pseudo-line arrangements are *equivalent* if there is a bijection between their pseudo-lines such that the order in which a pseudo-line intersects the other pseudo-lines remains the same. A *pseudo-halfplane arrangement* is a pseudo-line arrangement, with a side of each pseudo-line selected.

Facts about pseudo-line arrangements [19]

I. (Levi Enlargement Lemma) Given a pseudo-line arrangement, any two points of the plane can be connected by a new pseudo-line (if they are not connected already).

II. Given a pseudo-line arrangement, we can find a pseudo-line arrangement in which every pair of pseudo-lines intersects exactly once, and the order in which a pseudo-line intersects the other pseudo-lines remains the same (ignoring the new intersections).

III. Given a pseudo-line arrangement, we can find an equivalent graphic pseudo-line arrangement.

We also recommend [40] where generalizations of classical theorems are proved for *topological affine planes.*

From these facts it follows that in the definition of a pseudo-halfplane we can (and will) suppose that the underlying pseudo-line arrangement is a graphic pseudo-line arrangement.

Notice that ABA-free hypergraphs are in a natural bijection with (graphic) pseudo-line arrangements and sets of points, such that each hyperedge corresponds to the subset of points *above* a pseudo-line.

Proposition 7.49. *Given in the plane a set of points S (with all different x-coordinates) and a graphic pseudo-line arrangement L, define the hypergraph $\mathcal{H}_{S,L}$ with vertex set S such that for each pseudo-line $\ell \in L$ the set of points above ℓ is a hyperedge of $\mathcal{H}_{S,L}$. Then $\mathcal{H}_{S,L}$ is ABA-free with the order on the vertices defined by the x-coordinates.*

Conversely, given an ABA-free hypergraph $\mathcal{H}$, there exists a set of points S and a graphic pseudo-line arrangement L such that $\mathcal{H} = \mathcal{H}_{S,L}$.

Proof. The first part is almost trivial, suppose that there are two hyperedges A, B in $\mathcal{H}_{S,L}$ having an ABA-sequence on the vertices corresponding to the points $a, b, c \in S$. The pseudo-lines corresponding to the hyperedges A and B are denoted by ℓ_A and ℓ_B. The pseudo-line ℓ_A intersects the vertical line through a below a, the vertical line through b above b and the vertical line through c above c, while ℓ_A intersects these in the opposite way (above/below/above). Thus these lines must intersect in the vertical strip between a and b and also in the strip between b and c, thus having two intersections, a contradiction.

The second part of the proof is also quite natural. Given an ABA-free hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ with an ordering on $\mathcal{V}$, we want to realize it with a planar point set S and a graphic pseudo-line arrangement L. Let S be $|\mathcal{V}|$ points on the x axis corresponding to the vertices in $\mathcal{V}$ such that the order on $\mathcal{V}$ is the same as the order given by the x-coordinates on S. From now on we identify the vertices of $\mathcal{V}$ with the corresponding points of $\mathcal{V}$.

For a given $A \in \mathcal{H}$ it is easy to draw an ℓ_A graphic curve for which the points of S above ℓ_A are exactly in A. Draw a pseudo-line ℓ_A for every $A \in \mathcal{H}$, such that there are finitely many intersections among these pseudo-lines, all of them crossings. What we get is an arrangement of graphic curves, but it can happen that they intersect more than twice. Now among such drawings take one which has the minimal number of intersections, we claim that this is a pseudo-line arrangement.

Assume on the contrary, that there are two curves ℓ_A and ℓ_B intersecting (at least) twice. Let two consecutive (in the x-order) intersection points be p and q, where p has smaller x-coordinate than q. Without loss of generality, ℓ_A is above ℓ_B close to the left of p and close to the right of q, while ℓ_A is below ℓ_B in the open vertical strip between p and q. This structure is usually called a lens, and we want to eliminate it in a standard way, decreasing the number of intersections. We can change the part of ℓ_A and ℓ_B to the left of p (and to the right from the intersection p' next to and left of p if there is any) and change their drawing locally around p (and p' if it exists) such that we get rid of the intersection at p, see Figure 14. If there are no points of S between ℓ_A and ℓ_B and to the left of p (and to the right of p'), then this redrawing does not change the hyperedges defined by ℓ_A and ℓ_B, so we get a representation of $\mathcal{H}$ with less intersections, a contradiction. Thus there is a point $(p' <)a < p$ below ℓ_A and above ℓ_B. Similarly, there must be a point $p < b < q$ above ℓ_A and below ℓ_B and finally a point $q < c$ below ℓ_A and above ℓ_B, otherwise we could redraw the pseudo-lines with less intersections. These three points $a < b < c$ contradict the ABA-freeness of $\mathcal{H}$ as by the definition of the pseudo-lines, $b \in A \setminus B$ and $a, c \in B \setminus A$. □

From these, it follows that the hypergraphs defined by points contained in pseudo-halfplanes are exactly the pseudo-halfplane hypergraphs.

Having spent this much effort to prove results about pseudo-halfplane hypergraphs that are mostly already known for halfplane hypergraphs, it is worthwhile to investigate by what extent is the former a larger family compared to the latter. It is known that there are (simple) pseudo-line arrangements that are non-stretchable already with 9 pseudo-lines (based on the Pappus configuration), that is, which cannot be realized with line arrangements, moreover, almost all of them are such (see, e.g., Chapter 5 of [58] by Felsner and Goodman). This suggests that pseudo-halfplane hypergraphs are a much richer family than halfplane hypergraphs, but it might not be immediately obvious if there is a direct connection as arrangements encode geometric realizations while hypergraphs are strictly combinatorial structures. The aim of this section is to prove that the implication does hold.

One can regard a pseudo-line arrangement as a plane graph: the vertices of an arrangement of pseudo-lines are the intersection points of the pseudo-lines, the edges are the maximal connected parts of the pseudo-lines 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.[18] We say that two pseudo-line arrangements are (combinatorially) *equivalent* if there is a one-to-one adjacency-preserving correspondence between their pseudo-lines, vertices, edges and faces. We need the following:

Theorem 7.50. *Given a simple pseudo-line arrangement $\mathcal{A}$ such that every pair of pseudo-lines intersect, let P be a set of points which has exactly one point in each face of $\mathcal{A}$. Let $\mathcal{H}$ be a pseudo-halfplane hypergraph whose vertex set is P and for each pseudo-line of $\mathcal{A}$ it has a hyperedge which contains the points on one side of this pseudo-line.*[19] *Then in every realization of $\mathcal{H}$ with pseudo-halfplanes the arrangement of the boundary pseudo-lines is equivalent to $\mathcal{A}$.*

Proof. We claim that in any realization of $\mathcal{H}$ by pseudo-halfplanes, the arrangement of the boundary pseudo-lines is equivalent to $\mathcal{A}$. To see this, take the arrangement $\mathcal{A}'$ of the boundary pseudo-lines of any realization. In it we take the pseudo-lines in an arbitrary order: $\ell'_1, \ell'_2, \ldots, \ell'_m$. The pseudo-lines in $\mathcal{A}$ that are the boundaries of the hyperedges in the same order is denoted by $\ell_1, \ell_2, \ldots, \ell_m$. Denote $\mathcal{A}_i$ (resp. $\mathcal{A}'_i$) the arrangement defined by the first i boundary pseudo-lines of $\mathcal{A}$ (resp. $\mathcal{A}'$).

The sub-arrangement $\mathcal{A}'_1$ of $\mathcal{A}'$ defined by ℓ'_1 is unique (has two infinite faces with a bi-infinite curve separating them), and trivially equivalent to $\mathcal{A}_1$. Assume now that for some i we already know that $\mathcal{A}'_{i-1}$ is equivalent to $\mathcal{A}_{i-1}$ (and so we can identify their faces). We shall show that this holds also for i instead of $i-1$.

To see this we add ℓ'_i, the ith pseudo-line to $\mathcal{A}'_{i-1}$. The faces of the arrangement $\mathcal{A}'_{i-1}$ partition the vertex set. Clearly, ℓ'_i must cross those faces of $\mathcal{A}'_{i-1}$ for which the corresponding part of the partition is non-trivially intersected by the respective ith hyperedge (that is, the intersection of that part with the ith hyperedge is neither empty nor equal to that part), call these faces *active*. Note that ℓ_i crosses in $\mathcal{A}_{i-1}$ the active faces and no other face. Further, ℓ'_i must also cross the active faces in $\mathcal{A}'_{i-1}$, although ℓ'_i may intersect further faces (we will see that this cannot happen).

Let F be an active face of the arrangement $\mathcal{A}'_{i-1}$. We claim that topologically it is unique how ℓ'_i can cross F, and it is the same way as ℓ_i crosses F in $\mathcal{A}_i$. Notice that no two edges on the boundary of F can be on the same pseudo-line. Let F' be an active face neighboring F, i.e., it shares a common edge e with F. We claim that ℓ'_i must cross e. Indeed, the (at most) two pseudo-lines which contain the edges that are consecutive with e on the boundary of F (see the bold pseudo-lines on Figure 15) separate the plane into at most 4 parts. F and F' lie in the same part, thus a curve going from F to F' must cross both of these pseudo-lines an even number of times. As ℓ'_i crosses these pseudo-lines at most once, the only possibility is that a part of ℓ'_i connecting F and F' intersects both of these pseudo-lines zero times. The only way for this is if ℓ'_i crosses the edge e, as we claimed. After crossing e, ℓ'_i cannot cross again the pseudo-line supporting e and thus it cannot return to F. The same argument holds for ℓ_i.

We call the ℓ_i-order of the active faces the order in which ℓ_i crosses them. If F is not the first nor the last active face in the ℓ_i-order then our arguments so far show that both ℓ_i and ℓ'_i cross F once and the same way.

[18] The vertices of an arrangement should not be confused with the vertices of a hypergraph.

[19] For each pseudo-line we can choose the side arbitrarily. If for every pseudo-line we choose the side above the pseudo-line then $\mathcal{H}$ is also ABA-free.

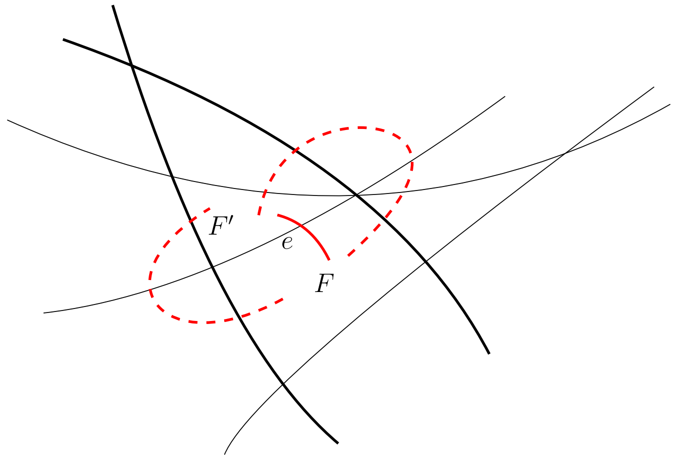

Figure 15: Proof of Theorem 7.50.

About the first and last face we only know that ℓ'_i leaves them once the same way as ℓ_i. We are left to prove that these must be the two faces where ℓ'_i goes into infinity, the same way as ℓ_i. Indeed, having determined the crossings on the boundaries of all the other active faces, we have found a crossing point for every pair of pseudo-lines (as in these faces ℓ'_i and ℓ_i have the same crossings and ℓ_i has no other crossings) and thus it cannot cross any other pseudo-line anymore. That is, the first and last face in the ℓ_i-order is also the first and last one in the order in which ℓ'_i crosses the faces, as claimed.

Thus $\mathcal{A}'_i$ is equivalent to $\mathcal{A}_i$. Repeatedly applying this for $i = 2, 3, \dots m$ we get that the final arrangement $\mathcal{A}'$ must be equivalent to $\mathcal{A}$, finishing the proof. □

Note that in Theorem 7.50 we need that the arrangement of the pseudo-lines is not loose, i.e., every pair of pseudo-lines intersects exactly once.

Now we can take an arbitrary non-stretchable simple pseudo-line arrangement $\mathcal{A}$. By Theorem 7.50 we have an ABA-free hypergraph $\mathcal{H}$ such that in every realization of $\mathcal{H}$ with pseudo-halfplanes the boundary pseudo-lines form an arrangement equivalent to $\mathcal{A}$. Thus, $\mathcal{H}$ cannot be realized with halfplanes as such a realization with halfplanes would also give the arrangement $\mathcal{A}$, contradicting that $\mathcal{A}$ was non-stretchable.

8 ABAB-free hypergraphs

The results of this chapter appeared mostly in [4], a joint work Ackerman and Pálvölgyi, with the exception of the result in Section 8.1, which appeared in [81], a joint work with Pálvölgyi.

In this chapter we deal with results about ABAB-free hypergraphs. In Section 8.2 we prove Theorem 2.32 by showing a family of m-uniform ABAB-free hypergraphs for each m which is not proper 2-colorable and we prove Theorem 2.33, which claims that on the other hand every ABAB-free hypergraph is proper 3-colorable. In Section 8.3 we consider some variants of ABAB-free hypergraphs and the notion of balanced colorings. Then, in Section 8.4, among others we consider ABABA-free hypergraphs and prove that for every $c \geq 2$ there are such hypergraphs which are non-c-colorable. The equivalence between ABAB-free hypergraphs and hypergraphs defined by stabbed pseudo-disks is discussed in Section 8.5, while we omit the proofs for this,

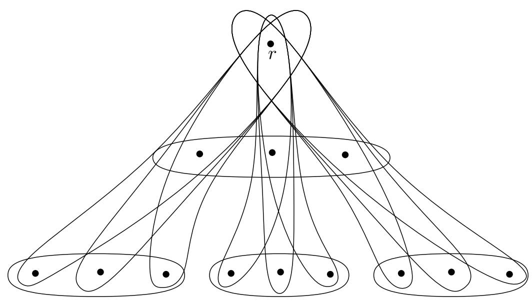

Figure 16: $\mathcal{H}_3$

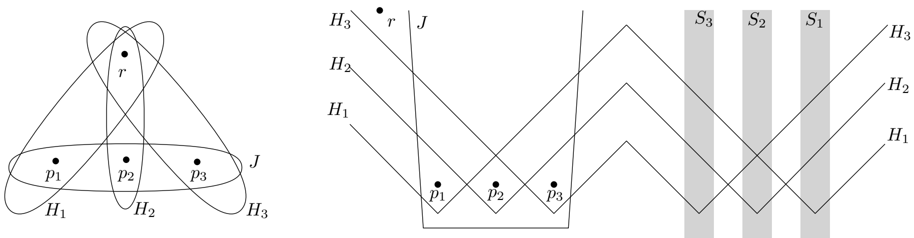

Figure 17: $\mathcal{H}'_2$ and its realization with pseudo-parabolas (for $k = 3$)

we provide an alternative proof of the upper bound in Theorem 2.35 that does not use this equivalence.

8.1 ABAB-free hypergraphs that are not two-colorable

Proof of Theorem 2.32. We show that there are ABAB-free hypergraphs that do not have a proper 2-coloring. We prove this by ordering the vertices of a non-2-colorable hypergraph $\mathcal{H}_k$ in a tricky way to give an ABAB-free hypergraph. First we define this hypergraph $\mathcal{H}_k$ often used in counterexamples, e.g., [105].

Definition 8.1. *Let G_k be the complete k-ary tree of depth k, i.e., the rooted tree such that its root r has k children, each vertex of G_k in distance at most $k-2$ from r has k children and the vertices in distance $k-1$ from r are the leafs (without children).*

$\mathcal{H}_k$ is the k-uniform hypergraph which has two types of hyperedges. First, for every non-leaf vertex the set of its children form an hyperedge. Second, the vertices of every descending path starting in r and ending in a leaf form an hyperedge.

It is easy to see that $\mathcal{H}_k$ is not two-colorable. Now we show how to realize $\mathcal{H}_k$ such that its vertices correspond to points in the plane and its hyperedges correspond to the points *above* pseudo-parabolas (simple curves such that any two intersect at most *twice*). This implies that the x-coordinates define an ordering of the vertices of $\mathcal{H}_k$ showing that $\mathcal{H}_k$ is ABAB-free. We fix k and define $\mathcal{H}'_\ell$ (resp. G'_ℓ) to be the hypergraph (resp. graph) induced by $\mathcal{H}_k$ (resp. G_k) and the subset of the vertices that are in distance at most $\ell-1$ from the root r in G_k ($\mathcal{H}'_\ell$ is a simple hypergraph, i.e., if multiple hyperedges induce the same hyperedge, we take it only once). Thus in particular G'_1 has one vertex and no hyperedges while $\mathcal{H}'_1$ has one vertex and one hyperedge containing it, while $\mathcal{H}'_k = \mathcal{H}_k$ and $G'_k = G_k$. Note that in G'_ℓ every non-leaf vertex has k children, and $\mathcal{H}'_\ell$ has hyperedges of size ℓ corresponding to descending paths (which we usually denote

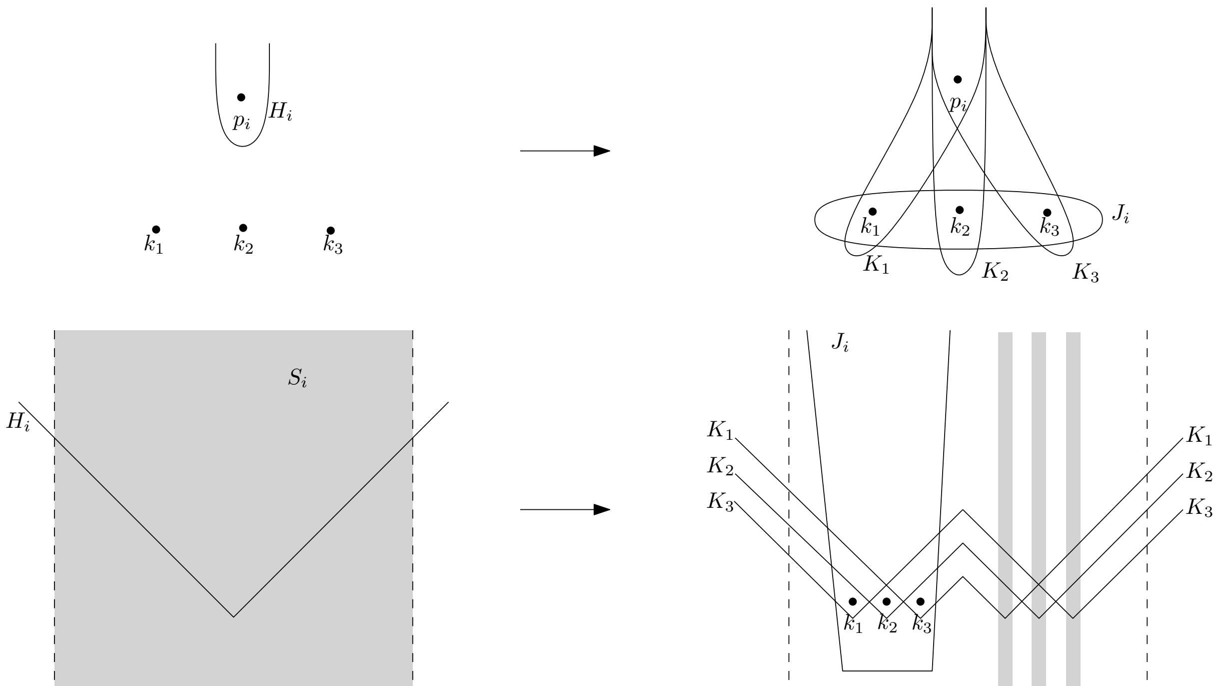

Figure 18: Recursive realization of $\mathcal{H}'_\ell$: adding k children to a leaf

by H_i for some i) and hyperedges of size k corresponding to the set of children of some vertex (which we usually denote by J_i for some i). See Figure 16.

In our realization, to simplify the presentation, points corresponding to vertices will be denoted with the same label, and similarly hyperedges and the corresponding pseudo-parabolas will have the same label.

We will recursively realize $\mathcal{H}'_\ell$, for an illustration see Figure 18. We additionally maintain that each hyperedge (pseudo-parabola) H_i corresponding to a descending path has a vertical strip S_i associated to it, such that inside S_i there are no points and H_i has the lowest boundary (thus no other hyperedge intersects H_i inside S_i). For $\ell = 1$, this is trivial to do as $\mathcal{H}'_1$ has one vertex and one hyperedge containing this vertex. For $\ell = 2$, Figure 17 shows a way to achieve this (for $k = 3$). Now suppose that for some ℓ we have $\mathcal{H}'_\ell$ and we want to construct $\mathcal{H}'_{\ell+1}$. Take the construction of $\mathcal{H}'_\ell$, and for each hyperedge H_i corresponding to a descending path P_i with endvertex p_i, do the following. First make k vertically translated copies of H_i very close to each other. Denote these by $K_1, K_2, \dots K_k$. Next, using these k copies of H_i, realize $\mathcal{H}'_2$ (except the root r) in an appropriately small area inside S_i, by adding k more points $k_1, k_2, \dots k_k$ such that for every i, k_i is above K_i and below every other pseudo-parabola. These points correspond to the children of p_i. Finally, define the pseudo-parabola J_i, which corresponds to the hyperedge containing all the k_i's but no other vertex, as a parabola very close to the vertical strip containing the k_i's. For each i, the vertical strip that belongs to K_i in the inner copy of $\mathcal{H}'_2$ is the strip corresponding to the descending hyperedge that ends at k_i. Therefore all properties are maintained, and by repeating the above procedure for each of the leafs p_i of $\mathcal{H}'_\ell$ we get a realization of $\mathcal{H}'_{\ell+1}$. $\square$

We are not aware of any nice characterization for the dual of ABAB-free hypergraphs, like we had for ABA-free hypergraphs in Proposition 7.18. $\square$

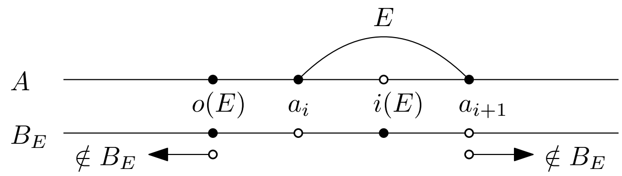

Figure 19: E is an only left-sided pair with witness hyperedge B_E.

8.2 Coloring ABAB-free hypergraphs

In this section we prove Theorem 2.33 which says that every ABAB-free hypergraph is 3-colorable, that is, its vertices can be 3-colored such that every hyperedge that contains at least two vertices, contains at least two different colors. Therefore, for the rest of this section we may suppose that each hyperedge contains at least two vertices.

Let H be an ABAB-free hypergraph on an ordered vertex set. A pair of vertices of H is called *unsplittable* if by adding this pair as a hyperedge of size two to H we get an ABAB-free hypergraph. For a pair of vertices $E = \{p, q\}$ we say that a hyperedge B *splits* this pair if E and B form an $EBEB$- or $BEBE$-sequence.

Lemma 8.2. *Let H be an ABAB-free hypergraph on an ordered vertex set. Then every hyperedge of H contains a pair of vertices that is unsplittable.*

Proof. Let A be a hyperedge of H. If A is of size two, then its vertices form an unsplittable pair, for otherwise there would be a hyperedge B that splits A and this would contradict that H is ABAB-free.

Thus we may assume that A is of size at least 3. For two distinct vertices a and b we write $a < b$ if a precedes b according to the given order of the vertices and say that a is to the *left of* b and b is to the *right of* a. Denote the vertices of A according to their order by $A = \{a_1, a_2, \ldots, a_k\}$. Two such vertices are called *consecutive* if one follows the other in this order. We will prove that one of the consecutive pairs of vertices of A is an unsplittable pair.

Assume on the contrary that none of the consecutive pairs is unsplittable. A consecutive pair $E = \{a_i, a_{i+1}\}$ is *left-sided* (resp., *right-sided*) if there exists a hyperedge $B \in H$ such that they together form a $BEBE$-sequence (resp., $EBEB$-sequence). By our assumption every consecutive pair is either left-sided or right-sided or both. A consecutive pair is called *one-sided* if it is not both left-sided and right-sided. Notice that the leftmost consecutive pair $C = \{a_1, a_2\}$ cannot be left-sided. Indeed, a hyperedge B left-splitting it would also form a $BABA$-sequence, which is a contradiction. Similarly, the rightmost consecutive pair cannot be right-sided and is therefore only left-sided.

For each only left-sided pair $E = \{a_i, a_{i+1}\}$ let B_E be a hyperedge that together with E forms a $B_E E B_E E$-sequence, see Figure 19. The existence of this sequence implies that $a_i, a_{i+1} \in E \setminus B_E$ and that there is a vertex $i(E) \in B_E \setminus A$ among the vertices of H between a_i and a_{i+1} (in the left-to-right order of the vertices of H). The leftmost vertex of B_E is denoted by $o(E)$. As E is left-sided and B_E is a witness for that, it follows that $o(E) < a_i$. Also, $o(E) \in A \cap B_E$ since if $o(E) \notin A$ then $o(E), a_i, i(E), a_{i+1}$ would form a $B_E A B_E A$-sequence, a contradiction. Note that there is no vertex in B_E to the left of $o(E)$ by definition and there is no vertex in B_E to the right of a_{i+1}, for otherwise E would also be a right-sided pair.

Similarly, for each only right-sided pair $E = \{a_i, a_{i+1}\}$ take a witness hyperedge B_E with which it forms an $E B_E E B_E$-sequence. Thus $a_i, a_{i+1} \in A \setminus B_E$ and there is a vertex $i(E) \in B_E \setminus A$

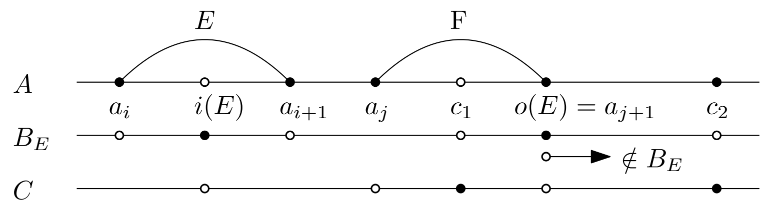

Figure 20: If F is right-sided then $i(E), c_1, o(E), c_2$ form a $B_E C B_E C$-sequence.

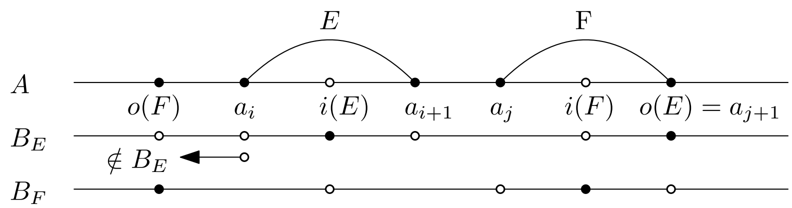

Figure 21: The vertices $o(F), i(E), i(F), o(E)$ form a $B_F B_E B_F B_E$-sequence.

among the vertices of H between a_i and a_{i+1}. In this case denote by $o(E)$ the rightmost vertex of B_E. Therefore, $a_{i+1} < o(E)$ and, as before, we have that $o(E) \in A \cap B_E$.

Among all one-sided pairs of A let $E = \{a_i, a_{i+1}\}$ be the pair with the least number of vertices of A between $i(E)$ and $o(E)$. Without loss of generality we may assume that E is only right-sided.

As $o(E) \in A$ and $a_{i+1} < o(E)$, $o(E) = a_{j+1}$ for some $j > i$. Consider the pair $F = \{a_j, a_{j+1}\}$ (note that a_j may coincide with a_{i+1}). We claim that F cannot be a right-sided pair. Indeed, assume to the contrary that there exists a hyperedge C and two vertices $c_1, c_2 \in C \setminus F$ such that $a_j < c_1 < a_{j+1} < c_2$ (and therefore a_j, c_1, a_{j+1}, c_2 form an $FCFC$-sequence), see Figure 20. Since $o(E) = a_{j+1} < c_2$ and $o(E)$ is the rightmost vertex of B_E, we also have $c_2 \notin B_E$. Also, $i(E) \notin C$, otherwise $i(E), a_j, c_1, a_{j+1}$ would form a $CACA$-sequence, a contradiction. Similarly, $c_1 \notin B_E$, otherwise $a_i, i(E), a_{i+1}, c_1$ would form an $AB_E AB_E$-sequence. However, then $i(E), c_1, a_{j+1}, c_2$ form a $B_E C B_E C$-sequence, which is again a contradiction.

Therefore, F is an only left-sided pair and thus $o(F) \leq a_j$. See Figure 21. Furthermore, $o(F) \leq a_i$ for otherwise there would be less vertices of A between $o(F)$ and $i(F)$ than there are between $o(E)$ and $i(E)$, contradicting our choice of E. We have that $o(F) \in B_F \cap A$ for otherwise $o(F), a_j, i(F), a_{j+1}$ would be a $B_F A B_F A$-sequence. Furthermore, $o(F) \notin B_E$ since $o(F) \leq a_i$ and no vertex of B_E is left of a_i. Similarly, $i(E) \notin B_F$ as otherwise $i(E), a_j, i(F), o(E)$ would form a $B_F A B_F A$-sequence. Finally, $i(F) \notin B_E$ for otherwise $a_i, i(E), a_{i+1}, i(F)$ would form an $AB_E AB_E$-sequence.

Thus, the vertices $o(F), i(E), i(F), o(E)$ form a $B_F B_E B_F B_E$-sequence, leading to the final contradiction. □

Proof of Theorem 2.33. Let H be an ABAB-free hypergraph on an ordered vertex set. We call a hyperedge of size at least 3 *unhit* if it does not contain as a subset a hyperedge of size 2. Starting from H we create a series of hypergraphs as follows. If the current hypergraph contains an unhit hyperedge, then by Lemma 8.2 this hyperedge contains an unsplittable pair which we add as a new hyperedge and obtain the next hypergraph in our series. Since H has a finite number of hyperedges and every hypergraph has one less unhit hyperedge than its preceding hypergraph,

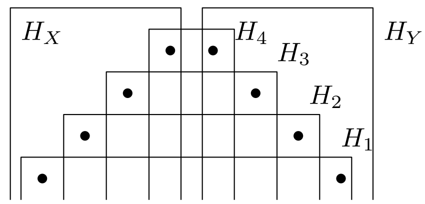

Figure 22: A containment-free bottomless rectangle family without a shallow hitting set

we get a finite series of hypergraphs. Let H' be the last hypergraph in this series.

Let G be the graph that is induced by the hyperedges of H' of size two. Note that every hyperedge of H' contains at least one edge of G. Therefore, a proper coloring of G is a proper coloring of H. The graph G also has the ABAB-free property. Consider the following drawing of G. Its vertices are represented by distinct points on a horizontal line according to their ABAB-free order and its edges are drawn as circular arcs above the line. Since G is ABAB-free its drawing does not contain crossing edges. Furthermore, this drawing of G is outerplanar. Since every outerplanar graph is 3-colorable, this completes the proof. □

As mentioned in the introduction, using Theorem 2.34 this also proves the upper bound of Theorem 2.35.

8.3 Bottomless rectangles and balanced polychromatic colorings

Every hypergraph given by a set of points and a collection of bottomless rectangles is ABAB-free, but not necessarily ABA-free. In fact, it is not hard to see that such hypergraphs would correspond exactly to "aBAb"-free hypergraphs, which can be defined similarly to Definition 2.27 as follows.

Definition 8.3. *A hypergraph whose vertices are real numbers is* aBAb-free *if for any two of its hyperedges, A and B, and vertices $x_1 < x_2 < x_3 < x_4$ it does* not *hold that $x_1 \in A$, $x_2 \in B \setminus A$, $x_3 \in A \setminus B$, $x_4 \in B$.*

It was shown in [15] that any finite set of points can be colored with k colors such that any bottomless rectangle that contains at least $3k-2$ points contains a point of every color. Unfortunately, we were not able to prove this using our methods, because containment-free bottomless rectangle families do not have a shallow hitting set, as shown by the following example.

Example 8.4. *Consider the set of points $X = \{(i,i) \mid i = 1..k\}$ and $Y = \{(k+i, k+1-i) \mid i = 1..k\}$ and the bottomless rectangle family that consists of the following.*

I. A rectangle H_X containing X.

II. A rectangle H_Y containing Y.

III. Rectangles H_i containing (i,i) and $(2k+1-i,i)$ for $i = 1..k$.

Any hitting set for the H_i rectangles contains $k/2$ points from X or Y, thus it is not $(k/2-1)$-shallow for H_X or H_Y (for an illustration for $k = 4$ see Figure 22).

Example 8.4 generalizes easily to other families, such as the translates or homothets of a convex polygon. However, to make the proof method work, we do not need shallow hitting sets for all containment-free families, it is enough to have them for uniform families (that is, in which every bottomless rectangle contains the same number of points). Even in this case, we do not always have a 3-shallow hitting set, as proved by Bursics et al. [22]. However, recently Planken and Ueckerdt [117] managed to prove that 10-shallow hitting sets exist for uniform families. They also show further results about the existence of shallow hitting sets for other uniform families.

Instead of shallow hitting sets, we can ask whether a k-coloring exists for any containment-free bottomless rectangle family that satisfies a certain nice property, that can be achieved by repeatedly finding c-shallow hitting sets and making each of them a separate color class. In the proofs in earlier sections, after k shallow hitting sets were found and colored to different colors, we did not care about the remaining points, they were colored arbitrarily. Instead, we could find a $(k+1)$-st shallow hitting set for the remaining points and use the first color for them, then the second color for the $(k+2)$-nd shallow hitting set, and so on, until there are no more points left. In general in the i-th step the shallow hitting set is colored with color $i \pmod k$, where color 0 and color k denote the same color. This way we achieve a coloring that is not just polychromatic, but also has the following *balanced* property.

Definition 8.5. *We say that a k-coloring is* c-balanced *if for any given set (hyperedge) of our family denoting the sizes of any two color classes in it by n_1 and n_2, then we have $n_1 \le c(n_2+1)$.*

As we have seen above, if a family has a c-shallow hitting set, then it also has a c-balanced k-coloring for any k. For uniform families, a converse also holds; if every set has size n, then any color class of a c-balanced n/c-coloring is a c^2-shallow hitting set.

8.4 ABABA-free hypergraphs

In this section we prove Theorem 2.37.

We have seen that there are ABAB-free hypergraphs that are not 2-colorable. By extending the construction from Section 8.1, we prove Theorem 2.37 by showing that there are non-c-colorable ABABA-free hypergraphs for every $c \ge 2$.

We will use *depth first search* (DFS) to traverse the vertices (nodes) of a directed rooted tree. The order in which a DFS search visits the vertices is called a DFS order. The root of the tree is the first vertex that is visited in this search, and thus the first vertex in the DFS order. In each subsequent step, we take the last (already visited) vertex in the DFS order that has a yet unvisited child, and visit one such child.

Proof of Theorem 2.37. Let $T(a,b)$ denote a full a-ary tree of depth $b-1$. That is, a tree in which every internal vertex has a children and every leaf is at distance $b-1$ from the root of the tree (i.e., the path connecting the root to the leaf contains b vertices).

Let the hypergraph $H(a,b)$ be defined as follows (see also [105]). Its vertex set is the vertex set of $T(a,b)$, the set of children of each internal vertex is a *horizontal* hyperedge of size a and the set of vertices of every path from the root to a leaf is a *vertical* hyperedge of size b.

It is easy to see that in every coloring of the vertices of $H(a,b)$ with two colors there is either a monochromatic horizontal hyperedge (of size a) or a monochromatic vertical hyperedge (of size b). Therefore, $H_2 := H(m,m)$ is an m-uniform non-2-colorable hypergraph. Let n_2 denote the number of vertices of H_2.

For $c > 2$ we define a non-c-colorable m-uniform hypergraph H_c recursively. The vertices of H_c are the vertices of $T(n_{c-1}, m)$. The hyperedges of H_c are defined as follows. For each set of m vertices that lie on a path from the root of the tree to one of its leaves define a *vertical* hyperedge; for each set of n_{c-1} children of an internal vertex of $T(n_{c-1}, m)$ define a hypergraph isomorphic to H_{c-1} and add all of its hyperedges as *horizontal* hyperedges of H_c.

It follows from the definition that H_c is an m-uniform hypergraph for every $c \geq 2$. It remains to show that H_c is non-c-colorable and ABABA-free.

Proposition 8.6. *H_c is non-c-colorable for every $c \geq 2$.*

Proof. We prove by induction on c. We observed above that H_2 is non-2-colorable. Suppose that $c > 2$ and H_{c-1} is non-$(c-1)$-colorable. Assume by contradiction that H_c is c-colorable and consider a proper coloring of its vertices by c colors. Recall that the vertices of H_c are the vertices of $T(n_{c-1}, m)$ and assume, without loss of generality, that the color of the root is red. Then, the color of one of its children must also be red, for otherwise, there is a $(c-1)$-coloring of the copy of H_{c-1} induced by these children. Similarly, one of the children of that red child must also be colored red, etc., yielding a path of red vertices from the root to one of the leaves. However, this implies a monochromatic red hyperedge of H_c which is a contradiction. □

Proposition 8.7. *H_c is ABABA-free.*

Proof. We prove by induction on c. First, we show that H_2 is ABABA-free,[20] the general case will follow similarly. Recall that the vertices of H_2 are the vertices of $T(m, m)$. We claim that any DFS-order of these vertices is ABABA-free. Indeed, let $A = \{a_1, a_2, \ldots, a_m\}$ and $B = \{b_1, b_2, \ldots, b_m\}$ be two hyperedges of H_2, such that the vertices of these hyperedges are listed in DFS order. Now we distinguish some cases.

If both A and B are vertical, then their first few elements are equal, and after these the remaining elements of one precedes the other, i.e., their vertices are ordered as, say, $a_1 = b_1, \ldots, a_i = b_i$, $a_{i+1}, \ldots, a_m$, $b_{i+1}, \ldots, b_m$ for some i. In this case they are even ABA-free.

If both A and B are horizontal, then either $A = B$ or they are disjoint. In the latter case, their vertices are ordered as, say, $a_1, \ldots, a_i$, $b_1, \ldots, b_m$, $a_{i+1}, \ldots, a_m$ for some i. In this case they are ABAB-free.

If A is vertical and B is horizontal, then they intersect in at most one element. Their vertices can be ordered as $a_1, \ldots, a_{i-1}$, $b_1, \ldots, b_j (=) a_i$, $a_{i+1}, \ldots, a_m$, $b_{j+1}, \ldots, b_m$ for some i, j, where b_j and a_i might be equal. In this case they are BABA-free, thus also ABABA-free.

If A is horizontal and B is vertical, then the same argument gives that they are even ABAB-free.

The proof of the induction step is quite similar. Suppose that $c > 2$ and H_{c-1} is ABABA-free. Recall that the vertices of H_c are the vertices of the tree $T(n_{c-1}, m)$. A vertical edge and another (vertical or horizontal) edge behave exactly the same way as for H_2, thus they will be ABABA-free using the same arguments. The same holds for two horizontal edges if they are not the descendants of the same vertex, i.e., they are not from the same copy of H_{c-1}. The only interesting case to check is if A and B are horizontal and they are from the same copy of H_{c-1}.

[20]We note that H_2 is in fact even ABAB-free, as shown in Theorem 2.32, but not when we take the vertices in DFS order. Instead we would need to take what we could call *siblings first order*. Quite surprisingly, however, if we use this "siblings first order" for $c > 2$, we can have an '*ABABA*', i.e., this would only show that H_c is ABABAB-free.

But then they are ABABA-free using induction, if during the DFS search we order the children in the order given by the induction. □

This concludes the proof of Theorem 2.37. □

Corollary 8.8. *For every triple of integers $2\ell \geq 5$, $c \geq 2$ and $m \geq 2$ there exists an $(AB)^\ell$-free m-uniform hypergraph which is not c-colorable.*

An interesting connection to Radon-partitions is the following. Given three points in $\mathbb{R}^1$, they have a unique Radon-partition into two sets, A and B, whose convex hulls intersect; the points must follow each other in the order A, B, A, so this cannot happen for ABA-free families. Given four points in $\mathbb{R}^2$, there are two possible Radon-partitions; the first is when three points of A contain the only point of B inside their convex hull, while the second is when there are two points in each of A and B such that their connecting segments intersect. Note that none of these configurations are possible for points from the symmetric difference of convex pseudo-disks, i.e., if A and B are convex pseudo-disks, then we cannot pick points from $A \setminus B$ and $B \setminus A$ that form a Radon-partition. We wonder whether this has some higher dimensional generalizations, or is just a coincidence.

8.5 Stabbed pseudo-disks

Theorem 2.34, that is, the equivalence between hypergraph families defined by stabbed pseudo-disks and ABAB-free hypergraphs can be proved in several steps [21]. First, we prove that stabbed pseudo-disks and internally stabbed pseudo-disks define the same hypergraphs. Then we prove that internally stabbed pseudo-disks and *pseudo-parabolas* define the same hypergraphs. Recall that a family of x-monotone bi-infinite Jordan curves[22] is a family of pseudo-parabolas if every two curves in the family intersect at most twice. A set of points S and a (finite) family of pseudo-parabolas $\mathcal{C}$ naturally define a hypergraph $H(S, \mathcal{C})$ whose vertex set is S and whose hyperedge set consists of every subset $S' \subset S$ that is exactly the points from S that lie on or above some pseudo-parabola in $\mathcal{C}$. The equivalence between hypergraphs defined by internally stabbed pseudo-disks and hypergraphs defined by pseudo-parabolas was already mentioned in [9] relying on a result of Snoeyink and Hershberger [131] by which a family of stabbed pseudo-disks can be swept by a ray. We can reprove this equivalence using more elementary tools. Finally, we can prove that hypergraphs defined by pseudo-parabolas are exactly the ABAB-free hypergraphs.

The details of these series of equivalences is non-trivial yet mostly technical and therefore omitted from this dissertation, and can be found in [4]. Instead, we restrict ourselves to provide an alternative and direct proof for the upper bound in Theorem 2.35. Before that we need some preparations.

We will need to use the fact that the intersection and union of stabbed pseudo-disks are *simply connected* (see also [9] for such a statement), see Theorem 8.9.

Theorem 8.9. *The union and the intersection of every finite family of stabbed pseudo-disks are both simply connected.*

[21] Two hypergraph families are equivalent if for every hypergraph in one family there is an isomorphic hypergraph in the other family.

[22] Let us call an x-monotone bi-infinite curve an 'x-monotone bi-infinite Jordan curve' if for every $x_1 < x_2$ the points of the curve which have x-coordinate between x_1 and x_2 form a Jordan arc.

As in the literature there seems to be a bit of confusion (and even false claims) as for why Theorem 8.9 holds, we provided a new and simple proof for it (in a more general form). As the proof is rather technical, we omit it from this dissertation, the details can be found in [4].

Proof of the upper bound in Theorem 2.35. Let $\mathcal{F}$ be a family of pseudo-disks whose intersection is non-empty and let S be a finite set of points. We wish to show that it is possible to color the points in S with three colors such that any pseudo-disk in $\mathcal{F}$ that contains at least two points from S contains two points of different colors.

Consider a finite subfamily $\mathcal{F}' \subset \mathcal{F}$ that defines the same hypergraph on S as $\mathcal{F}$. By applying Corollary 4.7 for every pseudo-disk (with $\ell = 2$ and an arbitrary point inside the pseudo-disk) we extend $\mathcal{F}'$ such that every pseudo-disk with at least two points from S contains a pseudo-disk with exactly two points from S. The pairs of points for which there is a pseudo-disk containing exactly these two points form the edges of the Delaunay graph of S with respect to $\mathcal{F}$. It follows that by properly 3-coloring this Delaunay graph one obtains a proper 3-coloring of the hypergraph $H(S, \mathcal{F})$.

We draw each edge of the Delaunay graph G such that it lies in one pseudo-disk that contains its two endpoints. This defines a drawing of G in which edges may intersect, however, by Lemma 4.8 independent edges intersect an even number of times.

Consider the subdivision of the plane into faces that the drawing of G defines. We claim that every point of S is incident to the unbounded face. Indeed, otherwise there is a cycle in G whose drawing separates a point $q \in S$ from infinity. However, then the union of the pseudo-disks corresponding to these edges is not simply connected (as it separates q from infinity), contradicting Theorem 8.9.

This implies that G is an outerplanar graph (note that the embedding of G that we consider is not necessarily a plane embedding). Indeed, connect all the points in S to a new point p' in the unbounded face such that the new edges do not cross each other and the original edges. Denote the resulting graph by G' and note that we get a drawing of G' such that independent edges cross an even number of times. Therefore by the Hanani-Tutte Theorem G' is a planar graph. Consider a plane embedding of G' and delete p' from this embedding. We obtain a plane embedding of G such that one face is incident to all the vertices. Therefore G is outerplanar. Since outerplanar graphs are 3-colorable, this completes the proof. □

9 Tangencies among a single family of curves

The results of this chapter appeared in [83] (Section 9.1, joint work with Pálvölgyi) and in [6] (Section 9.2, joint work with Ackerman).

In Section 9.1 we prove Theorem 2.40 about t-intersecting curves and in Section 9.2 we prove Theorem 2.39 about x-monotone curves that pairwise intersect exactly once.

9.1 A single family of t-intersecting curves

Lemma 9.1. *For every positive integer t there exists a constant $c = c(t)$ such that given a family of at most t-intersecting curves, the tangency graph of the curves cannot contain $K_{t+3,c}$ as a subgraph.*

Proof. Let c be a fixed number, yet explicitly chosen only later. Suppose for a contradiction that there exists a family of curves such that its tangency graph contains $K_{t+3,c}$ as a subgraph. Call the $t+3$ curves that form one part of this bipartite graph *red*, and the c curves that form the other part of the bipartite graph *blue*. Each red curve is cut by the other $t+2$ red curves into at most $1+(t+2)t$ parts, and by the pigeonhole principle there are at least $c_2 = c/(1+(t+2)t)^{t+3}$ blue curves that intersect the same part from each red curve. Choosing these red parts, from now on we assume that we have a family of red and blue curves whose tangency graph contains a K_{t+3,c_2} and the red curves are pairwise disjoint. Moreover, at least $c_3 = c_2/(t+3)!$ blue curves touch these $t+3$ red curves in the same order, where we consider each curve as a continuous image of $[0,1]$.

Denote the red curves by $\gamma_1, \ldots, \gamma_{t+3}$ in the order in which they are touched by the blue curves. At least $c_4 = c_3/2^{t+1}$ blue curves turn in the same (clockwise or counterclockwise, as in Figure 23) direction at each of their $t+1$ tangency points with the red curves $\gamma_2, \ldots, \gamma_{t+2}$, i.e., all red curves that are not the first or last that they touch. We choose c such that $c_4 \geq 2$.

Consider two blue curves. Going from γ_1 to γ_2, and then from γ_2 to γ_3, at γ_2 both blue curves either turn clockwise, or counterclockwise. See Figure 23. Let the parts of the two blue curves that go from the tangency with γ_1 to the tangency with γ_2 be δ_1 and δ_1'. Then either δ_1 and δ_1' intersect, or there is a simply connected region R_1 whose boundary is formed by δ_1, δ_1', and the parts of γ_1 and γ_2 connecting the endpoints of these two subcurves. In the latter case, using that the two blue curves turn in the same way when touching γ_2, one of them continues inside R_1, while the other outside R_1. Without loss of generality, we can assume that γ_3 is outside R_1, as the inside case is analogous, in fact, even equivalent if considered on a sphere. Thus, the blue curve that continues inside R_1 needs to intersect the boundary of R_1 to reach γ_3. As it cannot intersect itself or the two red curves, it has to intersect the other blue curve, i.e., δ_1 or δ_1', whichever is part of the other curve.

In either case there is an intersection of the two blue curves which is either on δ_1, or δ_1', or on both. Thus, we can delete γ_1 and the parts of the two blue curves up to the point of tangency with γ_2, and proceed by induction to get $t+1$ intersections between the two blue curves, contradicting our assumption. □

Proof of Theorem 2.40. The Kővári–Sós–Turán theorem [89] concerning the well-known Zarankiewicz problem states that if a bipartite graph G on $2n$ vertices with parts of size n does not contain $K_{a,b}$ as a subgraph, then it has $O(n^{2-1/a})$ edges. By Lemma 9.1, the tangency graph does not contain $K_{a,b}$ as a subgraph for $a = t+3$ and $b = c$, thus this theorem implies that tangency graph has $O(n^{2-\frac{1}{t+3}})$ edges. □

9.2 A single family of x-monotone = 1-intersecting curves

In this section we prove Theorem 2.39. We do this by considering two types of tangencies according to whether a tangency point is between two curves such that their projections on the x-axis are nested (i.e., one of them is a subset of the other) or non-nested. In each case we consider the tangencies graph whose vertices represent the curves and whose edges represent tangent pairs of curves. In the latter case we show that it is possible to disregard some constant proportion of the edges using the pigeonhole principle and the dual of Dilworth's Theorem and then order the remaining edges such that there is no long monotone increasing path with respect to this order. In the first case, we show that after disregarding some constant proportion of the edges the remaining edges induce a forest.

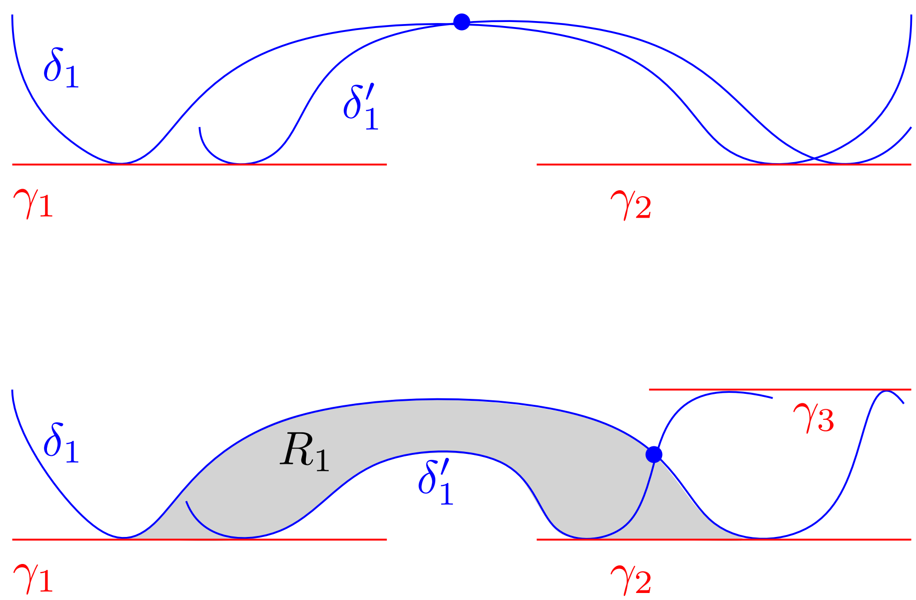

Figure 23: Either δ_1 and δ'_1 intersect each other (top figure) or one of them intersects the boundary of R_1 (bottom figure), where they intersect the other blue curve. Note that it does not matter in which order or from which side they touch the red segments as those could also be replaced by red disks without changing the combinatorial layout of the tangencies.

Let $\mathcal{C}$ be a set of n x-monotone curves such that no three curves in $\mathcal{C}$ intersect at a single point and every pair of curves in $\mathcal{C}$ intersect at exactly one point which is either a crossing or a tangency point. By slightly extending the curves if needed, we may assume that every intersection point of two curves is an interior point of both of them and that all the endpoints of the curves are distinct.

Let $p = (x_1, y_1)$ and $q = (x_2, y_2)$ be two points. We write $p <_x q$ if $x_1 < x_2$ and we write $p <_y q$ if $y_1 < y_2$. We mainly consider the order of points from left to right, so when we use terms like 'before', 'after' and 'between' they should be understood in this sense. For a curve $c \in \mathcal{C}$ we denote by $L(c)$ and $R(c)$ the left and right endpoints of c, respectively. If $p, q \in c$, then $c(p, q)$ denotes the part of c between these two points. We denote by $c(-, p)$ (resp., $c(p, +)$) the part of c between $L(c)$ (resp., $R(c)$) and p. For another curve $c' \in \mathcal{C}$ we denote by $I(c, c')$ the intersection point of c and c'. We may also write, e.g., $c(c', q)$ instead of $c(I(c, c'), q)$

Suppose that an x-monotone curve c_1 lies *above* another x-monotone curve c_2, that is, the two curves are non-crossing (but might be touching) and there is no vertical line ℓ such that $I(c_1, \ell) <_y I(c_2, \ell)$. Assuming the endpoints of c_1 and c_2 are distinct there are four possible cases: (1) $L(c_1) <_x L(c_2) <_x R(c_2) <_x R(c_1)$; (2) $L(c_2) <_x L(c_1) <_x R(c_1) <_x R(c_2)$; (3) $L(c_1) <_x L(c_2) <_x R(c_1) <_x R(c_2)$; and (4) $L(c_2) <_x L(c_1) <_x R(c_2) <_x R(c_1)$. We denote by $c_2 \prec_i c_1$ the relation that corresponds to case i, for $i = 1, 2, 3, 4$. It is not hard to see that each $\prec_i$ is a partial order.

Proposition 9.2. *For every* $i = 1, 2, 3, 4$ *there are no three curves* $c_1, c_2, c_3 \in \mathcal{C}$ *such that* $c_1 \prec_i c_2 \prec_i c_3$.

Proof. It is easy to see that if $c_1 \prec_i c_2 \prec_i c_3$ then c_1 and c_3 do not intersect. See Figure 24 for an illustration of the case $i = 4$. □

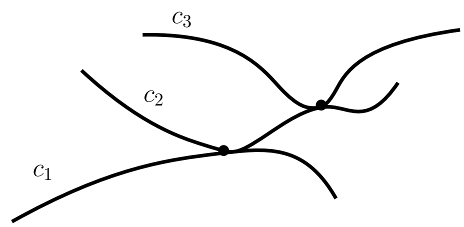

Figure 24: If $c_1 \prec_4 c_2 \prec_4 c_3$ then c_1 and c_3 do not intersect.

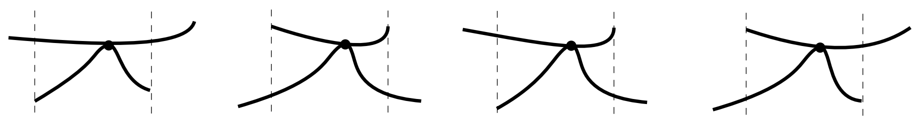

(a) Type 1 (b) Type 2 (c) Type 3 (d) Type 4

Figure 25: Types of tangency points.

We say that the tangency point of two touching curves $c_1, c_2 \in \mathcal{C}$ is of *Type i* if $c_1 \prec_i c_2$ (see Figure 25). We will count separately tangency points of Types 1 and 2 and tangency points of Types 3 and 4.

9.2.1 Bounding touching pairs of Type 1 or 2

We first describe the main idea before going into details: By symmetry it is enough to consider just tangency points of Type 2 which are to the right of a vertical line that intersects all the curves in $\mathcal{C}$. We will prove that there are linearly many such tangencies by showing that the tangencies graph that corresponds to such tangencies is a forest, if one ignores the rightmost touching point on every curve that touches another curve from above.

Lemma 9.3. *There are at most $8n - 4$ tangency points of Type 1 or 2.*

Proof. Since all the curves in $\mathcal{C}$ are pairwise intersecting and x-monotone there is a vertical line ℓ that intersects all of them. By slightly shifting ℓ if needed, we may assume that no two curves intersect ℓ at the same point. We assume without loss of generality that at least half of all the tangency points of Types 1 and 2 are to the right of ℓ, for otherwise we may reflect all the curves about ℓ. We may further assume that at least half of the tangency points of Types 1 and 2 to the right of ℓ are of Type 2, for otherwise we may reflect all the curves about the x-axis. Henceforth, we consider only Type 2 tangency points to the right of ℓ.

By Proposition 9.2 a curve cannot touch one curve from above and another curve from below at Type 2 tangency points. Thus, we may partition the curves into *blue* curves and *red* curves such that at every tangency point a blue curve touches a red curve from below (we ignore curves that contain no tangency points among the ones that we consider).

Proposition 9.4. *Every pair of blue curves cross each other.*

Proof. Suppose that a blue curve b_1 touches another blue curve b_2 from below (the tangency point may be of any type). Since b_2 is a blue curve there is a red curve r which it touches from below. Since the tangency point of b_2 and r is of Type 2 it follows that b_1 and r do not intersect. □

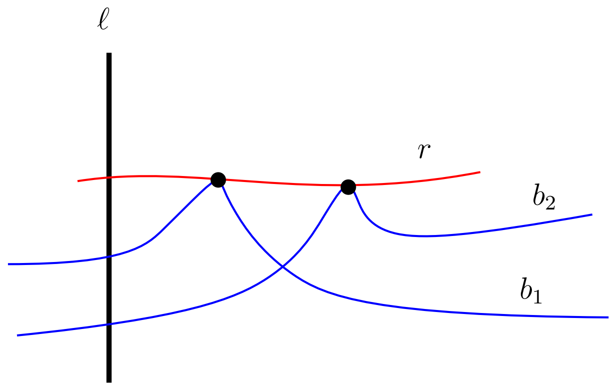

Figure 26: An illustration for the proof of Proposition 9.5: If two blue curves touch the same red curve, then they cross at a point between these two tangency points.

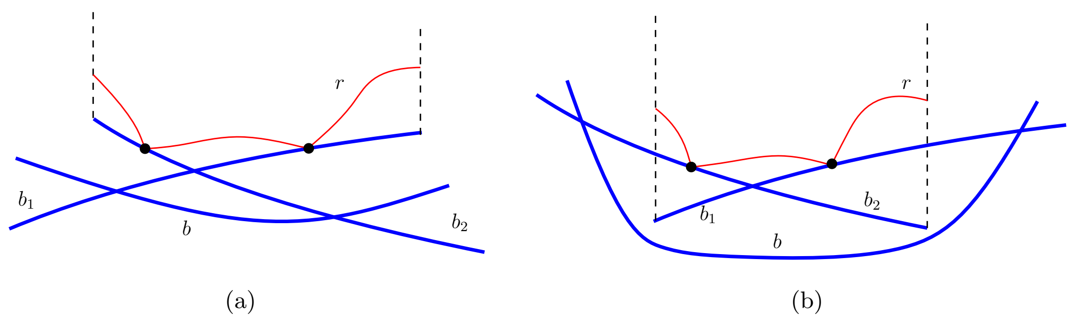

Figure 27: An illustration for the proof of Proposition 9.6. If $I(b_1, b_2)$ is above b, then r and b do not intersect.

Proposition 9.5. *Let r be a red curve and let b_1 and b_2 be two blue curves that touch r, such that $I(r, b_1) <_x I(r, b_2)$. Then $I(r, b_1) <_x I(b_1, b_2) <_x I(r, b_2)$.*

Proof. Since $L(b_i) <_x L(r) <_x R(r) <_x R(b_i)$, for $i = 1, 2$, it is easy to see that the blue curves cross at a point between their tangency points with r, see Figure 26. □

Proposition 9.6. *Let b_1 and b_2 be two blue curves both touching a red curve r. Let b be another blue curve such that $I(b_1, b_2)$ is between $I(b_1, b)$ and $I(b_2, b)$. Then $I(b_1, b_2)$ is below b.*

Proof. Observe that r lies above the upper envelope of b_1 and b_2. Furthermore, since we consider Type 2 tangencies $L(b_i) <_x L(r) <_x R(r) <_x R(b_i)$, for $i = 1, 2$. Therefore if $I(b_1, b_2)$ is above b then r and b do not intersect, see Figure 27. □

We proceed by marking the rightmost Type 2 tangency point on every red curve. Clearly, at most n tangency points are marked. Henceforth, we consider only unmarked tangency points.

Proposition 9.7. *Let b be a blue curve and let r_1 and r_2 be two red curves that touch b, such that $I(b, r_1) <_x I(b, r_2)$ (and both of these tangency points are unmarked). Then $I(b, r_1) <_x I(r_1, r_2) <_x I(b, r_2)$.*

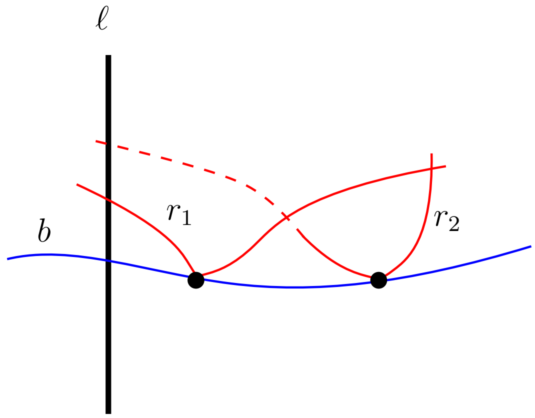

(a) If $I(r_1, r_2) >_x I(b, r_2)$ then r_2 must also intersect $r_1(b, r_2)$.

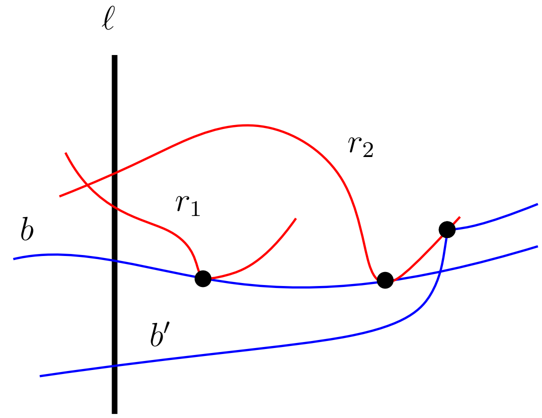

(b) If $I(r_1, r_2) <_x I(b, r_1)$ then b' and r_1 do not intersect.

Figure 28: Illustrations for the proof of Proposition 9.7: if r_1 and r_2 touch b then $I(r_1, r_2)$ is between these two tangency points.

Proof. If $I(r_1, r_2) >_x I(b, r_2)$ then r_2 must intersect $r_1(b, r_2)$ since it intersects ℓ, see Figure 28a. Suppose now that $I(r_1, r_2) <_x I(b, r_1)$. This implies that $R(r_1) <_x I(b, r_2)$ for otherwise r_1 and r_2 intersect also to the right of $I(b, r_1)$. Since $I(b, r_2)$ is not the rightmost tangency point on r_2, there is a blue curve b' that touches r_2 to the right of $I(b, r_2)$ at a Type 2 tangency point. However, it follows from Proposition 9.5 that $I(b, r_2) <_x I(b, b') <_x I(b', r_2)$ which implies that b' lies below b to the left of $I(b, b')$. Therefore, b' and r_1 do not intersect (recall that $L(b) <_x L(r_1)$ and see Figure 28b). □

Let G be the (bipartite) *tangencies graph* of the blue and red curves where edges correspond to unmarked tangency points between a red and a blue curve of Type 2 to the right of ℓ. We will show that G is a forest and hence has at most $n - 1$ edges.

Proposition 9.8. *Let $b_0 - r_0 - b_1 - r_1 - b_2$ be a path in G such that b_0, b_1 and b_2 correspond to distinct blue curves and r_0 and r_1 correspond to distinct red curves. If $I(b_1, \ell) <_y I(b_0, \ell) <_y I(b_2, \ell)$, then b_0 and b_2 intersect to the left of ℓ.*

Proof. It follows from Proposition 9.5 that $I(b_0, b_1)$ and $I(b_1, b_2)$ are to the right of ℓ. Suppose for contradiction that $I(b_0, b_2)$ is to the right of ℓ. Recall that b_2 intersects ℓ above b_0 and b_0 intersects ℓ above b_1. If, going from left to right, b_2 intersects first b_1 (necessarily to the right of $I(b_0, b_1)$) and then it intersects b_0, then $I(b_1, b_2)$ is above b_0 which contradicts Proposition 9.6 (see Figure 29a).

Therefore, b_2 intersects b_0 and then b_1 which implies that b_1 intersects b_2 and then b_0. The curve r_1 lies above the upper envelope of b_1 and b_2, therefore it may touch b_1 at a point which is either in $b_1(b_2, b_0)$ or in $b_1(b_0, +)$. Consider the first case. It follows from Proposition 9.7 that $I(r_0, r_1)$ is between $I(b_1, r_1)$ and $I(b_1, r_0)$ and therefore r_1 must cross b_0 after it touches b_1, since r_0 is above the upper envelope of b_0 and b_1. However, in this case r_1 must cross b_0 once more to the left of $I(b_1, r_1)$, since it also touches b_2 and therefore must lie above it (see Figure 29b).

Consider now the case that r_1 touches b_1 at a point in $b_1(b_0, +)$. Then r_1 must intersect $b_0(b_2, b_1)$ since this is the only part of b_0 which lies above the upper envelope of b_1 and b_2. Since

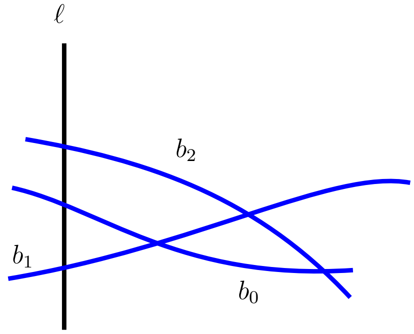

(a) b_2 intersects b_1 and then b_0.

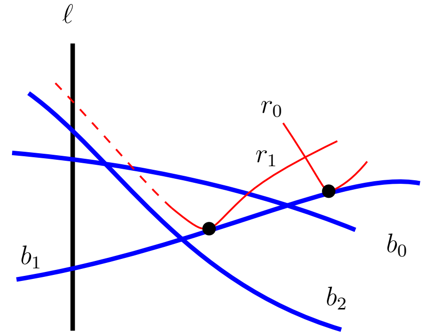

(b) If $I(b_1, r_1) \in b_1(I(b_1, b_2), I(b_1, b_0))$ then r_1 must cross b_0 twice.

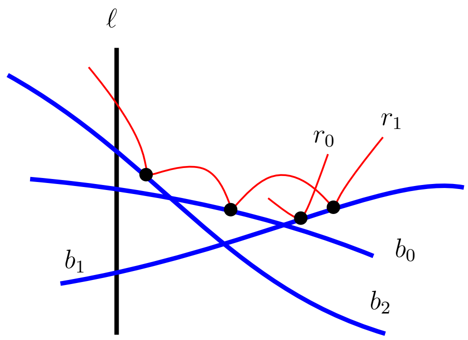

(c) If $I(b_1, r_1) \in b_1(I(b_1, b_0), +)$, then r_1 touches b_0. If $I(b_1, r_0) <_x I(b_1, r_1)$, then r_0 must intersect r_1 twice or crosses b_0 to be able to intersect ℓ.

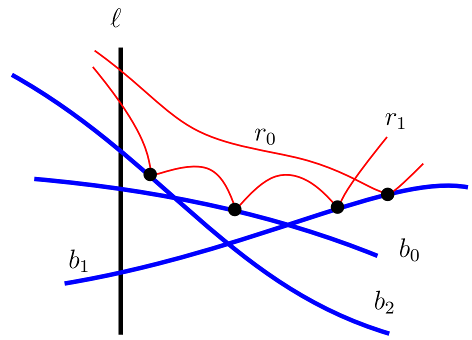

(d) If $I(b_1, r_1) <_x I(b_1, r_0)$, then r_0 does not intersect b_0 to the right of ℓ.

Figure 29: Illustrations for the proof of Proposition 9.8. $I(b_0, b_2)$ is to the right of ℓ.

r_1 may not cross b_2 or intersect b_0 twice, it follows that r_1 must touch b_0 (see Figure 29c). Note that r_0 also touches b_1 at $b_1(b_0, +)$. If $I(b_1, r_0)$ precedes $I(b_1, r_1)$ on b_1, then $r_0(b_1, +)$ must intersect r_1 by Proposition 9.7 (see Figure 29c). However, then r_0 must intersect r_1 once more or cross b_0 which is impossible. If, on the other hand, $I(b_1, r_1)$ precedes $I(b_1, r_0)$ on b_1, then $r_1(b_1, +)$ must intersect r_0 by Proposition 9.7. However, then r_0 cannot touch b_0 to the right of ℓ, see Figure 29d. □

Suppose that G contains a cycle and let $C = b_0 - r_0 - b_1 - r_1 - \ldots - b_k - r_k - b_0$ be a shortest cycle in G, such that b_i corresponds to a blue curve and r_i corresponds to a red curve, for every $i = 0, 1, \ldots, k$. We may assume without loss of generality that b_1 has the lowest intersection point with ℓ among the blue curves in C and that $I(b_0, \ell) <_y I(b_2, \ell)$.

Proposition 9.9. *For every $i \geq 1$ the curve r_i intersects ℓ above r_0 and intersects $b_0(-, \ell)$, $r_0(b_0, +)$ and $b_1(b_0, +)$. See Figure 30 for an illustration.*

Proof. We prove the claim by induction. Consider the case $i = 1$. Before showing that r_1 satisfies the claim, we first look at b_2. It intersects ℓ above b_0 and crosses b_1 to the right of ℓ by

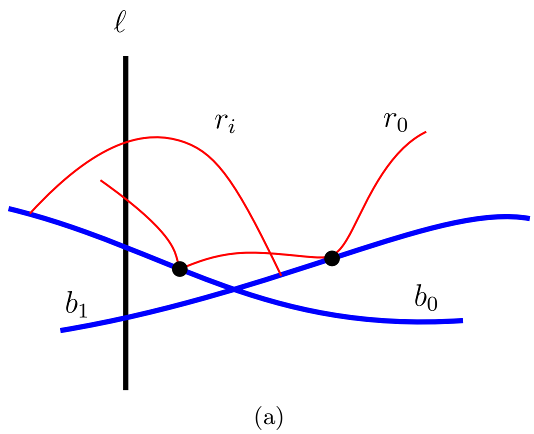

Figure 30: Illustrations for the statement of Proposition 9.9: r_i intersects ℓ above r_0 and intersects $b_0(-,\ell)$, $r_0(b_0,+)$ and $b_1(b_0,+)$.

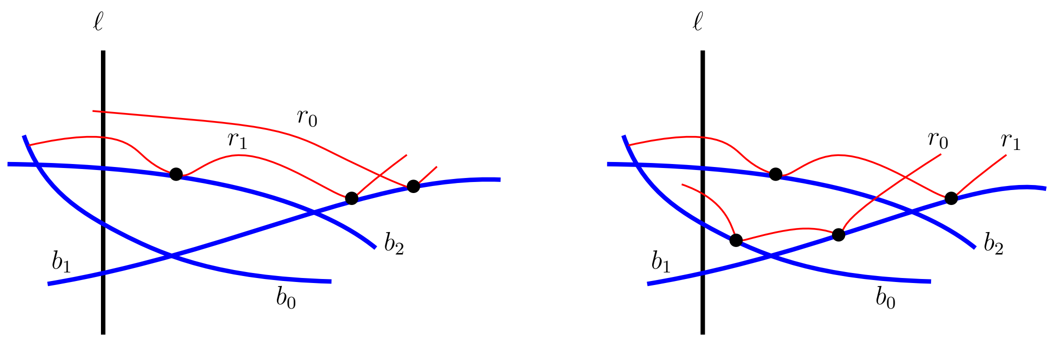

(a) If $I(b_1, r_1) <_x I(b_1, r_0)$ then r_0 cannot touch b_0.

(b) $I(b_1, r_0) <_x I(b_1, r_1)$. r_1 satisfies the properties of Proposition 9.9

Figure 31: Illustrations for the proof of Proposition 9.9: the induction base.

Propositions 9.4 and 9.5. It follows from Proposition 9.8 that b_0 and b_2 intersect to the left of ℓ (refer to Figure 31a). On the other hand the intersection of b_1 and b_2 is between their touchings with r_1, hence to the right of ℓ. Therefore, b_2 intersects $b_1(b_0,+)$. Since r_1 is above b_2 it must also intersect b_0 to the left of ℓ and intersect $b_1(b_0,+)$. It remains to show that r_1 intersects ℓ above r_0 and intersects $r_0(b_0,+)$.

Consider r_0 and note that it must touch $b_0(\ell, b_1)$ and (further to the right) touch $b_1(b_0,+)$. If $I(b_1, r_1) <_x I(b_1, r_0)$ then $I(r_0, r_1)$ is between these points by Proposition 9.7 and it follows that $r_0(\ell, r_1)$ is above $r_1(\ell, r_0)$ and hence r_0 cannot touch b_0 (see Figure 31a). Therefore, $I(b_1, r_0) <_x I(r_0, r_1) <_x I(b_1, r_1)$ which implies that $I(r_0, \ell) <_y I(r_1, \ell)$. Since r_0 touches b_0 before touching b_1, this also implies that r_1 intersects $r_0(b_0,+)$. Therefore, r_1 satisfies the properties above, see Figure 31b (note that r_1 cannot intersect b_1 to the left of $I(b_1, r_0)$ as in Figure 30(a)).

Suppose now that the claim holds for r_i, $i \geq 1$. Observe that b_{i+1} intersects ℓ above b_1 (as all the blue curves in C) and recall that r_i and r_{i+1} touch b_{i+1} at Type 2 tangency points. We consider two cases based on whether b_{i+1} intersects ℓ above or below b_0.

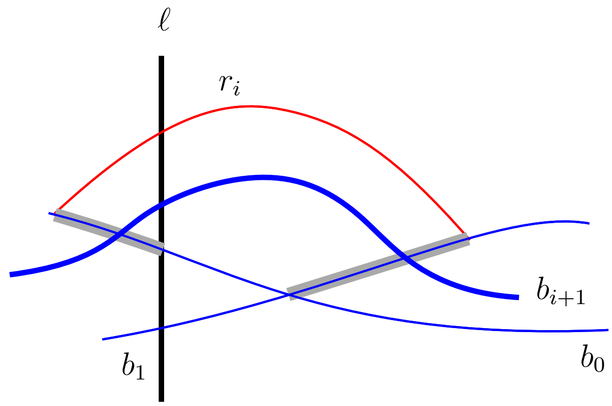

Figure 32: Proposition 9.9 Case 1: $I(b_0, \ell) <_y I(b_{i+1}, \ell) <_y I(r_i, \ell)$. Since $L(b_{i+1}) <_x L(r_i) <_x R(r_i) <_x R(b_{i+1})$ it follows that b_{i+1} must cross $b_0(r_i, \ell)$ and $b_1(b_0, r_i)$.

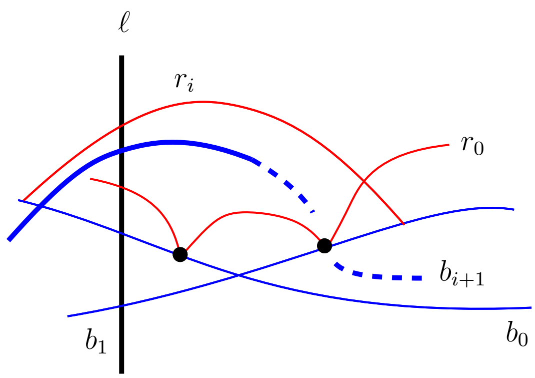

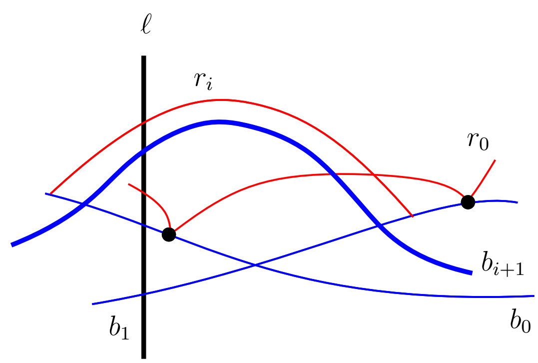

Figure 33: $I(r_0, \ell) <_y I(b_{i+1}, \ell)$. Since b_{i+1} has the desired properties so does r_{i+1}

Case 1: b_{i+1} intersects ℓ above b_0. Since $L(b_{i+1}) <_x L(r_i) <_x R(r_i) <_x R(b_{i+1})$ it follows that b_{i+1} must cross $b_0(r_i, \ell)$ and $b_1(b_0, r_i)$, see Figure 32. Since r_{i+1} lies above b_{i+1} it follows that, as b_{i+1}, it intersects b_0 to the left of ℓ and $b_1(b_0, +)$. It remains to show that $I(r_0, \ell) <_y I(r_{i+1}, \ell)$ and that r_{i+1} intersects $r_0(b_0, +)$. We proceed by considering two subcases.

<u>Case 1a</u>: $I(r_0, \ell) <_y I(b_{i+1}, \ell)$. Since $I(b_{i+1}, \ell) <_y I(r_{i+1}, \ell)$ it follows that $I(r_0, \ell) <_y I(r_{i+1}, \ell)$. Furthermore, b_{i+1} must intersect $r_0(b_0, +)$ since $R(r_i) <_x R(b_{i+1})$ and by the induction hypothesis $r_0(b_0, +)$ intersects r_i. Therefore r_{i+1} intersects $r_0(b_0, +)$ as well (see Figure 33).

<u>Case 1b</u>: $I(b_{i+1}, \ell) <_y I(r_0, \ell)$. Observe that r_0 touches b_0 at $b_0(\ell, b_1)$. We have noted that b_{i+1} intersects $b_0(r_i, \ell)$, hence b_{i+1} must intersect $r_0(\ell, b_0)$, see Figure 34. Therefore $b_{i+1}(r_0, +)$ does not intersect $r_0(b_0, +)$. Since r_i lies above b_{i+1} it follows that r_i does not intersect $r_0(b_0, +)$ which is a contradiction.

Case 2: b_{i+1} intersects ℓ below b_0. Recall that $I(b_1, \ell)$ is the lowest among the intersection points of blue curves with ℓ. Since $I(b_0, b_1) <_x I(b_1, r_i) <_x R(b_{i+1})$ the curve b_{i+1} must cross either $b_0(\ell, b_1)$ or $b_1(\ell, b_0)$. In the latter case b_{i+1} cannot intersect $b_0(b_1, +)$ by Proposition 9.6, therefore it must intersect $b_0(-, \ell)$. But then $L(r_i) <_x L(b_{i+1})$ (see Figure 35a) Therefore, b_{i+1}

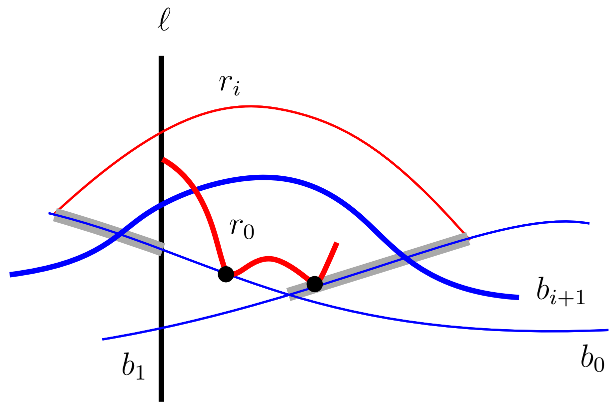

Figure 34: If $I(b_{i+1}, \ell) <_y I(r_0, \ell)$ then r_i does not intersect $r_0(b_0, +)$.

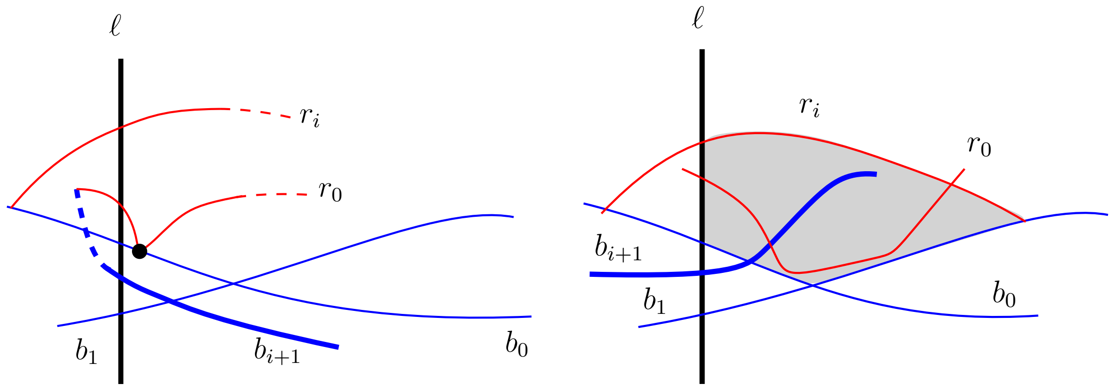

(a) If b_{i+1} intersects $b_1(\ell, b_0)$ then $L(r_i) <_x L(b_{i+1})$.

(b) b_{i+1} intersects $b_0(\ell, b_1)$. If b_{i+1} crosses r_0, then it must cross it twice.

Figure 35: $I(b_1, \ell) <_y I(b_{i+1}, \ell) <_y I(b_0, \ell)$.

intersects $b_0(\ell, b_1)$.

Consider the region bounded by $\ell(r_i, b_0)$, $b_0(\ell, b_1)$, $b_1(b_0, r_i)$ and $r_i(\ell, b_1)$, see Figure 35b. Then b_{i+1} enters this region at $b_0(\ell, b_1)$ and leaves it at $b_1(b_0, r_i)$. Note that b_{i+1} must intersect r_0 at this region since only within this region b_{i+1} has a part above the upper envelope of b_0 and b_1 (where r_0 lies). Furthermore, b_{i+1} must touch r_0, for otherwise it must cross r_0 twice (see Figure 35b).

It follows that b_{i+1} crosses $b_0(r_0, b_1)$, then touches r_0 and then crosses $b_1(b_0, r_i)$. One property that we wish to show is that $I(r_0, \ell) <_y I(r_{i+1}, \ell)$. Suppose that $I(r_{i+1}, \ell) <_y I(r_0, \ell)$. Since r_{i+1} lies above b_{i+1} it may intersect b_1 only at $b_1(b_{i+1}, +)$. It follows that r_{i+1} crosses $r_0(\ell, b_{i+1})$ and intersects $b_1(r_0, +)$, see Figure 36a. However, this implies that $R(r_0) <_x R(r_{i+1}) <_x R(b_{i+1})$. Furthermore, by induction r_i intersects $r_0(b_0, +)$ and so it cannot intersect $r_0(-, b_0)$ and so $L(b_{i+1}) <_x L(r_i) <_x L(r_0)$, thus $I(b_{i+1}, r_0)$ is a Type 2 tangency point. Since $I(b_{i+1}, r_0) <_x I(b_1, r_0)$, it also holds that $I(b_{i+1}, r_0)$ is not the rightmost tangency point on r_0 and therefore (b_{i+1}, r_0) is an edge in G. But then $r_0 - b_1 - r_1 - \ldots - b_{i+1} - r_0$ is a shorter cycle than C.

Thus $I(r_0, \ell) <_y I(r_{i+1}, \ell)$. If r_{i+1} intersects $r_0(-, b_{i+1})$ then it must touch it for otherwise r_{i+1} cannot intersect $b_1(b_{i+1}, +)$ (the only part of b_1 that lies above b_{i+1} and may intersect

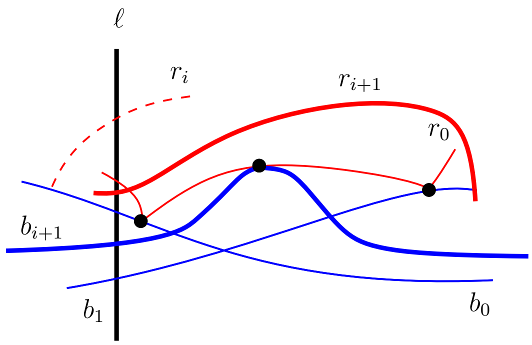

(a) If $I(r_{i+1}, \ell) <_y I(r_0, \ell)$ then (b_{i+1}, r_0) is an edge in G that creates a shorter cycle than C.

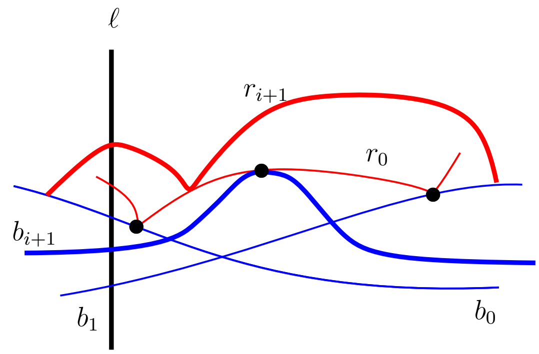

(b) If $I(r_0, \ell) <_y I(r_{i+1}, \ell)$ and r_{i+1} intersects $r_0(-, b_{i+1})$ then (b_{i+1}, r_0) is an edge in G that creates a shorter cycle than C.

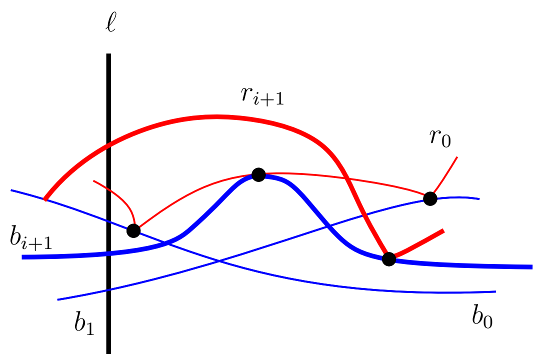

(c) If $I(r_0, \ell) <_y I(r_{i+1}, \ell)$ then r_i has the desired properties.

Figure 36: Concluding cases in the proof of Proposition 9.9.

r_{i+1}), see Figure 36b. However, then $R(r_0) <_x R(r_{i+1}) <_x R(b_{i+1})$ which implies as before that $I(b_{i+1}, r_0)$ is an unmarked Type 2 tangency point and there is a shorter cycle than C. Therefore r_{i+1} intersects $r_0(b_{i+1}, +)$ (and hence, $r_0(b_0, +)$), $b_0(-, \ell)$ and $b_1(b_0, +)$, as desired (see Figure 36c). □

We now return to the proof of Lemma 9.3 and consider the cycle C. It follows from Proposition 9.9 that r_k intersects b_0 to the left of ℓ and therefore (b_0, r_k) cannot be an edge in G. Thus G is a forest and has at most $n - 1$ edges. This implies that there are at most $2n - 1$ Type 2 tangency points to the right of ℓ and at most $8n - 4$ tangency points of Types 1 and 2. □

9.2.2 Bounding touching pairs of Type 3 or 4

We first describe the main idea before going into details: By symmetry it is enough to consider just tangency points of Type 4 which are to the right of a vertical line that intersects all the curves in $\mathcal{C}$. We will prove that there are linearly many such tangencies by showing that after some additional pruning we can order the edges of the tangencies graph such that there is no monotone increasing path of length 7. A linear bound on the size of this graph then follows by the next result about edge-ordered graphs, attributed to Rödl [122] in [55]. For completeness and since the latter reference is not easily accessible we reprove it.

Lemma 9.10. *[122] Let $G = (V, E)$ be an n-vertex graph and let $<$ be a total order of its edges. Let k be an integer such that G has at least $\binom{k+1}{2}n$ edges. Then G contains a monotone increasing path of k edges, that is, a path $e_1 - e_2 - \ldots - e_k$ such that $e_1 < e_2 < \ldots < e_k$.*

Proof. We prove by double induction. For any n and $k = 1$ the claim trivially holds as well as for every k and $n = 1$. For the induction step, we assume that the claim holds for n and $k - 1$ and also for k and $n - 1$ and consider an n-vertex graph with at least $\binom{k+1}{2}n$ totally ordered edges. If the graph contains a vertex of degree smaller than k, then we remove this vertex and obtain a graph with $n - 1$ vertices and at least $\binom{k+1}{2}n - (k - 1) \geq \binom{k+1}{2}(n - 1)$ edges. Thus, by the induction hypothesis a monotone increasing path of length k exists. Otherwise, if the degree of every vertex is at least k, then we remove for every vertex the k highest incident edges (with respect to $<$). There are at least $\binom{k+1}{2}n - nk \geq \binom{k}{2}n$ remaining edges, therefore by the induction hypothesis the graph that we get has a monotone increasing path of $k - 1$ edges and k vertices. Denote this path by p and let v be the last vertex on p. Then at least one of the k removed edges at v has an endpoint which is not on p. By adding this edge to p we get a monotone increasing path of k edges. □

Lemma 9.11. *There are at most* $1152n$ *tangency points of Type 3 or 4.*

Proof. Since all the curves in $\mathcal{C}$ are pairwise intersecting and x-monotone there is a vertical line ℓ that intersects all of them. By slightly shifting ℓ if needed we may assume that no two curves intersect ℓ at the same point. We assume without loss of generality that at least half of all the tangency points of Type 3 or 4 are to the right of ℓ, for otherwise we may reflect all the curves about ℓ. We may further assume that at least half of the tangency points of Type 3 or 4 to the right of ℓ are of Type 4, for otherwise we may reflect all the curves about the x-axis. Henceforth, we consider only Type 4 tangency points to the right of ℓ.

By Proposition 9.2 a curve cannot touch one curve from above and another curve from below at Type 4 tangency points. Thus, we may partition the curves into *blue* curves and *red* curves such that at every tangency point a blue curve touches a red curve from below (we ignore curves that contain no tangency points among the ones that we consider).

Clearly, there are no Type 4 tangencies among the blue curves, however, there might be tangencies of other types among them. Next we wish to obtain a subset of the blue curves such that every pair of them are crossing and they together contain a percentage of the tangency points that we consider. It follows from Proposition 9.2 that the largest chain in the partially ordered set of the blue curves with respect to $\prec_1$ is of length two. Therefore, by Mirsky's Theorem (the dual of Dilworth's Theorem) the blue curves can be partitioned into two antichains with respect to $\prec_1$. The blue curves of one of these antichains contain at least half of the tangency points that we consider. By continuing with this set of blue curves and applying the same argument twice more with respect to $\prec_2$ and $\prec_3$ we obtain a set of pairwise crossing blue curves that together contain at least 1/8 of the tangency points of Type 4 to the right of ℓ. Henceforth we consider these blue curves and the red curves that touch at least one of them at a Type 4 tangency point to the right of ℓ.

Let $G = (B \cup R, E)$ be the (bipartite) *tangencies graph* of these blue and red curves. That is, B corresponds to the blue curves, R corresponds to the red curves and E corresponds to pairs of touching blue and red curves (at Type 4 tangency points to the right of ℓ). We order the edges of G according to the order of their corresponding tangency points from left to right.

The lemma follows from Lemma 9.10 and the next claim.

Proposition 9.12. *G does not contain a monotone increasing path of* 7 *edges starting at B.*

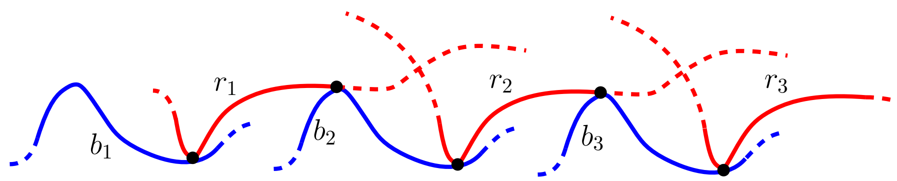

Figure 37: Since the curves intersect ℓ and $R(b_i) <_x R(r_{i-1})$, it follows that $I(b_2, r_1) <_x I(r_1, r_2) < I(b_2, r_2)$ and $I(b_3, r_2) <_x I(r_2, r_3) < I(b_3, r_3)$.

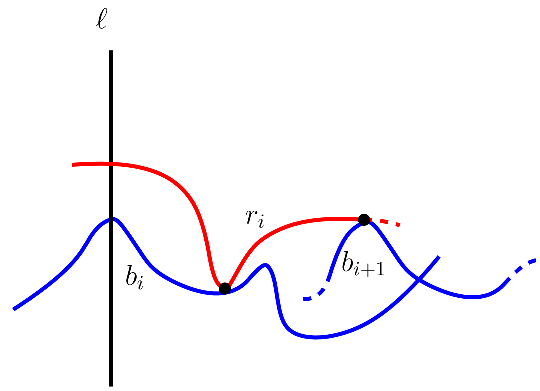

Figure 38: $I(b_i, b_{i+1})$ cannot be to the right of $I(b_{i+1}, r_i)$.

Proof. Suppose that G contains a monotone increasing path $b_1 - r_1 - \ldots - b_4 - r_4$, such that $b_i \in B$ and $r_i \in R$, for $i = 1, 2, 3, 4$. Since all the curves intersect ℓ and $R(b_i) <_x R(r_{i-1})$, we have:

Observation 9.13. *For $i = 1, 2, 3$ we have that $I(b_{i+1}, r_i) <_x I(r_i, r_{i+1}) < I(b_{i+1}, r_{i+1})$ and $r_{i+1}(-, r_i)$ lies above $r_i(-, r_{i+1})$ (see Figure 37).*

Considering consecutive blue curves in the path, observe that $I(b_i, b_{i+1})$ cannot be to the right of $I(b_{i+1}, r_i)$, since in such a case b_{i+1} must intersect b_i or r_i twice to be able to intersect ℓ, see Figure 38.

Observation 9.14. *For $i = 1, 2, 3$ we have that $I(b_i, b_{i+1}) <_x I(b_{i+1}, r_i)$.*

Thus, $I(b_1, b_2)$ is to the left of $I(b_2, r_1)$. We consider two cases based on its location with respect to $I(b_1, r_1)$.

Case 1: $I(b_1, b_2) <_x I(b_1, r_1)$. This implies that $R(b_1) <_x I(b_2, r_1)$. Since $I(b_2, r_1) <_x I(r_1, r_2)$ and $r_2(-, r_1)$ lies above $r_1(-, r_2)$ which lies above b_1, it follows that r_2 and b_1 may intersect only to the left of $L(r_1)$. We must also have $I(b_1, b_2) <_x I(b_1, r_2)$ since $L(b_2) < L(r_2)$ and r_2 lies above b_2 which must be above b_1 to the left of $I(b_1, b_2)$ (see Figure 39a).

Considering b_3, we observe that it cannot intersect $b_2(r_2, +)$ since then it does not intersect r_1. Indeed, suppose that b_3 intersects $b_2(r_2, +)$ and refer to Figure 39a. b_3 lies below r_2 and $R(b_3) <_x R(r_2)$, therefore b_3 may not intersect $r_1(r_2, +)$ which lies above r_2. Since $I(r_1, r_2) <_x I(b_2, r_2) <_x I(b_2, b_3)$ it follows that $b_3(b_2, +)$ cannot intersect r_1. The other part of b_3, $b_3(-, b_2)$, lies below b_2 which lies below r_1 and has its left endpoint to the left of $L(r_1)$. Therefore $b_3(-, b_2)$ cannot intersect r_1 as well.

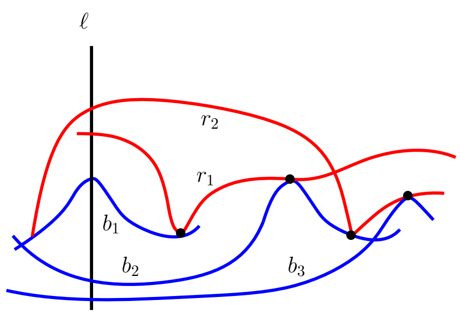

(a) If b_3 intersects $b_2(r_2,+)$ then it cannot intersect r_1.

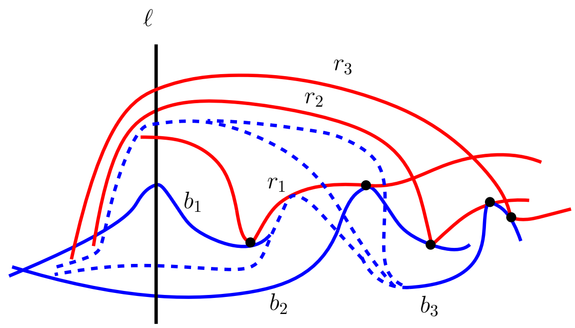

(b) If b_3 intersects $b_2(-,r_2)$ then b_2 and r_3 do not intersect.

Figure 39: Case 1: If $I(b_1,b_2) <_x I(b_1,r_1)$ then $I(b_1,b_2) <_x I(b_1,r_2) <_x L(r_1)$.

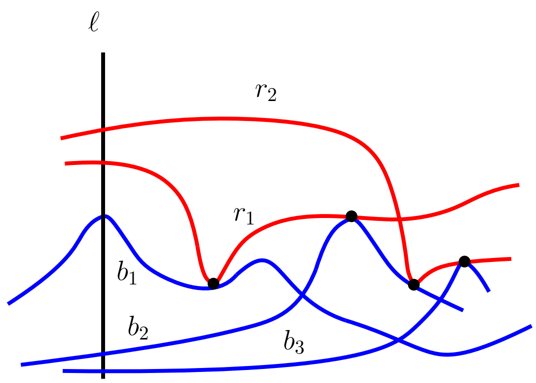

Figure 40: Case 2: $I(b_1,r_1) <_x I(b_1,b_2) <_x I(b_2,r_1)$. If $I(b_2,r_2) <_x I(b_2,b_3)$ then b_3 and r_1 do not intersect.

Therefore, b_3 crosses $b_2(-,r_2)$. This implies that $R(b_2) <_x I(b_3,r_2)$, see Figure 39b. We claim that b_3 and r_2 'block' r_3 from intersecting b_2. Indeed, since $R(b_2) <_x I(b_3,r_2) <_x I(r_2,r_3)$ and $r_3(-,r_2)$ lies above $r_2(-,r_3)$ which lies above b_2 it follows that r_3 may intersect b_2 only to the left of $L(r_2)$. However, $r_2(r_1,+)$ 'blocks' b_3 from intersecting r_1 to the right of $I(r_1,r_2)$, therefore b_3 must intersect $r_1(-,r_2)$, which implies that b_3 must cross b_2 to the right of $L(r_1)$ which is to the right of $L(r_2)$, see Figure 39b. But then b_3 is above b_2 to the left of $L(r_2)$, and since $L(b_3) <_x L(r_2)$ and $L(b_3) <_x L(r_3)$ it follows that b_3 'blocks' b_2 from intersecting r_3 to the left of $L(r_2)$. Therefore, b_2 and r_3 do not intersect.

This concludes Case 1. Note that we have not used the existence of b_4 and r_4, that is, we only considered the path $b_1 - r_1 - b_2 - r_2 - b_3 - r_3$ in G.

Case 2: $I(b_1,r_1) <_x I(b_1,b_2) <_x I(b_2,r_1)$. We claim that $I(b_2,b_3) <_x I(b_2,r_2)$. Indeed, suppose that $I(b_2,r_2) <_x I(b_2,b_3)$ and refer to Figure 40. Then $b_3(-,b_2)$ must lie below $b_2(-,b_3)$. Since b_2 lies below r_1 and $L(b_2) <_x L(r_1)$ it follows that $b_3(-,b_2)$ cannot intersect r_1. Considering the other part of b_3, $b_3(b_2,+)$, it lies below r_2 and its right endpoint is to the left of $R(r_2)$. Since r_2 is below r_1 to the right of $I(r_1,r_2)$ and $I(r_1,r_2) <_x I(b_2,r_2) <_x I(b_2,b_3)$, it follows that $b_3(b_2,+)$ cannot intersect r_1.

Therefore $I(b_2,b_3) <_x I(b_2,r_2)$. However, then we are in Case 1 for the path $b_2 - r_2 - b_3 -$

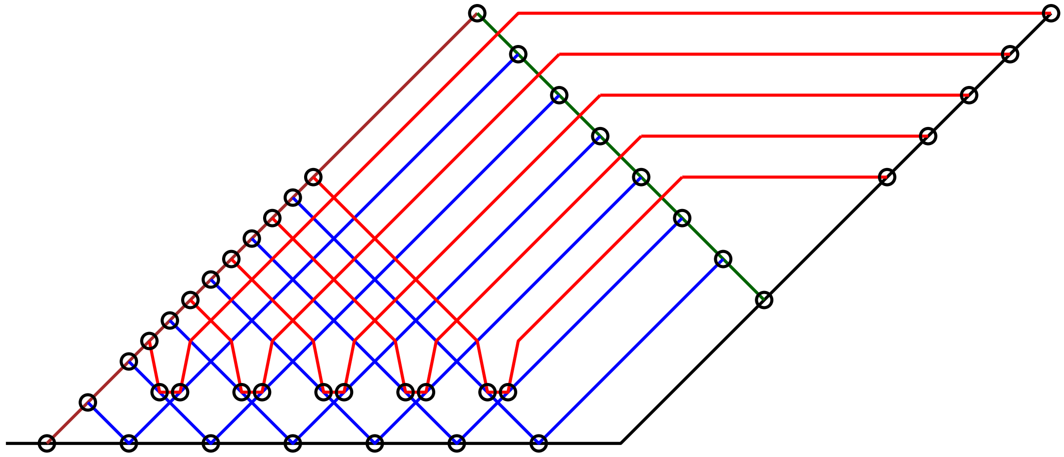

Figure 41: n x-monotone pairwise intersecting 1-intersecting curves can determine $\lfloor 3.5n \rfloor - 8$ tangencies (one can slightly elongate the curves in a way that there is no tangency at an endpoint).

$r_3 - b_4 - r_4$, which is impossible. □

Returning to the proof of Lemma 9.11 by Proposition 9.12 G has no monotone increasing path of length 8 (starting with an edge of any color) and then by Lemma 9.10 G has less than $36n$ edges. This in turn implies that there are at most $8 \cdot 2 \cdot 2 \cdot 36n = 1152n$ tangency points of Types 3 and 4 together (in fact at most $1152n - 32$). □

By Lemmata 9.3 and 9.11 there are at most $1160n - 4$ tangency points among the curves in $\mathcal{C}$. This concludes the proof of Theorem 2.39.

9.3 Discussion

We have shown the best known upper bound on the number of tangencies among t-intersecting curves, yet even already for $t = 1$ there remains a polynomial gap between the lower and upper bounds.

Further, we have shown that n x-monotone pairwise intersecting 1-intersecting curves determine $O(n)$ tangencies. The constant hiding in the big-O notation is rather large, since, for simplicity, we did not make much of an effort to get a smaller constant. In particular, our upper bound can be improved by considering more cases. For example, in the proof of Lemma 9.11 we may consider tangencies among blue curves and avoid using the dual of Dilworth's Theorem. It is also enough to forbid a monotone increasing path of 5 edges in Proposition 9.12, again by considering more cases. It would be interesting to determine the exact maximum number of tangencies among a set of n x-monotone curves each two of which intersect at exactly one point. The best lower bound we have is $\lfloor 3.5n \rfloor - 8$, see Figure 41.

Suppose that we allow more than two curves to intersect at a single point but count the number of *tangency points* rather than the number of tangent pairs of curves. Is it still true that there are linearly many tangency points? For n 1-intersecting curves which are not necessarily x-monotone one can get as many as $\Omega(n^{4/3})$ tangency points via the construction of that many point-line incidences, see Figure 42 for an illustration.

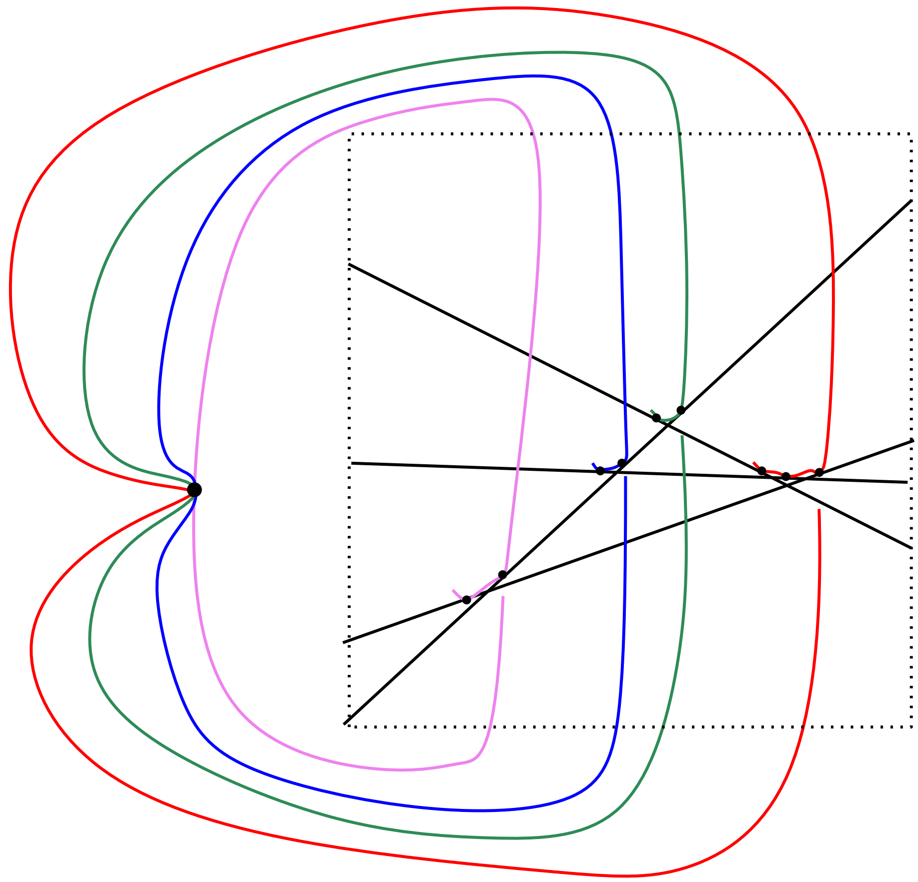

Figure 42: n pairwise intersecting 1-intersecting curves might determine $\Omega(n^{4/3})$ tangency points.

10 Tangencies among 1-intersecting red and blue curves

The results of this chapter appeared in [7], a joint work with Ackerman and Pálvölgyi.

The organization of this chapter is as follows. In Section 10.1 we prove Theorem 2.42. In Section 10.2 we prove a weaker version of Theorem 2.47 which is still stronger than Lemma 10.13 and suffices for proving Theorem 10.14. Theorem 2.47 is proved in Section 10.3. Then in Section 10.4 we give an outline of the conflict-free coloring result from [70] (Theorem 10.12) and explain how by plugging into it Theorem 2.47 or its weaker version we obtain an improved result (Theorem 10.14).

10.1 Proof of Theorem 2.42

In order to avoid technicalities and pathological cases we assume henceforth that all the curves that we consider are non-self-intersecting polygonal chains consisting of finitely many segments and that every pair of curves intersect at finitely many points. Each intersection point involves exactly two curves and is either a proper crossing of these curves or an endpoint of one of them that belongs to the interior of the other curve.[23]

We begin with the following definitions which we will use throughout the chapter. Let $\mathcal{C}$

[23] For L-shapes these properties are trivially satisfied after an appropriate small perturbation.

be a set of curves in the plane. Then $\mathcal{C}$ induces a partition of the plane, which is called the *arrangement* of $\mathcal{C}$, $\mathcal{A}_\mathcal{C}$, and consists of *vertices*, *edges*[24] and *faces*. A *vertex* is either an endpoint of a curve or an intersection point of (two) curves; an *edge* is a maximal sub-curve that does not contain any vertices in its interior; and a *face* is the closure of a maximal connected region of $\mathbb{R}^2 \setminus \mathcal{C}$. The vertices and the edges of $\mathcal{A}_\mathcal{C}$ naturally induce a plane graph $G_\mathcal{C}$. The size of a face F, denoted by $|F|$, is the number of edges that are adjacent to F, where the cut-edges (bridges) in $G_\mathcal{C}$ are counted with multiplicity two. We denote by E_F the set of edges that bound F. Note that if E_F contains a cut-edge, then $|E_F| < |F|$.

We will use the following simple lemma, which essentially follows from 'contracting' some curves into points to get a plane graph.

Lemma 10.1 ([70, Lemma 2.6])**.** *Let $\mathcal{S}$ and $\mathcal{C}$ be two sets of curves such that the curves within each set are pairwise disjoint and every curve in $\mathcal{C}$ intersects exactly two curves from $\mathcal{S}$. Then $\mathcal{D}(\mathcal{S},\mathcal{C})$ is a planar graph.*

Next we prove Theorem 2.46 which is equivalent to Theorem 2.42. Let $\mathcal{S}$ be a set of 1-intersecting red and blue curves, such that no two curves of the same color intersect, and let $\mathcal{C}$ be another set of curves such that each curve in $\mathcal{C}$ intersects a pair of disjoint curves from $\mathcal{S}$ and no other curve from $\mathcal{S} \cup \mathcal{C}$. We will show that $|\mathcal{C}| = O(n)$. Observe first that it follows from Lemma 10.1 that there are at most $3n - 3$ curves in $\mathcal{C}$ such that each of them intersects two curves from $\mathcal{S}$ of the same color.[25]

It remains to bound the number of curves in $\mathcal{C}$ that intersect curves of different colors. Next we consider only this set of curves $\mathcal{C}'$. Observe that each curve $c \in \mathcal{C}'$ contains a sub-curve c' whose endpoints are on edges of different colors that bound the same face of $\mathcal{A}_\mathcal{S}$ and whose interior is contained in that face. We replace every curve in $\mathcal{C}'$ with one of its sub-curves with these properties and, by a slight abuse of notation, keep using $\mathcal{C}'$ to denote the new set of (sub-)curves.

By slightly extending each curve in $\mathcal{S}$, every intersection point of two curves becomes a proper crossing of them. Denote by x the number of intersection points of two curves from $\mathcal{S}$. Thus $\mathcal{A}_\mathcal{S}$ has $2n + x$ vertices and $n + 2x$ edges. Let $\mathcal{F}$ be the face set of $\mathcal{A}_\mathcal{S}$ and let e denote the number of edges of $\mathcal{A}_\mathcal{S}$. Then by Euler's formula we have $e = n + 2x \leq (2n + x) + |\mathcal{F}| - 2$, and therefore, $x \leq n + |\mathcal{F}| - 2$.

For each crossing point we now add 4 dummy curves that connect the four pairs of neighboring edges of $\mathcal{A}_\mathcal{S}$ around this crossing point in its small neighborhood. These $4x$ curves are drawn such that they do not intersect the other curves. For a face $F \in \mathcal{F}$ let $\mathcal{C}'_F$ be the curves in $\mathcal{C}'$ whose interiors are inside F. Then it follows from Lemma 10.1 that $\mathcal{D}(E_F, \mathcal{C}'_F)$ is a bipartite planar graph and therefore the number of its edges is at most $2|E_F| - 4$ (since a red curve and a blue curve may intersect at most once, $\mathcal{A}_\mathcal{S}$ does not contain size-two faces). Since $\sum_{F\in\mathcal{F}} |E_F| \leq 2e$ we have: $|\mathcal{C}'| + 4x \leq \sum_{F\in\mathcal{F}}(2|E_F| - 4) \leq 4e - 4|\mathcal{F}| = 4n + 8x - 4|\mathcal{F}|$. Hence, $|\mathcal{C}'| \leq 4n - 4|\mathcal{F}| + 4x \leq 4n - 4|\mathcal{F}| + 4n + 4|\mathcal{F}| - 8 = 8n - 8$, and the number of edges in $\mathcal{D}(\mathcal{S},\mathcal{C})$ is at most $11n - 11$. This finishes the proof.

Remarks. (1) we have only used that curves from $\mathcal{C}$ do not connect *neighboring* edges of the arrangement of red and blue curves, that is, edges that are consecutive along some face of the

[24]Not to be confused with vertices and edges of a graph.

[25]We sometimes use the weaker bound $3n - 3$ (resp., $2n - 2$) for the maximum size of an n-vertex (resp., bipartite) planar graph, since it holds for every n.

arrangement—this is weaker than requiring that each connected pair of red and blue curves involves disjoint curves.

However, if $\mathcal{S}$ is the union of three sets of pairwise disjoint curves instead of two, then we can no longer claim that $\mathcal{D}(\mathcal{S},\mathcal{C})$ has linearly many edges when the curves in $\mathcal{C}$ may connect only non-neighboring edges (of possibly intersecting curves). Indeed, consider a triangular grid formed by n segments ($n/3$ segments of each of three directions). By slightly shifting the segments of one orientation parallel to themselves, each of the $\Omega(n^2)$ intersection points is replaced by a triangle which is adjacent to at least one hexagon. For each such hexagon it is possible to connect two non-neighboring, non-parallel edges by a curve that intersects no other edges. Thus the number of these curves is $\Omega(n^2)$.

(2) The only place where we used that each pair of a red curve and a blue curve intersects at most once is for claiming that $\mathcal{A}_{\mathcal{S}}$ does not contain size-two faces. Therefore Theorem 2.46 remains true even when each such pair intersects finitely many times, as long as there are no size-two faces in $\mathcal{A}_{\mathcal{S}}$.

10.2 Improving Lemma 10.13

In this section we improve the bound in Lemma 10.13 by proving the following weak version of Theorem 2.47. This bound suffices for obtaining the upper bound in Theorem 10.14.

Lemma 10.2. *Let $\mathcal{S}$ be a set of n axis-parallel line-segments and let $\mathcal{C}$ be a set of pairwise disjoint curves grounded with respect to $\mathcal{S}$. Then $\mathcal{D}(\mathcal{S},\mathcal{C})$ has at most $13n-11$ edges.*

By a slight perturbation of the segments if needed, we may assume that parallel segments do not fall on the same line—this does not decrease the number of edges of $\mathcal{D}(\mathcal{S},\mathcal{C})$, so from now on we assume this is the case.

Recall that E_F denotes the set of edges that bound a face F in the arrangement $\mathcal{A}_{\mathcal{S}}$, where cut-edges are counted once. Next we would like to bound $|E_F|$ where $\mathcal{S}$ is a set of axis-parallel segments. The following fact is probably known, however, we provide a proof for completeness and since we could not find any reference to it.

Lemma 10.3. *Let $\mathcal{S}$ be a set of n axis-parallel line-segments, $n > 1$, and let F be a face of $\mathcal{A}_{\mathcal{S}}$. Then $|E_F| \le 4n-4$. This bound is tight.*

Proof. By slightly extending every segment we may assume that if two segments intersect, then they intersect in their interiors. This step cannot decrease $|E_F|$.

The boundary of F consists of at most one outer closed walk W_0 (which surrounds F) and possibly also some inner closed walks $W_1, W_2, \ldots, W_k$ (which are surrounded by F). Denote by x_i the number of appearances of intersection points of two segments that one encounters while going along W_i, and by y_i the number of segment-endpoints along W_i. Note that an intersection point might be counted with multiplicity, whereas every endpoint is counted exactly once. Let e_i denote the number of edges along W_i, such that each edge is counted with multiplicity one, so $e_i \le |W_i|$. By definition $|E_F| = \sum_{i=0}^{k} e_i$.

Clearly, if $|W_i| = 2$, i.e., W_i consists of a single edge, then $e_i = 1 \le 4 = 2y_i$. Otherwise, $e_i \le x_i$. Indeed, associate every edge along W_i with its preceding vertex (thus, a cut-edge is associated with two vertices). If the preceding vertex of an edge e is a segment-endpoint, then the vertex before that endpoint is an intersection point of two segments which is also associated with e. Therefore, every edge is associated with an intersection point and every intersection point is associated to as many edges as its number of appearances along the walk to which it belongs.

Suppose that we traverse an inner walk W_i such that F is to our left. Then at each endpoint we turn 180° in the positive direction and at each intersection point we turn 90° in the negative direction. As the sum of positive and negative turns should be 360°, we have that $x_i = 2y_i - 4$, thus $e_i \leq 2y_i - 4$. If W_0 exists, then a similar calculation gives $x_0 = 2y_0 + 4$, thus $e_0 \leq 2y_0 + 4$.

Let $y = \sum_{i=0}^{k} y_i$ be the number of segment-endpoints on the boundary of F. If there is no outer walk, then either there are no intersection points, in which case $|E_F| \leq n \leq 4n - 4$, since $n > 1$. Otherwise, for at least one index j we have $|W_j| > 2$. In this case $|E_F| = \sum_{i=0}^{k} e_i \leq 2y_j - 4 + \sum_{i \neq j} 2y_i \leq 2y - 4$. As $y \leq 2n$ we get that $|E_F| \leq 4n - 4$.

If there is an outer walk, then we can only conclude that $|E_F| \leq 2y + 4$. However, in this case the topmost and bottommost horizontal segments and the leftmost and rightmost vertical segments do not contribute any endpoints to y, thus $y \leq 2n - 8$ and therefore $|E_F| \leq 4n - 12$.

It is easy to see that this bound is tight by considering the outer face in an arrangement of h $(1 < h < n - 1)$ horizontal segments and $n - h$ vertical segments in a grid-like arrangement. □

The last ingredient needed for proving Lemma 10.2 is the following.

Lemma 10.4. *Let $\mathcal{S}$ be a set of red and blue curves, such that any two curves of the same color do not intersect. Suppose that $\mathcal{C}$ is a set of pairwise disjoint curves such that each curve in $\mathcal{C}$ intersects exactly a distinct pair of curves from $\mathcal{S}$ (each of them possibly several times) and there is a face F of the arrangement $\mathcal{A}_{\mathcal{S}}$ that is intersected by every curve in $\mathcal{C}$. Then $|\mathcal{C}| \leq 5|\mathcal{S}| + 2|F| - 3$.*

Proof. It follows from Lemma 10.1 that there are at most $3|\mathcal{S}| - 3$ curves in $\mathcal{C}$ such that each of them intersects two curves of the same color.

Denote by E_r the red edges of F and let $\mathcal{C}_{rb} \subseteq \mathcal{C}$ be the curves that intersect a red edge of F and a blue curve. First we will bound $|\mathcal{C}_{rb}|$. Remove every red edge of $\mathcal{A}_{\mathcal{S}}$ which is not in E_r. Note that this might break a red curve into several pieces. Now we eliminate the possible multiple crossings of $\mathcal{C}_{rb}$ with members of $\mathcal{S}$. Every curve in $\mathcal{C}_{rb}$ has a sub-curve such that one of its endpoints is on an edge from E_r, its other endpoint is on a blue curve and its interior does not intersect any other curve or edge. Pick such a sub-curve for every curve in $\mathcal{C}_{rb}$ and denote this set of (sub-)curves by $\mathcal{C}'_{rb}$.

Clearly $\mathcal{D}(\mathcal{S}, \mathcal{C}_{rb})$ has as many edges as $\mathcal{D}(E_r \cup \mathcal{S}_b, \mathcal{C}'_{rb})$, where $\mathcal{S}_b \subseteq \mathcal{S}$ denotes the set of blue curves. Observe also that it follows from Lemma 10.1 that $\mathcal{D}(E_r \cup \mathcal{S}_b, \mathcal{C}'_{rb})$ is a bipartite planar graph. Thus, $\mathcal{D}(E_r \cup \mathcal{S}_b, \mathcal{C}'_{rb})$ has at most $2(|E_r| + |\mathcal{S}_b|)$ edges.

By applying the same argument for the similarly defined curves $\mathcal{C}_{br} \subseteq \mathcal{C}$ that intersect a blue edge of F and a red curve, we conclude that $|\mathcal{C}| \leq 2(|E_r| + |\mathcal{S}_b|) + 2(|E_b| + |\mathcal{S}_r|) + 3|\mathcal{S}| - 3 = 5|\mathcal{S}| + 2|F| - 3$. □

Lemma 10.2 now follows from Lemma 10.3 and Lemma 10.4. From these we get that the number of edges in $\mathcal{D}(\mathcal{S}, \mathcal{C})$ is upper bounded by $5|\mathcal{S}| + 2|F| \leq 5|\mathcal{S}| + 2(4|\mathcal{S}| - 4) = 13|\mathcal{S}| - 11$.

10.3 Proof of Theorem 2.47

In this section we prove Theorem 2.47. We will use the following bound on *Davenport-Schinzel sequences.*

Lemma 10.5 ([127]). *Let S be a sequence that consists of n symbols, such that no two consecutive symbols are the same and S does not contain a sub-sequence of the form a, b, a, b. Then the length of S is at most $2n - 1$.*

We will also need the following fact.

Proposition 10.6. *Let $S_1 = a_1, \ldots, a_n$ and $S_2 = b_1, \ldots, b_n$ be two sequences, such that there is no index j for which $a_j = a_{j+1}$ and $b_j = b_{j+1}$ (that is, if $a_j = a_{j+1}$, then $b_j \neq b_{j+1}$ and vice versa). Then there is $i \in \{1, 2\}$ such that S_i contains a sub-sequence of length at least $\lceil (n+1)/2 \rceil$ in which every element is different from its preceding element in the sub-sequence.*

Proof. We construct two sub-sequences S'_1 and S'_2 in an incremental and greedy manner. First, S'_1 contains only a_1 and S'_2 contains only b_1. Then, for every $j > 1$ we append a_j (resp., b_j) to the sub-sequence S'_1 (resp., S'_2) if it is different from the last element that was added to the sub-sequence. Clearly, for every $j > 1$ at least one sub-sequence is extended, otherwise there is j such that $a_j = a_{j+1}$ and $b_j = b_{j+1}$. Therefore, the total length of S'_1 and S'_2 is at least $2 + (n-1)$ and thus one of them is of length at least $\lceil (n+1)/2 \rceil$. □

Let $\mathcal{S}$ be a set of n red and blue curves such that no two curves of the same color intersect. Suppose that $\mathcal{C}$ is a set of pairwise disjoint curves grounded with respect to $\mathcal{S}$, such that each of them intersects exactly a distinct pair of curves from $\mathcal{S}$. Recall that we wish to show that $|\mathcal{C}| = O(n)$. Note that we cannot use Lemma 10.4 since the size of a face in $\mathcal{A}_{\mathcal{S}}$ can be arbitrarily large.

Observe first that it follows from Lemma 10.1 that there are at most $3n - 3$ curves in $\mathcal{C}$ such that each of them intersects two curves in $\mathcal{S}$ of the same color. We thus discard such curves from $\mathcal{C}$ and assume henceforth that every curve in $\mathcal{C}$ connects two curves from $\mathcal{S}$ of different colors.

Let F be a face of $\mathcal{A}_{\mathcal{S}}$ such that every curve in $\mathcal{C}$ intersects its boundary. By trimming every curve in $\mathcal{C}$ if necessary, we may assume that each such curve has one of its endpoints on an edge of the boundary of F that belongs to a red (resp., blue) curve $s \in \mathcal{S}$, its other endpoint is on a blue (resp., red) curve $s' \in \mathcal{S}$, and its interior does not intersect any curve in $\mathcal{C} \cup \mathcal{S}$ except for possibly s and s'.

Let $\mathcal{C}_1 \subseteq \mathcal{C}$ be the curves with exactly one endpoint on the boundary of F and let $\mathcal{C}_2 = \mathcal{C} \setminus \mathcal{C}_1$ be the curves with two endpoints on the boundary of F. Note that we may assume that the interior of every curve $c_1 \in \mathcal{C}_1$ intersects at most one curve, namely the curve that contains that edge of F that c_1 intersects. Furthermore, we may assume that the interior of every curve $c_2 \in \mathcal{C}_2$ does not intersect any curve (including s and s', as defined above). Indeed, otherwise c_2 contains a sub-curve that qualifies for $\mathcal{C}_1$ and we may replace c_2 with this sub-curve (see Figure 43 for an illustration).

Lemma 10.7. $|\mathcal{C}_1| \leq 8n - 4$.

Proof. We can assume that the interior of each curve in $\mathcal{C}_1$ is disjoint from the interior of F, otherwise we could take one of its sub-curves. The boundary of F may consist of several connected components. Let s be a curve in $\mathcal{S}$. Clearly, it is impossible that s contributes to edges on two different components. Furthermore, it is also impossible that there are curves $c, c' \in \mathcal{C}_1$ that connect s to two points on different connected components of F, since one of c and c' must then intersect the boundary of F at two points. Therefore, it is enough to consider one connected component. That is, let $\mathcal{C}'_1 \subseteq \mathcal{C}_1$ be the curves with an endpoint on a specific connected component of the boundary of F and let $\mathcal{S}' \subseteq \mathcal{S}$ be the curves that intersect at least

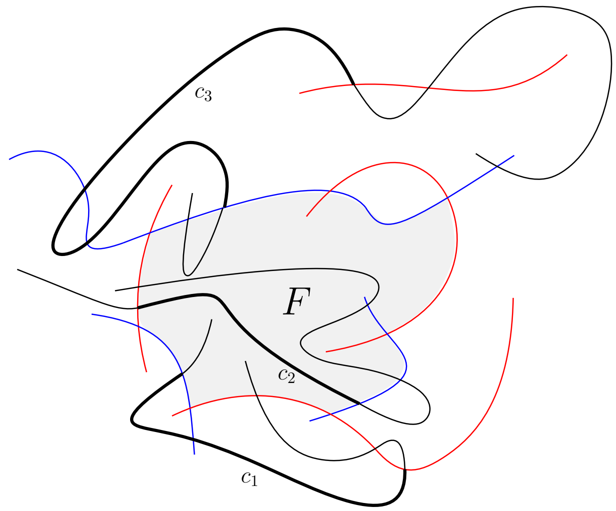

Figure 43: The sub-curves c_1 and c_3 belong to $\mathcal{C}_1$ whereas $c_2 \in \mathcal{C}_2$.

one curve from $\mathcal{C}_1'$. Then it is enough to show that $|\mathcal{C}_1'| = O(|\mathcal{S}'|)$, since by summing for every connected component of the boundary of F we get $|\mathcal{C}_1| = O(|\mathcal{S}|)$.

Let $\mathcal{C}_1'' \subseteq \mathcal{C}_1'$ be the curves whose endpoints on F belong to red curves. We may assume without loss of generality that $|\mathcal{C}_1''| \geq |\mathcal{C}_1'|/2$. Let $p_1, p_2, \ldots, p_k$ be the endpoints of the curves in $\mathcal{C}_1''$ on F listed in their cyclic order along the boundary of F. For $i = 1, 2, \ldots, k$, let $c_i \in \mathcal{C}_i''$ be the curve whose endpoint is p_i and let r_i and b_i be the red and blue curves, respectively, that are connected by c_i. Observe that it is possible that $b_i = b_j$ or $r_i = r_j$ for $i \neq j$. However, it is clearly impossible that for some i we have $b_i = b_{i+1}$ and $r_i = r_{i+1}$, otherwise c_i and c_{i+1} would connect the same pair of curves. Therefore, it follows from Proposition 10.6 that one of the sequences $S_1 = b_1, b_2, \ldots, b_k$ and $S_2 = r_1, r_2, \ldots, r_k$ contains a sub-sequence of length at least $\lceil (k+1)/2 \rceil \geq k/2$.

If neither S_1 nor S_2 does contain a sub-sequence $b_i, b_j, b_{i'}, b_{j'}$ such that $b_i = b_{i'} \neq b_j = b_{j'}$ then it follows from Lemma 10.5 that $k/2 \leq 2|\mathcal{S}'| - 1$ and hence $|\mathcal{C}_1''| \leq 4|\mathcal{S}'| - 2$ and $|\mathcal{C}_1| \leq 8n - 4$. The following two propositions thus finish the proof of Lemma 10.7.

Proposition 10.8. *S_1 does not contain a sub-sequence $b_i, b_j, b_{i'}, b_{j'}$ such that $b_i = b_{i'} \neq b_j = b_{j'}$.*

Proof. Suppose for contradiction that such a sub-sequence exists. Let α_i be the curve that consists of $c_i, c_{i'}$ and the sub-curve of b_i between the endpoints of c_i and $c_{i'}$ on b_i. Define α_j analogously. Consider a closed curve c' that consists of α_i and a curve within F that connects p_i and $p_{i'}$. Due to the alternating sub-sequence of symbols, the points p_j and $p_{j'}$ lie on different sides of c'. It follows that α_j, which connects these two points, must cross c', which is impossible since b_i and b_j do not intersect and the interiors of the curves in $\mathcal{C}$ are intersection-free. □

Proposition 10.9. *S_2 does not contain a sub-sequence $r_i, r_j, r_{i'}, r_{j'}$ such that $r_i = r_{i'} \neq r_j = r_{j'}$.*

Proof. Suppose for contradiction that such a sub-sequence exists. Consider a closed curve c' that consists of the subcurve of r_i between p_i and $p_{i'}$ and a curve within F that connects p_i and $p_{i'}$. Due to the alternating sub-sequence of symbols, the points p_j and $p_{j'}$ lie on different sides of c'. It follows that r_j, which connects these two points, must cross c', which is impossible. □

□

Lemma 10.10. $|\mathcal{C}_2| \leq 22n - 18$.

Proof. Recall that the boundary of F may consist of several connected components. We consider first the curves $\mathcal{C}_2' \subseteq \mathcal{C}_2$ which connect edges of $\mathcal{A}_{\mathcal{S}}$ that belong to different connected components of the boundary of F.

For each connected component choose either 'red' or 'blue' uniformly at random. If 'red' (resp., 'blue') is chosen for a certain connected component, then we delete all the red (resp., blue) curves that contain edges of this connected component. Every curve in $\mathcal{C}_2'$ survives (with probability 1/4) if both curves that contain its endpoints survive. Since red and blue curves from different components do not intersect, it follows from Lemma 10.1 that the surviving curves define a bipartite plane graph whose vertices correspond to surviving red and blue curves and its edges correspond to surviving curves in $\mathcal{C}_2'$. The expected number of vertices of this graph is at most $n/2$ and the expected number of edges of this graph is $|\mathcal{C}_2'|/4$. Thus we have $|\mathcal{C}_2'|/4 \leq 2 \cdot n/2$ and hence $|\mathcal{C}_2'| \leq 4n$.

It remains to bound the number of curves in $\mathcal{C}_2$ that connect edges of $\mathcal{A}_{\mathcal{S}}$ that belong to the same connected component of the boundary of F. Let $\mathcal{C}_2'' \subseteq \mathcal{C}_2$ be such curves for a certain connected component and let $\mathcal{S}'' \subseteq \mathcal{S}$ be the set of curves that contain edges of this component. Note that it is enough to show that $|\mathcal{C}_2''| = O(|\mathcal{S}''|)$, since then by summing for every connected component of the boundary of F we get $|\mathcal{C}_2| = O(|\mathcal{S}|)$.

Denote the edges of the certain connected component of the boundary of F that we consider by $e_1, e_2, \ldots, e_k$ ordered as they follow each other on the boundary of F (cut edges appear twice), and let $s_1, s_2, \ldots, s_k$ be the corresponding curves (with possible repetitions). We partition $e_1, e_2, \ldots, e_k$ into *alternating runs*: The first alternating run is the maximal sequence of edges $e_1, e_2, \ldots, e_i$ that belong to exactly two curves (necessarily a blue curve and a red curve). The next alternating run is the maximal sequence of edges $e_{i+1}, e_{i+2}, \ldots, e_j$ that belong to exactly two curves, and so on and so forth. Let m denote the number of alternating runs. The following proposition will be needed to finish the proof of Lemma 10.10.

Proposition 10.11. $m \leq 2|\mathcal{S}''| - 1$.

Proof. Let S_1 be the sequence we get by starting with s_1 and adding s_{2i+1} if it is different from s_{2i-1}, for $i = 1, 2, \ldots, \lfloor k/2 \rfloor$. Similarly, let S_2 be the sequence we get by starting with s_2 and adding s_{2i} if it is different from s_{2i-2}, for $i = 1, 2, \ldots, \lfloor k/2 \rfloor$. By definition, every element in these sequences is different from its preceding element. Furthermore, after each alternating run an element is added to one of S_1 and S_2, thus their total length is at least m. On the other hand, as in the proof of Proposition 10.9 we may conclude that none of S_1 and S_2 contains a sub-sequence of the form a, b, a, b. Therefore, by Lemma 10.5 the total length of S_1 and S_2 is at most $2|\mathcal{S}''| - 1$. Thus, $m \leq 2|\mathcal{S}''| - 1$. □

Clearly, $\mathcal{C}_2''$ contains at most m curves that connect two edges which belong to the same alternating run. Next, we bound the number of remaining curves $\mathcal{C}_2'''$, that is, those connecting edges from different alternating runs. To this end, for each alternating run choose either 'red' or 'blue' uniformly at random. If 'red' (resp., 'blue') is chosen for a certain alternating run, then we delete all red (resp., blue) edges of that run and the curves in $\mathcal{C}_2'''$ that have an endpoint on one of them. Thus, every curve in $\mathcal{C}_2'''$ survives with probability 1/4.

Consider the graph such that each of its vertices is the union of the remaining edges of a certain alternating run and whose edges correspond to surviving curves in $\mathcal{C}_2'''$. By Lemma 10.1

it is a planar (bipartite) graph. Since the expected number of surviving curves in $\mathcal{C}_2'''$ is $|\mathcal{C}_2'''|/4$, we conclude that $|\mathcal{C}_2''| \leq 4 \cdot 2m + m = 9m \leq 18|\mathcal{S}''| - 18$.

By summing for every connected component of the boundary of F and recalling that $|\mathcal{C}_2'| \leq 4n$, we have $|\mathcal{C}_2| \leq 22n - 18$. □

From Lemma 10.7 and Lemma 10.10 we have $|\mathcal{C}| \leq 3n - 3 + 8n - 4 + 22n - 18 = 33n - 25$. This concludes the proof of Theorem 2.47.

10.4 An application for conflict-free coloring of L-shapes

An *L-shape* consists of a vertical line-segment and a horizontal line-segment such that the left endpoint of the horizontal segment coincides with the bottom endpoint of the vertical segment (as in the letter 'L', hence the name). Whenever we consider a family L-shapes, we always assume that no pair of them have overlapping segments, that is, they have at most one intersection point.

A family of L-shapes is *grounded* if there is a horizontal line that contains the top point of each L-shape in the family. A (grounded) *L-graph* is a graph that can be represented as the intersection graph of a family of (grounded) L-shapes. Gonçalves et al. [57] proved that every planar graph is an L-graph. The line graph of every planar graph is also known to be an L-graph [49]. As for grounded L-graphs, McGuinness [93] proved that they are χ-bounded. Jelínek and Töpfer [65] characterized grounded L-graphs in terms of vertex ordering with forbidden patterns.

Keller, Rok and Smorodinsky [70] studied *conflict-free* colorings of string graphs[26] and in particular of grounded L-graphs. Recall that a coloring of the vertices of a hypergraph is *conflict-free*, if every hyperedge contains a vertex whose color is not assigned to any of the other vertices of the hyperedge. The minimum number of colors in a conflict-free coloring of a hypergraph $\mathcal{H}$ is denoted by $\chi_{CF}(\mathcal{H})$.

We have defined the intersection hypergraph of two families of geometric shapes. If the same regions appear in both families, in particular, when the two families coincide, there is a slightly different option. Given a family of shapes $\mathcal{S}$, the intersection hypergraph $\mathcal{I}(\mathcal{S},\mathcal{S})$ is equal to the so called *closed neighborhood hypergraph* of the intersection graph of $\mathcal{S}$, the one we define next will be equal to the *punctured* or *open neighborhood hypergraph* of the intersection graph of $\mathcal{S}$. Let the hypergraph $\mathring{\mathcal{I}}(\mathcal{S})$ be as follows: the vertex set is $\mathcal{S}$ and for every shape $S \in \mathcal{S}$ there is a hyperedge that consists of all the members of $\mathcal{S} \setminus \{S\}$ whose intersection with S is non-empty.[27] Similarly, for two families of shapes, $\mathcal{S}$ and $\mathcal{F}$, the hypergraph $\mathring{\mathcal{I}}(\mathcal{S},\mathcal{F})$ has $\mathcal{S}$ as its vertex set and has a hyperedge for every $F \in \mathcal{F}$ which consists of all the members of $\mathcal{S} \setminus \{F\}$ whose intersection with F is non-empty. Hence, $\mathring{\mathcal{I}}(\mathcal{S}) = \mathring{\mathcal{I}}(\mathcal{S},\mathcal{S})$.

Using Theorem 2.47 we can improve the following result.

Theorem 10.12 (Keller, Rok and Smorodinsky [70]). *$\chi_{CF}(\mathring{\mathcal{I}}(\mathcal{J})) = O(\log^3 n)$ for every set $\mathcal{J}$ of n grounded L-shapes. Furthermore, for every n there exists a set $\mathcal{J}$ of n grounded L-shapes such that $\chi_{CF}(\mathring{\mathcal{I}}(\mathcal{J})) = \Omega(\log n)$.*

In order to obtain this result and many other results considering (conflict-free) coloring of hypergraphs it is often enough to consider the chromatic number of the Delaunay graph. Recall that for two families of geometric shapes $\mathcal{S}$ and $\mathcal{F}$, the Delaunay graph of $\mathcal{S}$ and $\mathcal{F}$, denoted

[26] A *string graph* is the intersection graph of curves in the plane.

[27] If two shapes give rise to the same hyperedge, then this hyperedge appears only once in the hypergraph. And, as always, we omit hyperedges of size smaller than 2.

by $\mathcal{D}(\mathcal{S},\mathcal{F})$, is the graph whose vertex set is $\mathcal{S}$ and whose edge set consists of pairs of vertices such that there is a member of $\mathcal{F}$ that intersects exactly the shapes that correspond to these two vertices and no other shape. Note that if $\mathcal{S}$ is a set of planar points and $\mathcal{F}$ is the family of all disks, then $\mathcal{D}(\mathcal{S},\mathcal{F})$ is the standard Delaunay graph of the point set $\mathcal{S}$.

A key ingredient in the proof of Theorem 10.12 in [70] is the following lemma.

Lemma 10.13 ([70, Proposition 3.9])**.** *Let $\mathcal{J}\cup\mathcal{I}$ be a set of grounded L-shapes such that $|\mathcal{J}|=n$ and the L-shapes in $\mathcal{I}$ are pairwise disjoint. Then $\mathcal{D}(\mathcal{J},\mathcal{I})$ has $O(n\log n)$ edges.*

Theorem 2.47 clearly implies Lemma 10.13, since every family of n L-shapes consists of n pairwise disjoint horizontal (red) segments and n pairwise disjoint vertical (blue) segments. Thus, Theorem 2.47 is a twofold improvement of Lemma 10.13: we consider a more general setting and prove a better upper bound. Furthermore, our proof is simpler than the proof of Lemma 10.13 in [70] (especially the weak version of Theorem 2.47 that we prove in Section 10.2).

By replacing Lemma 10.13 with Theorem 2.47 (or its weaker version) in the proof of Theorem 10.12 we obtain a better upper bound for the number of colors that suffice to conflict-free color n grounded L-shapes.

Theorem 10.14. *Let $\mathcal{J}$ be a set of n grounded L-shapes. Then it is possible to color every L-shape in $\mathcal{J}$ with one of $O(\log^2 n)$ colors such that for each $\ell\in\mathcal{J}$ there is an L-shape with a unique color among the L-shapes whose intersection with ℓ is non-empty.*

Theorem 3.3 implies also that:

Corollary 10.15. *Let $\mathcal{S}$ be a set of n red and blue curves, such that no two curves of the same color intersect and let $\mathcal{C}$ be another set of pairwise disjoint curves which is grounded with respect to $\mathcal{S}$. Then the chromatic number of the intersection hypergraph $\mathcal{H}(\mathcal{S},\mathcal{C})$ and its VC-dimension are bounded by a constant, and for every k the number of hyperedges of size at most k in $\mathcal{H}(\mathcal{S},\mathcal{C})$ is $k^{O(1)}n$. Also there exists a randomized polynomial-time $O(1)$-approximation algorithm for the minimum weight hitting set problem for $\mathcal{H}$.*

Note that the upper bounds on the chromatic number and VC-dimension can be deduced easily in this case without using Theorem 3.1 (see the remarks at the end of this section).

In this section we outline the proof of the upper bound of Theorem 10.12 from [70], and indicate how using Theorem 2.47 (or just Lemma 10.2) improves on this by a $\log n$ factor, proving Theorem 10.14.

The proof in [70] uses a divide-and-conquer approach. The family of n grounded L-shapes is partitioned with respect to some vertical line into three sets, $\mathcal{J}_1$, $\mathcal{J}_2$ and $\mathcal{J}_3$, such that: $|\mathcal{J}_1|,|\mathcal{J}_3|\le n/2$; no L-shape from $\mathcal{J}_1$ intersects an L-shape from $\mathcal{J}_3$; each L-shape in $\mathcal{J}_i$ intersects some L-shape in $\mathcal{J}_i$, for $i=1,2,3$; and $\chi_{CF}(\mathring{\mathcal{I}}(\mathcal{J}_2))=O(\log^2 n)$. By using the same set of colors for $\mathcal{J}_1$ and $\mathcal{J}_3$, and applying induction the $O(\log^3 n)$ upper bound on $\chi_{CF}(\mathring{\mathcal{I}}(\mathcal{J}))$ follows.

In order to show that $\chi_{CF}(\mathring{\mathcal{I}}(\mathcal{J}_2))=O(\log^2 n)$, Keller et al. [70] show that $\chi(\mathring{\mathcal{I}}(\mathcal{J}_2'))=O(\log|\mathcal{J}_2'|)$ for any $\mathcal{J}_2'\subseteq\mathcal{J}_2$ and rely on the following known result.

Lemma 10.16. *[62, 70] If an n-vertex hypergraph $\mathcal{H}$ as well as any induced sub-hypergraph of $\mathcal{H}$ can be properly colored with t colors, then $\chi_{CF}(\mathcal{H})=O(t\log n)$.*

The proper coloring of $\mathring{\mathcal{I}}(\mathcal{J}_2)$ (and its induced sub-hypergraphs) is obtained by representing $\mathcal{J}_2$ as the union of four subsets, with their notation $\mathcal{J}_2=\mathcal{I}\cup\mathcal{F}_M\cup\mathcal{V}_R'\cup\mathcal{V}_L'$, where $\mathcal{I}$ is a set of pairwise disjoint L-shapes such that each of them intersects exactly two other L-shapes. Then it is basically shown that $\chi(\mathring{\mathcal{I}}(\mathcal{J}_2, \mathcal{F}_M \cup \mathcal{V}'_R \cup \mathcal{V}'_L)) = O(1)$ and that $\chi(\mathring{\mathcal{I}}(\mathcal{J}_2, \mathcal{I})) = \chi(\mathcal{D}(\mathcal{J}_2, \mathcal{I})) = O(\log n)$, thus concluding that $\chi(\mathring{\mathcal{I}}(\mathcal{J}_2)) = O(\log n)$. The bound $\chi(\mathcal{D}(\mathcal{J}_2, \mathcal{I})) = O(\log n)$ follows from the $O(n \log n)$ bound on the number of edges in $\mathcal{D}(\mathcal{J}_2, \mathcal{I})$ as stated in Lemma 10.13. By substituting this bound with our linear upper bound that follows from Lemma 10.2 or Theorem 2.47, we get that $\chi(\mathcal{D}(\mathcal{J}_2, \mathcal{I}))$ is upper bounded by a constant and therefore so is $\chi(\mathring{\mathcal{I}}(\mathcal{J}_2))$. Then by Lemma 10.16 we have $\chi_{\mathrm{CF}}(\mathring{\mathcal{I}}(\mathcal{J}_2)) = O(\log n)$ and the bound $\chi_{\mathrm{CF}}(\mathring{\mathcal{I}}(\mathcal{J})) = O(\log^2 n)$ follows by induction as before.

Remarks. Theorem 2.47 and Theorem 3.1 together imply Corollary 10.15. Observe also that Lemma 10.2 and Theorem 3.1 imply that there exists a proper 26-coloring of $\mathcal{H}(\mathcal{S}, \mathcal{C})$, for every family $\mathcal{S}$ of n axis-parallel segments and a family $\mathcal{C}$ of pairwise disjoint curves. Furthermore, this hypergraph has VC-dimension at most 26 and has $O(k^{25} n)$ hyperedges of size at most k.

However, it is easy to get better upper bounds on the chromatic number and the VC-dimension of these hypergraphs, even when $\mathcal{S}$ consists of red and blue curves instead of segments. Indeed, the VC-dimension is at most 8, since a shattered set of 9 vertices would imply $5 = \lceil \frac{9}{2} \rceil$ shattered curves of the same color, which would give a plane drawing of K_5. Furthermore, recall that the Delaunay graph of the red (resp., blue) curves with respect to the curves in $\mathcal{C}$ is planar, and therefore 8 colors suffice for coloring the red and blue curves such that the corresponding hypergraph is properly colored (four different colors are used for each of the two).

11 Tangencies among red and blue curves

The results of this chapter appeared in [83], a joint work with Pálvölgyi.

This chapter is organized as follows. In Section 11.1, we consider the case when the curves are x-monotone. In Section 11.2, we present our lower bound construction. In Section 11.3, we present our upper bounds for the doubly-grounded case.

11.1 Upper bound for the x-monotone case

For the proof, we will use the following result from [55].

Theorem 11.1 (Gerbner et al. [55])**.** *An n-vertex graph with a total ordering $<$ on its edges such that it does not contain as a subgraph a path on 5 vertices, $\{a, b, c, d, e\}$, whose edges are ordered as $ab < cd < bc < de$, can have at most $O(n \log n)$ edges.*

Proof of Theorem 2.41. Let $\mathcal{S}$ be the family of n red and blue x-monotone curves. After a slight perturbation we can assume that no two tangencies (as points in the plane) have the same x-coordinate. Let G be the graph whose vertices correspond to the curves and whose edges correspond to the tangent pairs. Let G_1 (resp. G_2) be the subgraph where we take an edge only if in a small neighborhood of the corresponding tangency the red curve is above (resp. below) the blue curve. Note that G is the edge-disjoint union of G_1 and G_2. We will show that the edges of G_1 can be ordered in such a way that it will satisfy the conditions of Theorem 11.1, thus G_1 has $O(n \log n)$ edges. Since the same holds for G_2 due to symmetry, in turn this implies the same for G itself, proving the theorem.

Order the edges of G_1 according to the x-coordinate of the corresponding tangencies. This gives a complete ordering of the edges. We need to show that G_1 cannot contain a path on 5

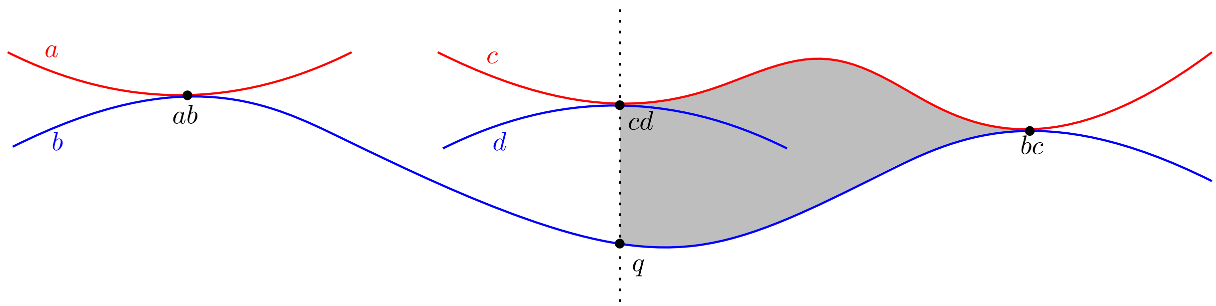

Figure 44: It is impossible to realize a certain path on 5 vertices in G_1.

vertices, $\{a, b, c, d, e\}$, such that $ab < cd < bc < de$ is the order of the edges (see Figure 44). Indeed, assume the contrary. Without loss of generality, we can assume that a corresponds to a red curve. Note that as ab (i.e., the corresponding tangency) is left to cd and cd is left to bc, the curve b must intersect the vertical line passing through cd in a point q. Moreover, as b touches c from below (in bc), q must be also below cd. This means that the part of the curve d that lies to the right of cd must be in the closed region determined by the part of c between cd and bc, the part of b between q and bc, and the vertical segment between cd and q. This implies that d has no point to the right of bc, contradicting that $bc < de$. □

11.2 Lower bound for the general case

Theorem 2.43 is a consequence of the following theorem which is also interesting in its own:

Theorem 11.2. *Given a family of (red) lines and (blue) points in the plane, there exist a corresponding doubly-grounded family of red and blue curves such that curves of the same color are disjoint and a red curve touches a blue curve if and only if the corresponding line and point are incident.*

We first show how it implies Theorem 2.43.

Proof of Theorem 2.43 from Theorem 11.2. Take a construction with n lines, $\ell_1, \ldots, \ell_n$, and n points, $p_1, \ldots, p_n$, that has $\Omega(n^{4/3})$ incidences, using the famous construction of Erdős and Purdy [43] (whose order of magnitude is in fact best possible, as shown by the Szemerédi–Trotter theorem [136]). Then apply Theorem 11.2 to get the required doubly-grounded family of curves. □

Proof of Theorem 11.2. We can suppose that none of the n lines is vertical or horizontal, and no two points form a vertical line. We can also suppose (after a possible affine transformation) that all lines have positive slopes.

Take a large axis-parallel box B such that all crossings lie within B and all lines intersect the left and right sides of B. We will convert the lines into red curves, and the points into blue curves, such that each incidence will become a tangency. Each point p_j is replaced by a small blue circle D_j whose bottommost point is p_j and a segment going upwards from the top of the circle to the top side of B (we can delete a small arc of the circle adjacent to its top so that together with the segment they form a simple curve, as required).

We will place the red curves, $\gamma_1, \ldots, \gamma_n$, one-by-one in the order given by the slopes of the corresponding lines, i.e., the order in which they intersect the left side of B, starting with the

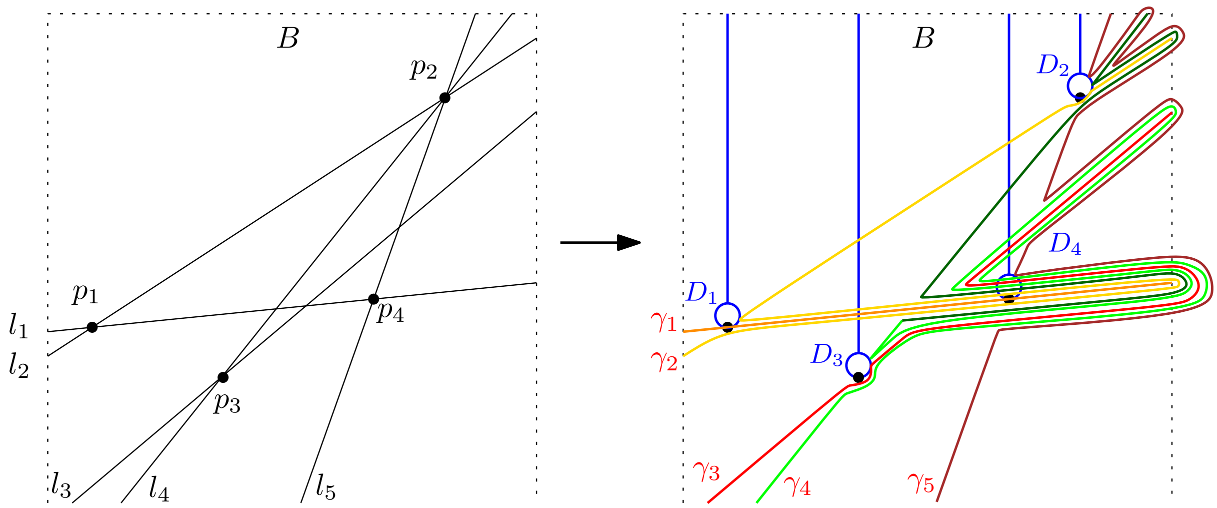

Figure 45: The lower bound construction; 'red' curves are drawn with various colors and γ_4 is drawn with two shades of green to ease readability.

topmost one (see Figure 45). The two ends of γ_i will be the intersection points of the sides of B with ℓ_i.

The curves never cross the left, top or bottom side of B but they may go slightly out at the right side. For every line ℓ_i we will also maintain the following for the corresponding red curve γ_i.

- The curve γ_i is disjoint from γ_j for every $j < i$.
- The curve γ_i lies in the closed halfplane whose boundary is ℓ_i and whose interior lies below ℓ_i.
- The curve γ_i has parts that lie on ℓ_i. Whenever γ_i leaves ℓ_i (when going from left to right) at some point $x \in \ell_i$, it later returns to ℓ_i at a point x' very close to and a bit to the right of x. Between x and x', all points of γ_i lie below ℓ_i and have an x-coordinate at least as large as the x-coordinate of x. The distance between x and x' can be chosen to be at most the diameter of the small blue circles.

In general, when we draw the next curve γ_i, it starts at the intersection point of ℓ_i and the left side of B, and starts to follow ℓ_i to the right.

Whenever ℓ_i is incident to some point p_j, it intersects the disk D_j in two points, one of which is p_j; denote the other intersection point by p_j^i. Since the lines were ordered by their slopes, the points p_j^i (which exist) follow each other in a counterclockwise order as i increases; see Figure 46. Recall also that all lines have positive slopes.

We draw the curve γ_i so that it touches D_j at p_j^i. If there is no already drawn curve touching D_j, then the curve γ_i leaves ℓ_i slightly before p_j and goes slightly outside of the disk along the $p_j p_j^i$ arc, then returns to ℓ_i at p_j^i, and continues along ℓ_i.

However, if there is some $h < i$ for which ℓ_h is incident to p_j (and thus the corresponding curve touches D_j), then instead of going alongside the arc, γ_i might be forced to make a much larger detour; the description of this detour will be detailed in the next paragraph. The only important part for us here is that whenever the first time γ_i returns from the detour so that it is not locally separated from D_j anymore, it touches D_j at p^i_j, and continues along ℓ_i; see Figure 46.

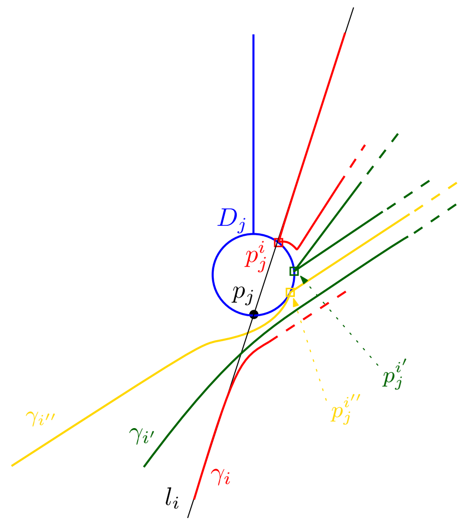

Figure 46: The way γ_i is drawn in the vicinity of p_j.

Whenever γ_i would intersect another red curve γ_h, $h < i$, at some point x, instead of creating an intersection, using that we ordered the lines by their slopes, we follow γ_h, going very close to it, until we reach the endpoint of γ_h on the right side of B. Then we go back on the other side of γ_h, again very close to it, until we arrive back at some point $x' \in \gamma_i \cap \ell_i$ in a close vicinity of x.

Notice that if a line ℓ_i goes through some p_j, then γ_i will go in the above defined way one-by-one around every curve drawn earlier that touches D_j, and only then will it touch D_j at p^i_j.

In case of going around some γ_h, the part of γ_h which γ_i went around lies below ℓ_h and to the right of $\ell_i \cap \ell_h$, thus also below ℓ_i. This maintains the condition for γ_i as well, and it also implies that while going around some γ_h, γ_i can never intersect any disk D_j that it is supposed to touch.

To see that this construction can be made doubly-grounded, we can easily deform it to lie in a vertical strip. We can transform the left side of B into a segment on the left side of the strip, and the top side of B into a segment on the right side of the strip. $\square$

11.3 Conditional upper bound for the doubly-grounded case

Our primary aim in this section is to prove Theorem 2.44.

We first define a certain redrawing of the curves to make some arguments easier (see Figure 47). For each curve γ we choose an arbitrary point on it, p_γ. Then, we draw a star-shape γ^* such that $p_\gamma \in \gamma^*$ will be a special point in γ^*, called the center, and for each tangency x of γ and some other curve of the family, we draw a simple arc from p_γ to x that goes very close to the original curve, such that the curves starting at p_γ are pairwise disjoint. Moreover, whenever γ touched another curve, the arc from p_γ to x should avoid that intersection (except of course for x), and whenever γ crossed another curve, the arc from p_γ to x should also cross it (with one intersection point). We repeat this for every curve transforming each into such a star, while preserving disjointnesses and tangencies.

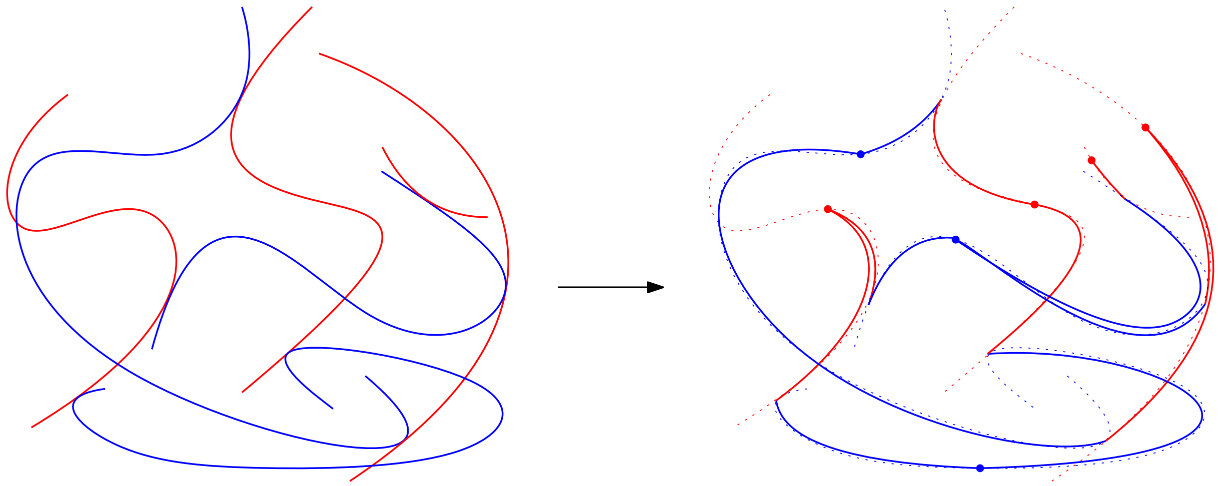

Figure 47: Drawing of a graph based on a set of curves.

This drawing can be regarded as a drawing D of a bipartite graph G in the plane, whose vertices correspond to the original curves, represented by the centers of the stars, and whose edges correspond to the tangencies. Note that each edge consists of one red and one blue part; we refer to these as *edge-parts*. Thus, the number of edges of G is exactly the number of tangencies in the original setting. Also, if two vertices are adjacent in G, then none of their neighboring edge-parts are allowed to intersect. Indeed, if uv is an edge, then the star-shapes that correspond to u and v can only meet at their point of tangency. Nor are two edge-parts of the same color allowed to intersect. These observations immediately imply the following.

Claim 11.3. *The drawing D of G obtained by the above redrawing process cannot contain a self-crossing path on five edge-parts. Therefore, it can contain neither a self-crossing C_4, nor a self-crossing path on two edges.*

Proof. Consider a path with 5 edge-parts and assume, without loss of generality, that the five edge-parts of a path are red, red, blue, blue, red. Denote the vertex between the first two red parts by u, the vertex between the two blue parts by v, and the vertex incident to the last red part by w. Recall that edge-parts of the same color do not intersect. Both blue edge-parts are incident to v, so they cannot intersect the edge-parts incident to a vertex adjacent to v. But each red edge-part is incident to either u or w, which are both adjacent to v. Therefore, the path cannot be self-crossing. □

In [91] it was proved that a graph G which can be drawn without a self-crossing C_4 has at most $O(n^{3/2}\log n)$ edges. As the number of edges of G equals the number of tangencies between curves, this implies the upper bound of $O(n^{3/2}\log n)$ on the number of tangencies between the curves that was mentioned in the introduction. Now we use this representation to prove our better upper bound in case the curves are doubly-grounded.

Our main result of this section shows an upper bound matching the lower bound for doubly-grounded curves, assuming a conjecture about forbidden 0-1 matrices holds. For more details about the following definitions see, e.g., [108]. Briefly, a 0-1 matrix M is said to contain another 0-1 matrix P if there exists a submatrix M' of M such that we can get P from M' by possibly replacing some 1-entries by 0-entries.[28] A 0-1 matrix is a cycle if it is the adjacency matrix of a cycle graph. We omit the definition of positive orthogonal cycles in general, we just say that positive orthogonal 6-cycles are those 3×3 matrices that are 6-cycles and their middle entry (i.e., the one in the second row and second column) is a 1-entry.

Pach and Tardos [108] proved the following.

Theorem 11.4 (Pach and Tardos [108], Theorem 3)**.** *If a* 0-1 *matrix avoids all positive orthogonal cycles, then it has* $O(n^{4/3})$ 1*-entries.*

In fact, it may be enough to forbid only positive orthogonal 6-cycles to achieve the same bound, as mentioned in [108].

Conjecture 11.5. *If a* 0-1 *matrix avoids all positive orthogonal* 6*-cycles, then it has* $O(n^{4/3})$ 1*-entries.*

Actually, it may be enough to forbid only one 6-cycle and it may not even be necessary that it is a positive orthogonal cycle, as conjectured in [45] and [94].[29]

Conjecture 11.6. *If a* 0-1 *matrix avoids a given* 6*-cycle, then it has* $O(n^{4/3})$ 1*-entries.*

Theorem 2.44 states that Conjecture 11.5, the weaker one of these two conjectures, would already imply that our lower bound is optimal for doubly-grounded curves.

Proof of Theorem 2.44 and Claim 2.45. Given a doubly-grounded family of curves, for each curve γ choose the point where it touches the appropriate boundary of the vertical strip to be p_γ, and do the redrawing defined at the beginning of Section 11.3. We get a bipartite graph G on n vertices whose vertices lie on the two boundaries of the strip and whose edges are drawn inside the strip.

We have seen that this graph G has no self-crossing C_4.[30] As in a doubly-grounded representation any C_4 is self-crossing, this already implies that G is C_4-free, thus has at most $O(n^{3/2})$ edges. This proves Claim 2.45.

Now we prove that additionally it does not contain a certain type of C_6 that we call *positive*, following [108]. A C_6 on vertices $u_1, u_2, u_3, v_1, v_2, v_3$ is positive if u_1, u_2, u_3 are on the left side of the strip in this order from top to bottom, v_1, v_2, v_3 are on the right side in this order from top to bottom, and u_2v_2 is an edge of the C_6.[31]

Assume on the contrary that such a positive C_6 exists. Recall that each edge u_iv_j has a red part incident to u_i and a blue part incident to v_j, and that no parts of the same color can cross.

Without loss of generality, we can assume that u_3v_1 is an edge of the C_6. The edges u_2v_2 and u_3v_1 must cross, denote (one of) their crossing point(s) by x. Without loss of generality, we can assume that the red part of u_2v_2 crosses the blue part of u_3v_1. This implies that u_2v_1 cannot be an edge, as otherwise xv_1u_2x would be a self-crossing path on 4 edge-parts contradicting Claim 11.3. So u_2v_3 and u_1v_1 need to be edges.

[28] Note that the order of the rows and columns is fixed in both M and P.

[29] In [45] they show for various subsets of the 6-cycles that forbidding them implies the $O(n^{4/3})$ upper bound on the edges, but unfortunately all of these subsets contain non-positive cycles as well, so they do not imply Conjecture 11.5.

[30] In case of a doubly-grounded family, this is trivial.

[31] In fact, it would be sufficient to assume that u_2v_2 is an edge of G, as since there is no self-crossing C_4, this can only happen if u_2v_2 is an edge of the C_6.

Then u_2x (contained in the red part of u_2v_2) and xv_1 (contained in the blue part of u_3v_1) separate u_1 from v_2 and from v_3 (within the strip). However, at least one of u_1v_2 and u_1v_3 must be an edge of G.
If u_1v_2 is an edge, then as it cannot intersect u_2v_2, it must intersect xv_1, so $v_2u_1v_1x$ is a self-crossing path on 5 edge-parts, see Figure 48(a).
If u_1v_3 is an edge, then either u_1v_3 intersects u_2x and then $u_1v_3u_2x$ is a self-crossing path on 5 edge-parts, see Figure 48(b); or u_1v_3 intersects xv_1 and then $v_3u_1v_1x$ is a self-crossing path on 5 edge-parts, see Figure 48(c).
All cases contradict Claim 11.3.

Thus, we have shown that positive C_6's are not in G. Take the adjacency 0-1 matrix of the bipartite graph G such that the rows and columns are ordered according to the order of the corresponding points along the boundaries of the strip. Recall that the positive 6-cycles are the ones in whose 3×3 matrix representation the middle entry is a 1-entry. Thus, Conjecture 11.5 implies that G has at most these many edges, which in turn implies the required upper bound on the number of tangencies in the original setting of curves, finishing the proof of Theorem 2.44. □

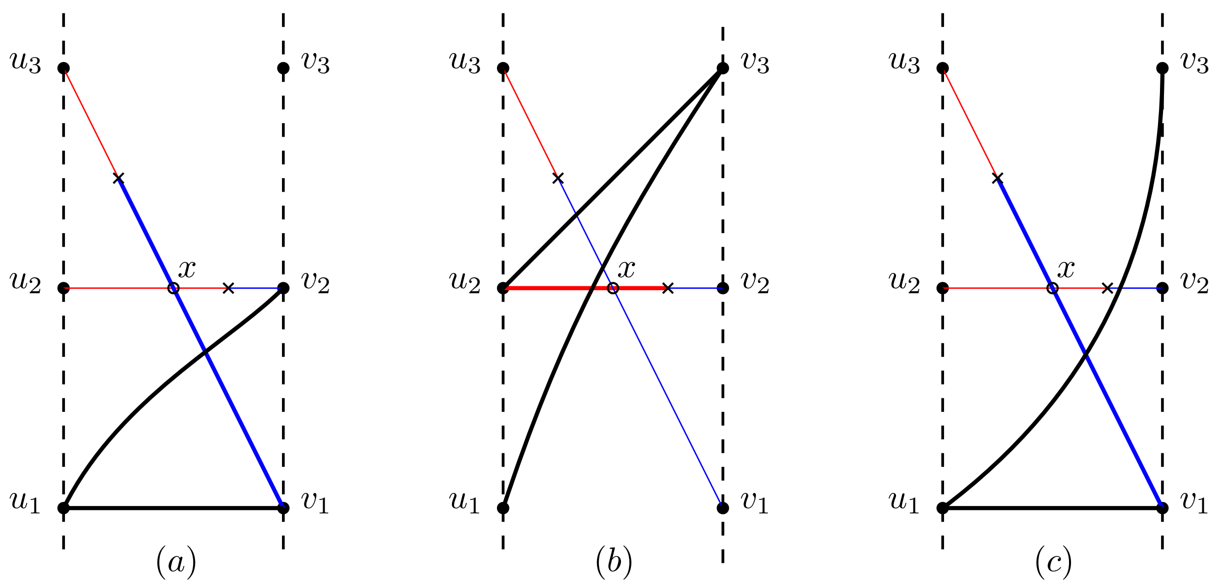

Figure 48: In every case there exists a crossing path on 5 edge-parts.

Theorem 2.44 implies the following slight extension in a standard way.

Theorem 11.7. *Given a family of n red and blue curves that lie inside a circle such that no two curves of the same color intersect and every curve touches the circle once, then the number of tangencies between the curves is $O(n^{4/3})$ if Conjecture 11.5 holds.*

Proof. We prove by induction that the number of tangencies is at most $dn^{4/3}$ for a suitable d. This certainly holds for small n. For a general n, by a continuity argument we can split the boundary circle into two arcs such that each arc touches at most $\lceil n/2 \rceil$ curves of each color. Then the number of tangencies between the red curves (resp. blue curves) touching the first arc and the blue curves (resp. red curves) touching the second arc is at most $cn^{4/3}$ by Theorem 2.44, where the c is hidden in the O notation. The tangencies between the red and blue curves touching the same arc is at most $d\lceil n/2\rceil^{4/3}$ by induction (after deforming the drawing such that the arc becomes a circle). Thus, the total number of tangencies is at most $2d\lceil n/2\rceil^{4/3} + 2cn^{4/3} \leq dn^{4/3}$ for d large enough (depending on c), finishing the proof. □

We have seen that both in the general case (see the paragraph after the proof of Claim 11.3) and in the doubly-grounded case (see Theorem 2.44) the best known upper bounds use only Claim 11.3. Besides improving the lower and upper bounds in the general case, it would be interesting to decide if it is possible to improve the general upper bound using only Claim 11.3. Also, this connection makes it even more interesting to solve Conjectures 11.5 and 11.6. Note that any upper bound for the respective problems that is better than $\tilde{O}(n^{3/2})$ would also improve our bound on the doubly-grounded case.

References

[1] Eyal Ackerman. *The Maximum Number of Tangencies Among Convex Regions with a Triangle-Free Intersection Graph*, pages 19–30. Springer New York, New York, NY, 2013.

[2] Eyal Ackerman, Gábor Damásdi, Balázs Keszegh, Rom Pinchasi, and Réka Blanka Raffay. On the number of digons in arrangements of pairwise intersecting circles. In *40th International Symposium on Computational Geometry (SoCG 2024)*, volume to be published, 2024.

[3] Eyal Ackerman, Gábor Damásdi, Balázs Keszegh, Rom Pinchasi, and Réka Blanka Raffay. On the number of digons in arrangements of pairwise intersecting pseudo-circles. Manuscript, 2024.

[4] Eyal Ackerman, Balázs Keszegh, and Dömötör Pálvölgyi. Coloring hypergraphs defined by stabbed pseudo-disks and ABAB-free hypergraphs. *SIAM Journal on Discrete Mathematics*, 34(4):2250–2269, 2020.

[5] Eyal Ackerman, Balázs Keszegh, and Mate Vizer. Coloring points with respect to squares. *Discrete & Computational Geometry*, jun 2017.

[6] Eyal Ackerman and Balázs Keszegh. On the number of tangencies among 1-intersecting x-monotone curves. *European Journal of Combinatorics*, 118:103929, 2024.

[7] Eyal Ackerman, Balázs Keszegh, and Dömötör Pálvölgyi. On tangencies among planar curves with an application to coloring l-shapes. *European Journal of Combinatorics*, page 103837, 2023.

[8] Eyal Ackerman and Rom Pinchasi. On coloring points with respect to rectangles. *Journal of Combinatorial Theory, Series A*, 120(4):811–815, 2013.

[9] Pankaj K Agarwal, Eran Nevo, János Pach, Rom Pinchasi, Micha Sharir, and Shakhar Smorodinsky. Lenses in arrangements of pseudo-circles and their applications. *Journal of the ACM (JACM)*, 51(2):139–186, 2004.

[10] Deepak Ajwani, Khaled Elbassioni, Sathish Govindarajan, and Saurabh Ray. Conflict-free coloring for rectangle ranges using o(n.382) colors. *Discrete & Computational Geometry*, 48(1):39–52, Jul 2012.

[11] Noga Alon, Hagit Last, Rom Pinchasi, and Micha Sharir. On the complexity of arrangements of circles in the plane. *Discret. Comput. Geom.*, 26(4):465–492, 2001.

[12] Greg Aloupis, Jean Cardinal, Sébastien Collette, Stefan Langerman, David Orden, and Pedro Ramos. Decomposition of multiple coverings into more parts. *Discrete & Computational Geometry*, 44(3):706–723, Oct 2010.

[13] Boris Aronov, Anirudh Donakonda, Esther Ezra, and Rom Pinchasi. On pseudo-disk hypergraphs. *Computational Geometry*, 92, 2021.

[14] Boris Aronov, Daniel Q. Naiman, János Pach, and Micha Sharir. An invariant property of balls in arrangements of hyperplanes. *Discret. Comput. Geom.*, 10:421–425, 1993.

[15] Andrei Asinowski, Jean Cardinal, Nathann Cohen, Sébastien Collette, Thomas Hackl, Michael Hoffmann, Kolja Knauer, Stefan Langerman, Michał Lasoń, Piotr Micek, Günter Rote, and Torsten Ueckerdt. Coloring hypergraphs induced by dynamic point sets and bottomless rectangles. In Frank Dehne, Roberto Solis-Oba, and Jörg-Rüdiger Sack, editors, *Algorithms and Data Structures*, pages 73–84, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

[16] Edgar Asplund and Branko Grünbaum. On a coloring problem. *Mathematica Scandinavica*, 8:181–188, Dec. 1960.

[17] Maya Bechler-Speicher. A crossing lemma for families of jordan curves with a bounded intersection number, 2019. `arXiv:1911.07287`.

[18] Claude Berge. Balanced matrices. *Mathematical Programming*, 2(1):19–31, Feb 1972.

[19] Anders Björner, Michel Las Vergnas, Bernd Sturmfels, Neil White, and Günter M. Ziegler. *Oriented Matroids*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2nd edition, 1999.

[20] Marcin Briański, James Davies, and Bartosz Walczak. Separating polynomial χ-boundedness from χ-boundedness. *Combinatorica*, 44(1):1–8, Feb 2024.

[21] James Perkins Burling. *On Coloring Problems of Families of Polytopes*. Phd thesis, University of Colorado, 1965.

[22] Balázs Bursics, Bence Csonka, and Luca Szepessy. Hitting sets and colorings of hypergraphs, 2023. `arXiv:2307.12154`.

[23] Sarit Buzaglo, Rom Holzman, and Rom Pinchasi. On s-intersecting curves and related problems. In *Proceedings of the Twenty-Fourth Annual Symposium on Computational Geometry*, SCG '08, page 79–84, New York, NY, USA, 2008. Association for Computing Machinery.

[24] Sarit Buzaglo, Rom Pinchasi, and Günter Rote. Topological hypergraphs. In *Thirty Essays on Geometric Graph Theory*, pages 71–81. Springer New York, oct 2012.

[25] Jean Cardinal, Kolja Knauer, Piotr Micek, Dömötör Pálvölgyi, Torsten Ueckerdt, and Narmada Varadarajan. Colouring bottomless rectangles and arborescences. *Computational Geometry*, 115:102020, 2023.

[26] Jean Cardinal and Matias Korman. Coloring planar homothets and three-dimensional hypergraphs. *Computational geometry*, 46(9):1027–1035, 2013.

[27] Parinya Chalermsook and Bartosz Walczak. Coloring and maximum weight independent set of rectangles. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 860–868.

[28] Jérémie Chalopin and Daniel Gonçalves. Every planar graph is the intersection graph of segments in the plane: extended abstract. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC '09, page 631–638, New York, NY, USA, 2009. Association for Computing Machinery.

[29] Timothy M. Chan. Conflict-free coloring of points with respect to rectangles and approximation algorithms for discrete independent set. In *Proceedings of the Twenty-Eighth Annual Symposium on Computational Geometry*, SoCG '12, page 293–302, New York, NY, USA, 2012. Association for Computing Machinery.

[30] 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 2012 Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1576–1585.

[31] Timothy M. Chan and Sariel Har-Peled. Approximation algorithms for maximum independent set of pseudo-disks. *Discrete & Computational Geometry*, 48(2):373–392, 2012.

[32] Bernard Chazelle, Leo J. Guibas, and D. T. Lee. The power of geometric duality. *BIT*, 25:76–90, 1985.

[33] Xiaomin Chen, János Pach, Mario Szegedy, and Gábor Tardos. Delaunay graphs of point sets in the plane with respect to axis-parallel rectangles. *Random Structures & Algorithms*, 34(1):11–23, 2009.

[34] Chaim Chojnacki. Über wesentlich unplättbare kurven im dreidimensionalen raume. *Fundamenta Mathematicae*, 23(1):135–142, 1934.

[35] Kenneth L. Clarkson and Peter W. Shor. Applications of random sampling in computational geometry, ii. *Discrete & Computational Geometry*, 4(5):387–421, 1989.

[36] Gábor Damásdi and Dömötör Pálvölgyi. Realizing an m-uniform four-chromatic hypergraph with disks. *Combinatorica*, 42(1):1027–1048, 2022.

[37] Gábor Damásdi and Dömötör Pálvölgyi. Three-Chromatic Geometric Hypergraphs. In Xavier Goaoc and Michael Kerber, editors, *38th International Symposium on Computational Geometry (SoCG 2022)*, volume 224 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 32:1–32:13, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[38] Harold Davenport and Andrzej Schinzel. A combinatorial problem connected with differential equations. *American Journal of Mathematics*, 87(3):684–694, 1965.

[39] James Davies. Improved bounds for colouring circle graphs. *Proceedings of the American Mathematical Society*, 150(12):5121–5135, 2022.

[40] Raghavan Dhandapani, Jacob E. Goodman, Andreas Holmsen, Richard Pollack, and Shakhar Smorodinsky. Convexity in topological affine planes. *Discrete & Computational Geometry*, 38(2):243–257, 2007.

[41] Jordan S. Ellenberg, József Solymosi, and Joshua Zahl. New bounds on curve tangencies and orthogonalities. *Discrete Anal.*, 11 2016.

[42] Paul Erdős. On sets of distances of n points. *Am. Math. Mon.*, 53(5):248–250, 1946.

[43] Paul Erdős. Problems and results in combinatorial geometry. Discrete geometry and convexity, Proc. Conf., New York 1982, Ann. N.Y. Acad. Sci. 440, 1–11, 1985.

[44] Paul Erdős and Branko Grünbaum. Osculation vertices in arrangements of curves. *Geom. Dedicata*, 1:322–333, 1973.

[45] Ervin Győri and Dániel Korándi and Abhishek Methuku and István Tomon and Casey Tompkins and Máté Vizer. On the Turán number of some ordered even cycles. *European Journal of Combinatorics*, 73:81–88, 2018.

[46] Guy Even, Zvi Lotker, Dana Ron, and Shakhar Smorodinsky. Conflict-free colorings of simple geometric regions with applications to frequency assignment in cellular networks. *SIAM Journal on Computing*, 33(1):94–136, jan 2003.

[47] Esther Ezra, János Pach, and Micha Sharir. On regular vertices of the union of planar convex objects. *Discrete & Computational Geometry*, 41(2):216–231, Mar 2009.

[48] S. P. Fekete and P. Keldenich. Conflict-free coloring of intersection graphs. *International Journal of Computational Geometry & Applications*, 28(3):289–307, 2018.

[49] Stefan Felsner, Kolja Knauer, George B. Mertzios, and Torsten Ueckerdt. Intersection graphs of L-shapes and segments in the plane. *Discrete Applied Mathematics*, 206:48 – 55, 2016.

[50] Stefan Felsner, Sandro Roch, and Manfred Scheucher. Arrangements of pseudocircles: on digons and triangles. In *Graph drawing and network visualization*, volume 13764 of *Lecture Notes in Computer Science*, pages 441–455. Springer, Cham, 2023.

[51] Stefan Felsner and Manfred Scheucher. Arrangements of pseudocircles: On circularizability. *Discret. Comput. Geom.*, 64(3):776–813, 2020.

[52] Jacob Fox, János Pach, and Csaba D. Tóth. Intersection patterns of curves. *Journal of the London Mathematical Society*, 83(2):389–406, 2011.

[53] Radoslav Fulek. Coloring geometric hyper-graph defined by an arrangement of half-planes, 2011. `arXiv:1002.4529`.

[54] Fănică Gavril. The intersection graphs of subtrees in trees are exactly the chordal graphs. *Journal of Combinatorial Theory, Series B*, 16(1):47–56, 1974.

[55] Dániel Gerbner, Abhishek Methuku, Dániel T. Nagy, Dömötör Pálvölgyi, Gábor Tardos, and Máté Vizer. Turán problems for edge-ordered graphs. *Journal of Combinatorial Theory, Series B*, 160:66–113, 2023.

[56] Matt Gibson and Kasturi Varadarajan. Optimally decomposing coverings with translates of a convex polygon. *Discrete & Computational Geometry*, 46(2):313–333, Sep 2011.

[57] Daniel Gonçalves, Lucas Isenmann, and Claire Pennarun. Planar graphs as L-intersection or L-contact graphs. In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 172–184. SIAM, 2018.

[58] Jacob E. Goodman, Joseph O'Rourke, and Csaba D. Tóth, editors. *Handbook of Discrete and Computational Geometry*. CRC Press, Inc., USA, 2017.

[59] Branko Grünbaum. Arrangements and spreads. In *CBMS Regional conference series in mathematics*. American Mathematical Society, 1972.

[60] Péter Györgyi, Bálint Hujter, and Sándor Kisfaludi-Bak. On the number of touching pairs in a set of planar curves. *Computational Geometry*, 67:29 – 37, 2018.

[61] A. Gyárfás. On the chromatic number of multiple interval graphs and overlap graphs. *Discrete Mathematics*, 55(2):161–166, 1985.

[62] Sariel Har-Peled and Shakhar Smorodinsky. Conflict-free coloring of points and simple regions in the plane. *Discret. Comput. Geom.*, 34(1):47–70, 2005.

[63] David Haussler and Emo Welzl. ϵ-nets and simplex range queries. *Discrete & Computational Geometry*, 2(2):127–151, Jun 1987.

[64] Andreas F. Holmsen, Hossein Nassajian Mojarrad, János Pach, and Gábor Tardos. Two extension of the Erdős-Szekeres problem. *J. Eur. Math. Soc.*, 22(12):3981–3995, 2022.

[65] Vít Jelínek and Martin Töpfer. On grounded L-graphs and their relatives. *Electron. J. Comb.*, 26(3):P3.17, 2019.

[66] Frederik Brinck Jensen, Aadi Joshi, and Saurabh Ray. Discrete Helly type theorems. In *Proceedings of the 30th Annual Canadian Conference on Computational Geometry, CCCG 2020, August 5-7, 2020, University of Saskatchewan, Saskatoon, Saskatchewan, Canada*, pages 332–335, 2020.

[67] T. Karthick and Frédéric Maffray. Vizing bound for the chromatic number on some graph classes. *Graphs and Combinatorics*, 32(4):1447–1460, Jul 2016.

[68] Klara Kedem, Ron Livne, János Pach, and Micha Sharir. On the union of jordan regions and collision-free translational motion amidst polygonal obstacles. *Discrete & Computational Geometry*, 1(1):59–71, Mar 1986.

[69] Chaya Keller, Balázs Keszegh, and Dömötör Pálvölgyi. On the Number of Hyperedges in the Hypergraph of Lines and Pseudo-Discs . *Electron. J. Comb.*, 29(3):P3.25, 2022.

[70] Chaya Keller, Alexandre Rok, and Shakhar Smorodinsky. Conflict-free coloring of string graphs. *Discrete & Computational Geometry*, 2020.

[71] Chaya Keller and Shakhar Smorodinsky. Conflict-Free Coloring of Intersection Graphs of Geometric Objects. *Discrete & Computational Geometry*, pages 1–26, 2019.

[72] Balázs Keszegh. Weak conflict-free colorings of point sets and simple regions. In Prosenjit Bose, editor, *Proceedings of the 19th Annual Canadian Conference on Computational Geometry, CCCG 2007, August 20-22, 2007, Carleton University, Ottawa, Canada*, pages 97–100. Carleton University, Ottawa, Canada, 2007.

[73] Balázs Keszegh. Coloring half-planes and bottomless rectangles. *Computational geometry*, 45(9):495–507, 2012.

[74] Balázs Keszegh. Coloring intersection hypergraphs of pseudo-disks. *Discrete & Computational Geometry*, 64(3):942–964, Oct 2020.

[75] Balázs Keszegh. Discrete Helly-type theorems for pseudohalfplanes. *European Journal of Combinatorics*, 101:103469, 2022.

[76] Balázs Keszegh, János Pach, and Dömötör Pálvölgyi. Drawing planar graphs of bounded degree with few slopes. In Ulrik Brandes and Sabine Cornelsen, editors, *Graph Drawing*, pages 293–304, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[77] Balázs Keszegh and Dömötör Pálvölgyi. Geometric hypergraph zoo.

[78] Balázs Keszegh and Dömötör Pálvölgyi. Octants are cover-decomposable. *Discrete & Computational Geometry*, 47(3):598–609, 2012.

[79] Balázs Keszegh and Dömötör Pálvölgyi. Convex polygons are self-coverable. *Discrete & Computational Geometry*, 51(4):885–895, Jun 2014.

[80] Balázs Keszegh and Dömötör Pálvölgyi. More on decomposing coverings by octants. *Journal of Computational Geometry*, 6(1):300–315, 2015.

[81] Balázs Keszegh and Dömötör Pálvölgyi. An abstract approach to polychromatic coloring: shallow hitting sets in ABA-free hypergraphs and pseudohalfplanes. *J. Comput. Geom.*, 10:1–26, 2019.

[82] Balázs Keszegh and Dömötör Pálvölgyi. Proper coloring of geometric hypergraphs. *Discrete & Computational Geometry*, 62(3):674–689, Oct 2019.

[83] Balázs Keszegh and Dömötör Pálvölgyi. The number of tangencies between two families of curves. *Combinatorica*, 43(5):939–952, Oct 2023.

[84] Balázs Keszegh. A new discrete theory of pseudoconvexity. *Discrete Mathematics & Theoretical Computer Science*, vol. 25:1, May 2023.

[85] Balázs Keszegh and Dömötör Pálvölgyi. Octants are cover-decomposable into many coverings. *Computational Geometry*, 47(5):585–588, 2014.

[86] Paul Koebe. Kontaktprobleme der konformen abbildung. *Ber. Saechs. Akad. Wiss. Leipzig, Math.-Phys. Kl.*, 88:141–164, 1936.

[87] János Komlós, János Pach, and Gerhard Woeginger. Almost tight bounds for ϵ-Nets. *Discrete & Computational Geometry*, 7(2):163–173, Feb 1992.

[88] István Kovács. Indecomposable coverings with homothetic polygons. *Discrete & Computational Geometry*, 53(4):817–824, 2015.

[89] Tamás Kővári, Vera T. Sós, and Paul Turán. On a problem of K. Zarankiewicz. *Colloquium Mathematicum*, 3:50–57, 1954.

[90] Peter Mani-Levitska and János Pach. Decomposition problems for multiple coverings with unit balls. Unpublished Manuscript, 1986.

[91] Adam Marcus and Gábor Tardos. Intersection reverse sequences and geometric applications. *J. Comb. Theory, Ser. A*, 113(4):675–691, 2006.

[92] Jiri Matousek. *Lectures on Discrete Geometry.* Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2002.

[93] Sean McGuinness. On bounding the chromatic number of L-graphs. *Discrete Mathematics*, 154(1):179 – 187, 1996.

[94] Abhishek Methuku and István Tomon. Bipartite Turán Problems for Ordered Graphs. *Combinatorica*, 42(6):895–911, Dec 2022.

[95] János Pach. personal communication.

[96] János Pach. Decomposition of multiple packing and covering. In *2. Kolloquium über Diskrete Geometrie*, number CONF, pages 169–178. Institut für Mathematik der Universität Salzburg, 1980.

[97] János Pach. Covering the plane with convex polygons. *Discrete & Computational Geometry*, 1:73–81, 1986.

[98] János Pach and Pankaj K. Agarwal. *Combinatorial Geometry*, chapter 11, pages 177–178. John Wiley and Sons Ltd, 1995.

[99] János Pach, Dömötör Pálvölgyi, and Géza Tóth. Survey on decomposition of multiple coverings. In *Geometry–Intuitive, Discrete, and Convex*, pages 219–257. Springer, 2013.

[100] János Pach, Natan Rubin, and Gábor Tardos. On the Richter-Thomassen conjecture about pairwise intersecting closed curves. *Combinatorics, Probability and Computing*, 25(6):941–958, 2016.

[101] János Pach, Natan Rubin, and Gábor Tardos. A crossing lemma for Jordan curves. *Advances in Mathematics*, 331:908–940, 2018.

[102] János Pach and Micha Sharir. On vertical visibility in arrangements of segments and the queue size in the Bentley-Ottmann line sweeping algorithm. *SIAM J. Comput.*, 20(3):460–470, 1991.

[103] János Pach and Micha Sharir. On the boundary of the union of planar convex sets. *Discrete & Computational Geometry*, 21(3):321–328, 1999.

[104] János Pach, Andrew Suk, and Miroslav Treml. Tangencies between families of disjoint regions in the plane. *Comput. Geom.*, 45(3):131–138, 2012.

[105] János Pach, Gábor Tardos, and Géza Tóth. Indecomposable coverings. In *Discrete Geometry, Combinatorics and Graph Theory*, pages 135–148. Springer, 2007.

[106] János Pach and Dömötör Pálvölgyi. Unsplittable coverings in the plane. *Advances in Mathematics*, 302:433–457, 2016.

[107] János Pach and Micha Sharir. *Sylvester-Gallai problem: The beginnings of combinatorial geometry*, pages 1–12. 12 2008.

[108] János Pach and Gábor Tardos. Forbidden paths and cycles in ordered graphs and matrices. *Isr. J. Math.*, 155:359–380, 2006.

[109] János Pach and Gábor Tardos. Coloring axis-parallel rectangles. *Journal of Combinatorial Theory, Series A*, 117(6):776–782, 2010.

[110] János Pach and Géza Tóth. Decomposition of multiple coverings into many parts. *Computational Geometry*, 42(2):127–133, 2009.

[111] Dömötör Pálvölgyi. Indecomposable coverings with concave polygons. *Discrete & Computational Geometry*, 44(3):577–588, Oct 2010.

[112] Dömötör Pálvölgyi and Géza Tóth. Convex polygons are cover-decomposable. *Discrete & Computational Geometry*, 43(3):483–496, Apr 2010.

[113] Arkadiusz Pawlik, Jakub Kozik, Tomasz Krawczyk, Michał Lasoń, Piotr Micek, William T. Trotter, and Bartosz Walczak. Triangle-free intersection graphs of line segments with large chromatic number. *Journal of Combinatorial Theory, Series B*, 105:6–10, 2014.

[114] Rom Pinchasi. The zone theorem revisited. Manuscript, 2011.

[115] Rom Pinchasi. A finite family of pseudodiscs must include a "small" pseudodisc. *SIAM Journal on Discrete Mathematics*, 28(4):1930–1934, jan 2014.

[116] Rom Pinchasi and Radoš Radoičić. Topological graphs with no self-intersecting cycle of length 4. In *Proceedings of the Nineteenth Annual Symposium on Computational Geometry*, SCG '03, pages 98–103, New York, NY, USA, 2003. Association for Computing Machinery.

[117] Tim Planken and Torsten Ueckerdt. Polychromatic colorings of geometric hypergraphs via shallow hitting sets, 2023. `arXiv:2310.19982`.

[118] Dömötör Pálvölgyi. Note on polychromatic coloring of hereditary hypergraph families, 2023. `arXiv:2309.04603`.

[119] Fulek Radoslav, Michael J Pelsmajer, Marcus Schaefer, and Daniel Štefankovič. Adjacent crossings do matter. *Journal of Graph Algorithms and Applications*, 16(3):759–782, 2012.

[120] Rajiv Raman and Saurabh Ray. Constructing planar support for non-piercing regions. *Discrete & Computational Geometry*, 64(3):1098–1122, Oct 2020.

[121] R. Bruce Richter and Carsten Thomassen. Intersections of curve systems and the crossing number of $C_5 \times C_5$. *Discrete Comput. Geom.*, 13(2):149–159, 1995.

[122] Vojtěch Rödl. Master's thesis, Charles University, 1973.

[123] Alexandre Rok and Bartosz Walczak. Outerstring graphs are χ-bounded. *SIAM Journal on Discrete Mathematics*, 33(4):2181–2199, 2019.

[124] Gelasio Salazar. On the intersections of systems of curves. *J. Comb. Theory, Ser. B*, 75(1):56–60, 1999.

[125] Alex Scott and Paul Seymour. A survey of χ-boundedness. *Journal of Graph Theory*, 95(3):473–504, 2020.

[126] Kim Seog-Jin, Alexandr Kostochka, and Kittikorn Nakprasit. On the Chromatic Number of Intersection Graphs of Convex Sets in the Plane. *Electron. J. Comb.*, 11(1):R52, 2004.

[127] Micha Sharir and Pankaj K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, USA, 2010.

[128] Shakhar Smorodinsky. On the chromatic number of geometric hypergraphs. *SIAM Journal on Discrete Mathematics*, 21(3):676–687, 2007.

[129] Shakhar Smorodinsky. Conflict-free coloring and its applications. In *Geometry–Intuitive, Discrete, and Convex*, pages 331–389. Springer, 2013.

[130] Shakhar Smorodinsky and Yelena Yuditsky. Polychromatic coloring for half-planes. *Journal of Combinatorial Theory, Series A*, 119(1):146–154, 2012.

[131] Jack Snoeyink and John Hershberger. Sweeping arrangements of curves. In *Discrete and computational geometry (New Brunswick, NJ, 1989/1990)*, volume 6 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 309–349. Amer. Math. Soc., Providence, RI, 1991.

[132] Joel Spencer, Endre Szemerédi, and William T. Trotter. *Unit distances in the Euclidean plane*, pages 294–304. Academic Press, United Kingdom, 1984.

[133] Raphael Steiner, Manfred Scheucher, Stefan Felsner, Pavel Valtr, Man-Kwun Chiu, and Patrick Schnider. On the average complexity of the k-level. *J. Comput. Geom.*, 11(1):493–506, 2020.

[134] Andrew Suk. *A Note on Geometric 3-Hypergraphs*, pages 489–498. Springer New York, New York, NY, 2013.

[135] László A. Székely. Crossing Numbers and Hard Erdős Problems in Discrete Geometry. *Comb. Probab. Comput.*, 6(3):353–358, 1997.

[136] Endre Szemerédi and William T. Trotter. Extremal problems in discrete geometry. *Combinatorica*, 3(3):381–392, Sep 1983.

[137] Gábor Tardos and Géza Tóth. Multiple coverings of the plane with triangles. *Discret. Comput. Geom.*, 38(2):443–450, 2007.

[138] William T. Tutte. Toward a theory of crossing numbers. *Journal of Combinatorial Theory*, 8(1):45–53, 1970.

[139] Gerhart Woeginger. Epsilon-nets for halfplanes. In J. van Leeuwen, editor, *Graph-Theoretic Concepts in Computer Science*, pages 243–252, Berlin, Heidelberg, 1989. Springer Berlin Heidelberg.

[140] Péter Ágoston, Gábor Damásdi, Balázs Keszegh, and Dömötör Pálvölgyi. Orientation of good covers, 2022. `arXiv:2206.01723`.