

Realizing an m -uniform four-chromatic hypergraph with disks

Gábor Damásdi
MTA-ELTE Lendület

Combinatorial Geometry Research Group

Dömötör Pálvölgyi
MTA-ELTE Lendület

Combinatorial Geometry Research Group

November 25, 2020

Abstract

We prove that for every m there is a finite point set \mathcal{P} in the plane such that no matter how \mathcal{P} is three-colored, there is always a disk containing exactly m points, all of the same color. This improves a result of Pach, Tardos and Tóth who proved the same for two colors. The main ingredient of the construction is a subconstruction whose points are in convex position. Namely, we show that for every m there is a finite point set \mathcal{P} in the plane in convex position such that no matter how \mathcal{P} is two-colored, there is always a disk containing exactly m points, all of the same color. We also prove that for unit disks no similar construction can work, and several other results.

1 Introduction

Coloring problems for hypergraphs defined by geometric range spaces have been studied a lot in different settings. A pair $(\mathcal{P}, \mathcal{S})$, where \mathcal{P} is a set of points in the plane and \mathcal{S} is a family of subsets of the plane (the *range space*), defines a (primal) hypergraph $\mathcal{H}(\mathcal{P}, \mathcal{S})$ whose vertex set is \mathcal{P} , and for each $S \in \mathcal{S}$ we add the edge $S \cap \mathcal{P}$ to the hypergraph. Given any hypergraph \mathcal{G} , a planar realization of \mathcal{G} is defined as a pair $(\mathcal{P}, \mathcal{S})$ for which $\mathcal{H}(\mathcal{P}, \mathcal{S})$ is isomorphic to \mathcal{G} . If \mathcal{G} can be realized with some pair $(\mathcal{P}, \mathcal{S})$ where \mathcal{S} is from some family \mathcal{F} , then we say that \mathcal{G} is realizable with \mathcal{F} .

A hypergraph is (properly) c -colorable if its vertices can be colored by c colors such that no edge is monochromatic. In this paper we focus on the c -colorability of hypergraphs realizable with disks. It is an easy consequence of the properties of Delaunay-triangulations and the Four Color Theorem that any hypergraph realizable with disks is four-colorable if every edge contains at least two vertices. Since K_4 is realizable with disks, this is sharp. But are less colors sufficient if all edges are required to contain at least m vertices for some large enough constant m ? Pach, Tardos and Tóth [21] have shown that two colors are not enough for any m , i.e., for any m , there exists an m -uniform hypergraph that is not two-colorable and that permits a planar realization with disks. Our main theorem is the following strengthening, which shows that three colors are also not enough, by realizing a four-chromatic hypergraph, completely resolving this question.

Theorem 1. *For any m , there exists an m -uniform hypergraph that is not three-colorable and that permits a planar realization with disks.*

The proof of Theorem 1 is based on two ideas. The first one is from [13], where colorings of so-called ABAB-free hypergraphs were considered. We will use a construction, based on one

from [13], to create a point set that is not two-colorable with respect to disks and, furthermore, the points are close to a prescribed set of points that lie on a circle. Then using ideas from [1] and [21], we will combine several copies of this non-two-colorable construction to create point sets that are not three-colorable. The non-two-colorable construction can be of independent interest.

Theorem 2. *For any m , there exists an m -uniform hypergraph that is not two-colorable and that permits a planar realization with disks that all contain some fixed point.*

Moreover, in the realization $(\mathcal{P}, \mathcal{D})$, the points \mathcal{P} can be placed arbitrarily close to some given points on a circle such that the boundary of each disk from \mathcal{D} is also arbitrarily close to this circle.

Previously, such a construction was only known for pseudo-disks containing a fixed point, and it was also shown that such hypergraphs are always three-colorable (already for $m = 2$) [1].

Note that if we require the points of \mathcal{P} to be placed close enough to the given points on the circle, then the points of \mathcal{P} will be in convex position.

The construction of Theorem 1 can be generalized as follows.

Theorem 3. *Let C be any closed convex set in the plane, that has two parallel supporting lines such that C is strictly convex in some neighbourhood of the two points of tangencies. For any m , there exists an m -uniform hypergraph that is not three-colorable and that permits a planar realization with homothets of C .*

To complement our results, we also show that no similar construction for unit disks exists.

Theorem 4. *For any k , any finite point set \mathcal{P} can be k -colored such that any unit disk, that contains some fixed point $o \notin \mathcal{P}$ and $8k - 7$ points from \mathcal{P} , will contain all k colors.*

Note that the condition $o \notin \mathcal{P}$ is in fact not required, we only state our results this way to emphasize that the common point o need not be from \mathcal{P} .

It is known that without requiring a common point o , the statement does not hold [19].

By using the well-known equivalence of the hypergraphs defined by primal and dual range spaces of the translates of any set [18, 20], we can conclude the following.

Corollary 5. *For any k , any finite collection of unit disks \mathcal{D} can be partitioned into k parts such that any point of a fixed unit disk $D_0 \notin \mathcal{D}$ that was covered by at least $8k - 7$ members of \mathcal{D} will be covered by all k parts.*

This statement is sharp in the sense that for a disk D_0 of larger radius it fails already for $k = 2$, even if the members of \mathcal{D} are required to be very close to each other; this follows from taking the dual of the construction from [19]. The function $8k - 7$ is unlikely to be sharp.

Theorem 4 can be extended in two different ways. Firstly, instead of stabbed unit disks we can consider stabbed translates of a given convex set.

Theorem 6. *For any plane convex set C and any integer k the following holds. Any finite point set \mathcal{P} can be k -colored such that any translate of C , that contains some fixed point $o \notin \mathcal{P}$ and $24k - 23$ points from \mathcal{P} , will contain all k colors.*

Secondly, it is also interesting to consider the case when our family contains stabbed homothetic copies of a given convex set. Theorem 2 implies that the natural analogue of Theorem 4 does not hold for homothetic copies of a disk. On the other hand we can show the following for polygons.

Theorem 7. *For any convex polygon P there is a t_P such that for any k the following holds. Any finite point set \mathcal{P} can be k -colored such that any homothetic copy of P , that contains some fixed point $o \notin \mathcal{P}$ and $t_P k$ points from \mathcal{P} , will contain all k colors.*

The rest of the paper is organized as follows. In Section 2 we prove Theorem 2, then building on this construction, in Section 3 we prove Theorem 1. In Section 4 we prove Theorem 4. The sketch of the proofs of the generalizations for other shapes (not disks), can be found in Section 5. Finally, in Section 6 we highlight the most important problems left open. We end this introduction with a bit more history, and a basic observation.

Related results

Pach proved in 1986 that any sufficiently thick covering of the plane [18] by the translates of a centrally symmetric open convex polygon can be partitioned into two disjoint coverings. The proof followed from showing that for every such polygon P there is an $m(P)$ for which any hypergraph, whose edges contain at least $m(P)$ vertices and can be realized with translates of P , is two-colorable. Several results followed, eventually showing that the similar statement is true for the translates of all convex polygons [7, 23], while counterexamples were given for non-convex polygons [21, 22] and convex shapes with a smooth boundary [19]. We know much less when instead of translates, homothetic (i.e., scaled and translated) copies are considered [2, 5, 10, 15] or about the problem of decomposing into multiple coverings/polychromatic colorings [4, 7, 24]. For a summary of the most closely related results, see Table 1; for more results, see the decade-old survey [20] or the webpage <https://coge.elte.hu/cogezoo.html>.

The above papers mainly focused on two- or polychromatic colorability. Proper three-colorability of geometric hypergraphs was studied in detail in [8] and [12]. In the latter paper it was shown that for every convex polygon P there is an $m(P)$ such that every $m(P)$ -uniform hypergraph realizable by homothetic copies of P is proper three-colorable. Our results imply that this is not the case for disks, disproving a conjecture from both of the above papers.

In this paper all disks are assumed to be open.¹ Let \mathcal{P} be a point set and let \mathcal{D} be a family of disks. An important folklore observation that we will use many times is that small perturbations of the points and the disks will not change the hypergraph $\mathcal{H}(\mathcal{P}, \mathcal{D})$. To put this into more precise terms, we will say that two points are ε -close if their distance is less than ε and two disks/circles are ε -close if their centers are ε -close and the difference of their radius is also smaller than ε .

Observation 8. *Suppose we have a finite point set \mathcal{P} and a finite set of open disks \mathcal{D} such that none of the points lie on the boundary of any of the disks. Then there is an ε such that replacing each disk with any ε -close disk and each point with any ε -close point will not change the hypergraph $\mathcal{H}(\mathcal{P}, \mathcal{D})$.*

¹But note that our results also hold for closed disks, as we could slightly shrink any finite system of open disks to contain the same points of a finite point set.

	translates	homothets
triangles	$\chi_{fat} = 2$ [25] $m_k = O(k)$ [7]	$\chi_{fat} = 2$ [9] $m_k = O(k^{4.09})$ [5] + [11]
convex polygons	$\chi_{fat} = 2$ [23] $m_k = O(k)$ [7]	$2 \leq \chi_{fat} \leq 3$ [12] $\chi_{fat} = 2$ for squares [2]
non-convex polygons*	$3 \leq \chi_{fat}$ [22]	$3 \leq \chi_{fat}$ [22]
stabbed convex polygons	$\chi_{fat} = 2$ [23] $m_k = O(k)$ [7]	$\chi_{fat} = 2$ [NEW] $m_k = O(k)$ [NEW]
stabbed disks	$\chi_{fat} = 2$ [NEW] $m_k = O(k)$ [NEW]	$\chi_{fat} \leq 3$ [1] $\chi_{fat} = 3$ [NEW]
disks	$3 \leq \chi_{fat} \leq 4$ [19]	$\chi_{fat} = 4$ [NEW]

Table 1: Summary of old and new results related to the topic of this paper. For a family, χ_{fat} denotes the smallest number k for which there is an m such that any finite set of points can be k -colored such that any member of the family with at least m points will contain at least two colors. For k colors, m_k denotes the smallest number m such that any finite set of points can be k -colored such that any member of the family with at least m points will contain all k colors. Results from this paper are marked with ‘[NEW]’.

* There are some very special non-convex polygons for which $\chi_{fat} = 2$ —for the complete classification, see [23] or [20].

2 A point set that is not two-colorable

Here we prove Theorem 2 by realizing for any m a non-two-colorable m -uniform hypergraph with disks that all contain some fixed point.

Our main lemma is the following. Combined with Observation 8, this gives us a way to perturb the points of \mathcal{P} with changing only a small part of the hypergraph $\mathcal{H}(\mathcal{P}, \mathcal{D})$.

Lemma 9. *If $\varepsilon > 0$, C is a circle, and a, b_1, \dots, b_t, c are points on C in this order, then there is a circle C' and points b'_1, \dots, b'_t on C' with the following properties.*

1. C' is ε -close to C .
2. b'_i is ε -close to b_i for each $i \in [t]$.
3. C' intersects C between a and b_1 , and between b_t and c .
4. Each b'_i is outside of C .

Proof. Choose points A and B between a and b_1 , and between b_t and c , respectively (see Figure 1). Choose O on the perpendicular bisector of AB , close to the center of C . C' will be the circle centered at O passing through A and B . Project b_1, \dots, b_t onto C' using O as center. If O is close enough to the center of C , this will clearly satisfy the requirements. \square




Figure 1: Lemma 9.

2.1 Hypergraphs based on rooted trees

The following hypergraph construction was used in [21] to create several counterexamples for coloring problems.

Definition 10. For any rooted tree T , let $\mathcal{H}(T)$ denote the hypergraph on vertex set $V(T)$, whose hyperedges are all sets of the following two types.

1. *Sibling hyperedges:* for each vertex $v \in V(T)$ that is not a leaf, take the set $S(v)$ of all children of v .
2. *Descendent hyperedges:* for each leaf $v \in V(T)$, take all vertices along the unique path $Q(v)$ from the root to v .

It is easy to see that $\mathcal{H}(T)$ is not two-colorable for any T . Either there is a monochromatic sibling edge, or we can follow the color of the root down to a leaf, finding a monochromatic descendent edge. We can create an m -uniform hypergraph by choosing T to be the complete m -ary tree of depth m . The non-two-colorable construction of Pach, Tardos and Tóth is also based on these hypergraphs.

Theorem 11 (Pach, Tardos and Tóth [21]). For any rooted tree T , the hypergraph $\mathcal{H}(T)$ permits a planar realization $(\mathcal{P}, \mathcal{D})$ with disks in general position such that every disk $D \in \mathcal{D}$ has a point on its boundary that does not belong to the closure of any other disk $D' \in \mathcal{D}$.

In order to be able to later build a point set that is not three-colorable, we will first extend Theorem 11 by showing that we can require the points to be close to a prescribed set of concyclic points, and require the disks to be close to the circle that contains the prescribed points. (We loose the property that every disk $D \in \mathcal{D}$ has a point on its boundary that does not belong to the closure of any other disk $D' \in \mathcal{D}$, so strict speaking Theorem 12 is not a generalization of Theorem 11.)

Theorem 12. If $\gamma > 0$, C is a circle and q_1, q_2, \dots, q_n are distinct points on C , then for any rooted tree T on n vertices, the hypergraph $\mathcal{H}(T)$ permits a planar realization $(\mathcal{P}, \mathcal{D})$ with disks such that

- (I) $\mathcal{P} = \{p_1, \dots, p_n\}$, and each p_i is γ -close to q_i .

(II) Every $D \in \mathcal{D}$ is γ -close to C .

An important property of rooted trees is that we can order their vertices in a special way. For a vertex v , let $Des(v)$ denote all descendants of v . Keszegh and Pálvölgyi [13] have shown that there is an ordering on the vertices of T such that

1. For each vertex $v \in V(T)$ the vertices in $S(v)$ are consecutive and they appear in the order later than v .
2. Furthermore, suppose $S(v) = \{r_1, \dots, r_k\}$ and they are in this order. Then the vertices $r_1, \dots, r_{k-1}, r_k, Des(r_k), Des(r_{k-1}), \dots, Des(r_1)$ are ordered like this, and the rest of the vertices of T are not in this interval. (The internal order of each $Des(r_i)$ is not specified by this statement.)

Call an order satisfying these properties a *siblings first order* [1] of T .

Proof of Theorem 12. In the planar realization of $\mathcal{H}(T)$, the vertices will correspond to the points q_i according to an (arbitrary) fix siblings first order. We start by showing that the sibling hyperedges can be easily realized and we only need to consider the descendent hyperedges.

Sibling hyperedges

From the properties of the siblings first order we know that for each v the vertices of $S(v)$ are consecutive, i.e., the points q_i corresponding to $S(v)$ are neighbours along the circle C . We apply Lemma 9 to find a disk that is γ -close to C , and contains exactly the points of $S(v)$. We can also ensure that no point q_i lies on the boundary of this disk. We repeat this for each $v \in V(\mathcal{H})$, until each sibling hyperedge is realized. Let \mathcal{D}_{Sib} denote the set of these disks. We apply Observation 8 to $(\{q_1, \dots, q_n\}, \mathcal{D}_{Sib})$ to get ε_{Sib} . That is, if each point p_i is ε_{Sib} -close to q_i , then $(\mathcal{P}, \mathcal{D}_{Sib})$ will still represent the sibling hyperedges. Therefore, it is enough to show that we can realize the descendent hyperedges for any γ .

Descendent hyperedges

It is useful to realize a slightly extended hypergraph. In $\mathcal{H}(T)$, we have a descendent hyperedge for each leaf. Now we will create a descendent hyperedge for non-leaf vertices too. So let $Q(v)$ denote the path from the root to v , and for each vertex v , we will realize the hyperedge $Q(v)$. The disk realizing $Q(v)$ will be denoted by $B(v)$. Let this extended hypergraph be denoted by $\mathcal{H}'(T)$. We will realize not only $\mathcal{H}(T)$, but $\mathcal{H}'(T)$. See Figure 2 for an example where we omitted the sibling edges.

We prove the existence of the required point set by describing an algorithm that produces a solution. Let $\mathcal{P} = \{p_1, p_2, \dots, p_n\}$ denote the vertices of T in a siblings first order. We will create the planar realization gradually, step by step. The algorithm starts by setting $p_i = q_i$ and in each step we will update the position of some of the p_i -s. That is, we identify the vertices of the hypergraph with planar points, and we will update the position of the vertices until they realize the hypergraph $\mathcal{H}'(T)$.

During the algorithm, we need to change the position of the points many times without altering the hyperedges that we have already realized. For this reason, we introduce a set of fixed points, \mathcal{P}_{fix} , and a set of disks, \mathcal{D}_{Des} , that corresponds to the descendent hyperedges. Once a point is in \mathcal{P}_{fix} we will not change its position any more. Every disk that we create




Figure 2: Disks realizing all paths from the root.

will be immediately added to \mathcal{D}_{Des} , and we will never change its position. The unfixed points, i.e., the points in $\mathcal{P} \setminus \mathcal{P}_{fix}$, will always be kept on the boundary of $\bigcup_{D \in \mathcal{D}_{Des}} D$. Furthermore, if we add $B(v)$ to \mathcal{D}_{Des} in the k -th step of the algorithm, $B(v) \cap \mathcal{P}$ will remain the same after we finished the k -th step, i.e., it will contain exactly the points of the path $Q(v)$ in its interior. Moreover, all descendants of v will be on the boundary of $B(v)$ after the k -th step, but later they will be moved off.

The structure of the algorithm is the following. We go through the vertices in the siblings first order and for each we do the following. By when we arrive at vertex p_k , the disk $B(p_k)$ representing the path $Q(p_k)$ from the root to p_k , will be already realized. For each child of p_k we will add a new disk, close to $B(p_k)$, that also realizes the same path from the root to p_k . Then we will move the points such that each new disk contains exactly one child of p_k . We have summarised the structure of the algorithm in the following pseudo code, while the phases of a step are depicted in Figure 3.

Algorithm 1: Structure of the algorithm

```

Set  $p_i = q_i$ ,  $\mathcal{P}_{fix} = \emptyset$ ,  $\mathcal{D}_{Des} = \emptyset$ .
Add the disk of  $C$ ,  $B(p_1)$ , to  $\mathcal{D}_{Des}$ .
Move  $p_1$  inside  $B(p_1)$  and add  $p_1$  to  $\mathcal{P}_{fix}$ .
for  $k = 1$  to  $n$  do
  for each child  $r_i$  of  $p_k$  do
    Add a disk  $B(r_i)$  representing  $Q(p_k)$  to  $\mathcal{D}_{Des}$ .          /* Using Lemma 9 */
    Move the required descendants of  $p_k$  to the boundary of  $B(r_i)$ .
  end
  for each child  $r_i$  of  $p_k$  do
    Move  $r_i$  inside  $B(r_i)$ .                                     /* now  $B(r_i)$  represents  $Q(r_i)$  */
    Add  $r_i$  to  $\mathcal{P}_{fix}$ .
  end
end

```

To be able to do these steps we also need a parameter δ that will ensure that the points

do not move too much and the disks are close to each other. There will be only three kinds of operations that we do during the algorithm.

- a) We update the position of a vertex by moving it at a distance less than the current value of δ . If it reached its final position, we add it to \mathcal{P}_{fix} .
- b) We add a new disk to \mathcal{D}_{Des} that is δ -close to one of the disks already in \mathcal{D}_{Des} .
- c) We decrease the value of δ .

We start with $\delta = \gamma/n^2$. (The reason for this is explained later.) Every time a disk is added or the position of a vertex is changed, we use Observation 8 for $(\mathcal{P}_{fix}, \mathcal{D}_{Des})$ to update δ to a smaller value, if needed. After any given update of δ , we only move points at distance less than δ . This ensures two things. Firstly, the points do not move far from their original position, and secondly, if we take a new disk that is δ -close to one of the disks, then it contains the same points of \mathcal{P}_{fix} .

The algorithm makes an initial adjustment on p_1 , and then there is one step (k -th step) for each point p_k . After each step the following properties will hold.

- (i) Each point p_i has either reached its final position or it lies on the boundary of the disk $B(w)$ where w is the lowest ancestor of p_i for which $B(w)$ is already defined. Also, in the second case p_i does not belong to the closure of any other disk $D' \in \mathcal{D}_{Des}$.
- (ii) Suppose p_j is the parent of p_i in T . Then in the j -th step the point p_i is added to \mathcal{P}_{fix} and the disk $B(p_i)$ is added to \mathcal{D}_{Des} .
- (iii) Each disk in \mathcal{D}_{Des} contains those points of P that correspond to the appropriate hyper-edge of $\mathcal{H}'(T)$.

During the initial adjustment we update p_1 to lie inside C but δ -close to q_1 . We add the disk corresponding to the circle C to \mathcal{D}_{Des} . We add p_1 to \mathcal{P}_{fix} and then we update δ by applying Observation 8 for $(\mathcal{P}_{fix}, \mathcal{D}_{Des})$. Clearly properties (i), (ii) and (iii) are satisfied.

In the k -th step we update the points in $Des(p_k)$. (If p_k is a leaf, we continue with the next step.) The process is depicted in Figure 3; we recommend the reader to consult it before reading on. From properties (i) and (ii) we know that at the start of the k -th step every point in $Des(p_k)$ lies on the boundary of $B(p_k)$, and they do not belong to any disk in \mathcal{D}_{Des} . Suppose $S(p_k) = \{r_1, \dots, r_l\}$. To maintain the three properties, we want to add the disks $B(r_1), \dots, B(r_l)$, and by the end of the k -th step we want to place the points of $Des(r_i)$ on the boundary of $B(r_i)$.

To achieve this, we apply Observation 8 and Lemma 9 for each child in the following way. First apply Observation 8 for the points in $Des(p_k)$ and disks in $\mathcal{D}_{Des} \setminus \{B(p_k)\}$. Since the points in $Des(p_k)$ do not belong to the boundary of any other disk this is possible. We update δ to the value of ε obtained from Observation 8 if $\varepsilon < \delta$.

Then we apply Lemma 9 for the boundary of the disk $B(p_k)$, such that the b_i -s are the points $\{r_1, \dots, r_l\} \cup Des(r_l) \cup \dots \cup Des(r_1)$ and ε is chosen to be the current value of δ . The points a and c have to be chosen carefully. We know that the points in $\{r_1, \dots, r_l\} \cup Des(r_l) \cup \dots \cup Des(r_1)$ lie on an arc of $B(p_k)$ that is not covered by any disk in \mathcal{D}_{Des} . We choose a and c on the two ends of this arc such that they are also not covered by any disk. Lemma 9 gives us a circle C' , this defines $B(r_1)$, which is added to \mathcal{D}_{Des} . The position of the points




Figure 3: Phases of the k -th step of the algorithm.

in $\{r_1, \dots, r_l\} \cup Des(r_l) \cup \dots \cup Des(r_1)$ is updated according to the result of Lemma 9. As usual, δ is also updated.

When we apply Observation 8 for the i -th time ($i > 1$) we apply it for the points $\{r_i, \dots, r_l\} \cup Des(r_l) \cup \dots \cup Des(r_i)$ and disks in $\mathcal{D}_{Des} \setminus \{B(r_{i-1})\}$.

When we apply Lemma 9 for the i -th time ($i > 1$), we apply it for the boundary of $B(r_{i-1})$, such that the b_i -s are the points in $\{r_i, \dots, r_l\} \cup Des(r_l) \cup \dots \cup Des(r_i)$ and ε is the current value of δ . The point a is chosen between r_{i-1} and r_i . The point c is chosen after the points of $\{r_i, \dots, r_l\} \cup Des(r_l) \cup \dots \cup Des(r_i)$ but before the points of $Des(r_{i-1}) \cup \dots \cup D(r_1)$. If $Des(r_{i-1}) \cup \dots \cup D(r_1)$ is empty, c is chosen before the arc of $B(r_{i-1})$ reaches any other disk. In the i -th case we get $B(r_i)$, which is added to \mathcal{D}_{Des} , and new positions for the points $\{r_i, \dots, r_l\} \cup Des(r_l) \cup \dots \cup Des(r_i)$ on $B(r_i)$. We update δ after each application of Lemma 9.

Finally, we finish the k -th step by moving r_1, r_2, \dots, r_l inside $B(r_1), B(r_2), \dots, B(r_l)$ respectively, but at most δ far. Since r_i was lying on the boundary of $B(r_1)$, we can also ensure that they are not moved into any other disk. Then we add r_1, \dots, r_l to \mathcal{P}_{fix} . We update δ again by applying Observation 8 to $(\mathcal{P}_{fix}, \mathcal{D}_{Des})$.

Let us see why properties (i), (ii), (iii) remain true during the run of the algorithm. Suppose they are true after the $(k-1)$ -th step.

The first part of property (i) and property (ii) are maintained since we have created the disks $B(r_1), \dots, B(r_l)$, and by the end of the k -th step the points of $Des(r_i)$ are on the boundary of $B(r_i)$. Points r_1, \dots, r_l were added to \mathcal{P}_{fix} .

The second part of property (i) could be violated in two ways. It could be that one of the new disks covers a point it should not. We have always chosen a and c such that this is avoided. The other possible violation is that we move a point into a disk. This is avoided, since if a point lies on the boundary of disk D before moving it, then we have updated δ for the disks in $\mathcal{D}_{Des} \setminus \{D\}$ right before moving the point. Also the movement is done by Lemma 9 so v cannot move into D .

As for property (iii), note that the only new disks in \mathcal{D}_{Des} are $B(r_1), \dots, B(r_l)$. Since (iii) was true before the step, $B(p_k)$ contains the vertices of $Q(p_k)$ and each of those are in

\mathcal{P}_{fix} . When we add $B(r_1)$, it is δ -close to $B(p_k)$, so it contains exactly the points in $Q(p_k)$. Similarly, when $B(r_i)$ is added, it is δ -close to $B(r_{i-1})$, so each of $B(r_1), \dots, B(r_l)$ contains the vertices of $Q(p_k)$. When finally we move the points r_1, r_2, \dots, r_l inside $B(r_1), B(r_2), \dots, B(r_l)$, respectively, we achieve that $B(r_i)$ contains the vertices of $Q(r_i)$.

We also need to check property (iii) for the disks that were already in \mathcal{D}_{Des} . Consider a disk $D \in \mathcal{D}_{Des}$. No point inside D was moved, since they are in \mathcal{P}_{fix} . The points in $\mathcal{P} \setminus \mathcal{P}_{fix}$ remain outside of D , since r_i only moves into one of the new disks and the rest of the points remain on the boundary of $\bigcup_{D \in \mathcal{D}_{Des}} D$. Hence property (iii) remains true.

Finally, we need to show that we get a solution for Theorem 1. Let \mathcal{D} contain the disks in \mathcal{D}_{Sib} and those disks in \mathcal{D}_{Des} that correspond to descendent edges that end at a leaf. Property (iii) and the argument for the sibling edges tells us that $(\mathcal{P}, \mathcal{D})$ is a planar representation of $\mathcal{H}(T)$.

For property (I), note that each point moves less than n^2 times since in the k -th step they move less than n times. We started with $\delta = \gamma/n^2$, so each point is at most $n^2 \frac{\gamma}{n^2} = \gamma$ far from their original position. The first disk was given by C and each disk was taken δ -close to a previous disk, in at most n steps reaching back to the first disk. Hence, all disks are γ/n close to C , implying (II). \square

3 A point set that is not three-colorable

Here we prove Theorem 1 by realizing for any m a non-three-colorable m -uniform hypergraph with disks.

We introduce a general operation for constructing hypergraphs that are not c -colorable, which is similar to prolonging $\mathcal{H}(T)$ with the children of a leaf. Essentially the same construction was used in [1].

Definition 13. *Suppose we have a hypergraph \mathcal{A} , a hypergraph \mathcal{B} and let F be an edge of \mathcal{A} . Then we define “ \mathcal{A} extended by \mathcal{B} through F ” as follows. Take $|V(\mathcal{B})|$ copies of the edge F and add a new vertex to each copy. Then we have $|V(\mathcal{B})|$ new vertices. Add a set of edges such that they form the hypergraph \mathcal{B} on these new vertices. The resulting hypergraph is denoted by $\mathcal{A} +_F \mathcal{B}$.*

Suppose \mathcal{A} is m -uniform and \mathcal{B} is $(m+1)$ -uniform. Then, if we extend \mathcal{A} by \mathcal{B} through each edge of \mathcal{A} , we get a hypergraph that is $(m+1)$ -uniform.

Claim 14. *Suppose \mathcal{A} is not c -colorable and \mathcal{B} is not $(c-1)$ -colorable. Then any extension of \mathcal{A} by \mathcal{B} is not c -colorable.*

Proof. Since any coloring of \mathcal{A} using c -colors has a monochromatic edge, and we have only “lost” the edge F , our only chance is to color the extended hypergraph such that the vertices of F are monochromatic. These are contained in each new edge, implying that none of the new vertices have the same color as the vertices of F . But then the copy of \mathcal{B} is $(c-1)$ -colored, and thus one of its edges is monochromatic. \square

Let \mathcal{G}_i denote the hypergraph that has i vertices and only one edge that contains all the vertices. Clearly, \mathcal{G}_1 is not c -colorable for any c and \mathcal{G}_i is not 1-colorable. Hence, we can build non- c -colorable hypergraphs starting from these trivial ones and using them in the extensions.

Observation 15. For any rooted tree T the hypergraph $\mathcal{H}(T)$ can be built with a series of extensions starting from \mathcal{G}_1 , where each extending hypergraph is one of the \mathcal{G}_i -s.

Now we are ready to define the non-three-colorable hypergraphs that we will realize with points and disks. For each m let $\mathcal{H}^2(m) = \mathcal{H}(T)$ where T is the m -ary tree of depth m . Clearly, $\mathcal{H}^2(m)$ is m -uniform and not two-colorable. We define non-three-colorable m -uniform hypergraphs $\mathcal{H}^3(m)$ based on them inductively.

First we create a sequence of hypergraphs. Let $\mathcal{F}_1(m) = \mathcal{G}_1$ and let v denote the single vertex of it. For $i > 1$, let $\mathcal{F}_i(m)$ be the hypergraph that we get if we extend each edge of $\mathcal{F}_{i-1}(m)$ that contains v by $\mathcal{H}^2(m)$. Note that $\mathcal{F}_i(m)$ has only two types of edges. The first type are the ones that contain v , each of these contains exactly i vertices. (They are like the descendent edges.) The second type are those that were added in a copy of $\mathcal{H}^2(m)$, these contain exactly m vertices. (They are a bit like the sibling edges.) Therefore, $\mathcal{F}_m(m)$ is m -uniform. Also, by Claim 14 and induction, no $\mathcal{F}_i(m)$ is three-colorable.

Let $\mathcal{H}^3(m) = \mathcal{F}_m(m)$. For example, $\mathcal{F}_1(2)$ is \mathcal{G}_1 and $\mathcal{H}^2(2)$ is a triangle graph. Extending \mathcal{G}_1 through its single edge by a triangle gives us the complete graph on 4 vertices, so $\mathcal{H}^3(2)$ is just K_4 .

3.1 Realizing $\mathcal{H}^3(m)$

Since $\mathcal{H}^3(m)$ was built by a series of extensions, it is enough to show that we can do each extension geometrically. This step is essentially the same as in [21].

Lemma 16. Suppose a hypergraph \mathcal{A} is realized with $(\mathcal{P}, \mathcal{D})$ such that every disk $D \in \mathcal{B} \subset \mathcal{D}$ has a point on its boundary that does not belong to the closure of any other disk $D' \in \mathcal{D}$. Then we can also realize $\mathcal{A} +_F \mathcal{H}(T)$ for any rooted tree T and any edge $F \in \mathcal{B}$, with a pair $(\mathcal{P}', \mathcal{D}')$, such that any disk $D \in \mathcal{B}$ and every new copy of F has a point on its boundary that does not belong to the closure of any other disk $D' \in \mathcal{D}'$.




Figure 4: Realizing an extension.

Proof. Suppose $D_F \in \mathcal{D}$ is the disk realizing the edge F . Then D_F has a point p on its boundary that does not belong to the closure of any other disk $D' \in \mathcal{D}$. Take a small circle C that is tangent to D at p . If we chose the radius of C to be small enough, then C will not intersect any of the disks. Now take $n = |V(T)|$ copies of D_F and rotate them slightly around the center of C . If all the rotations are small, the resulting disks will be ε -close to D_F and by Observation 8 they will contain the same points as D_F . Also, if the angles of the rotations are

different, then each new disk has a point on its boundary that does not belong to the closure of any other disk.

Then we enlarge each copy of D_F slightly, so that they intersect C but some part of their boundary remains uncovered by other disks. Place the points q_1, \dots, q_n in the intersections and use Theorem 12 to realize $\mathcal{H}(T)$. Since each disk in the realization of $\mathcal{H}(T)$ is close to C , they do not contain any point of \mathcal{P} . \square

4 Stabbed unit disks

Here we prove Theorem 4, that given a fixed origin o , any point set \mathcal{P} can be k -colored such that every unit disk that contains the origin and at least $8k - 7$ points of \mathcal{P} , contains all k colors.

We call unit disks containing o *stabbed* unit disks. Take a representative disk for every $\mathcal{P}' \subset \mathcal{P}$ that can be obtained as the intersection of a stabbed unit disk and \mathcal{P} . Since \mathcal{P} is finite, the collection of these representative disks, \mathcal{D} , is also finite. We divide the plane into four quarters by two perpendicular lines through o .

Lemma 17. *The boundaries of two stabbed unit disks intersect at most once in each quarter.*

(See Figure 5.)




Figure 5: From stabbed disks to pseudolines.

Proof. Denote the two intersection points of the boundary of the two disks, D_1 and D_2 , by x_1 and x_2 . Since D_1 and D_2 have equal radii, the length of the arc they contain from each other is less than half of their perimeter. Thus, using Thales's theorem, for any inner point p of $D_1 \cap D_2$ the angle $x_1 p x_2$ is at least 90° . This proves the statement as $o \in D_1 \cap D_2$. \square

This means that in each quarter the boundaries of the disks behave as pseudolines. Let us partition our point set \mathcal{P} into four parts, $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3, \mathcal{P}_4$, depending on which quarter the points are in. We will color each \mathcal{P}_i separately.

Denote the quarter that contains the points of \mathcal{P}_i by Q_i . Prolong the section of the boundaries falling into this quarter, i.e., $D \cap Q_i$ for each $D \in \mathcal{D}$, such that they do not intersect outside Q_i and all of them stretches from “left infinity” to “right infinity”. Without loss of generality, we can suppose that o is above all these curves. The (upward) regions

that are bounded by such curves form (upward) *pseudohalfplanes*. Generalizing a theorem of Smorodinsky and Yuditsky [24], the following polychromatic coloring result is known about them.

Theorem 18 (Keszegh-Pálvölgyi [13]). *Given a finite collection of points and pseudohalfplanes, the points can be k -colored such that every pseudohalfplane that contains at least $2k - 1$ points will contain all k colors.*

We can apply this theorem to \mathcal{P}_i (ignoring $\mathcal{P} \setminus \mathcal{P}_i$) and our regions. From this we get that any stabbed disk that contains at least $2k - 1$ points of \mathcal{P} from the quarter Q_i will contain all k colors. Repeating this for all four \mathcal{P}_i , we obtain by the pigeonhole principle that any stabbed disk containing at least $4(2k - 2) + 1 = 8k - 7$ points of \mathcal{P} contains all k colors.

For $k = 2$, this gives that 9 points per disk are enough, though this is probably far from being tight. From below, we could only find a construction showing that 3 points per disk are not enough.

5 Generalizations to other shapes and connection to clustering

Theorem 3: Generalizing Theorem 1

Let $(\mathcal{P}, \mathcal{S})$ be a range space where \mathcal{S} contains homothets of a given convex sets C . Theorem 1 shows that if C is a disk, then the hypergraph $\mathcal{H}(\mathcal{P}, \mathcal{S})$ might not be three-colorable. For which other C convex sets can we generalize this statement? Theorem 3 shows that the answer to this question is a surprisingly wide family. We omit the detailed proof of Theorem 3, the main ideas are the following.

It is not hard to see, that to achieve an appropriate version of Theorem 9, it is enough to assume that C is closed and strictly convex. Furthermore, since we can restrict ourselves to an arbitrarily small portion of the boundary of C , the non-two-colorable construction of Theorem 12 can be achieved by any closed C that is strictly convex in the neighbourhood of a point on the boundary.

The geometric extension in Lemma 16 requires a bit more. Let C be any closed convex set in the plane, that has two parallel supporting lines such that C is strictly convex in some neighbourhood of the two points of tangencies. When we extend through an edge, we can use one of these neighbourhoods to correctly place the new points and the other neighbourhood to create the extending hypergraph on the new points.

Theorems 6 and 7: Generalizing Theorem 4

The proofs of Theorems 6 and 7 follow the same steps as of Theorem 4. We divide the plane into regions using rays starting at o and then we show that in each region the boundaries of the sets behave as pseudolines. Here we provide the sketch of the proof of Theorem 6, but leave the details to the interested reader.

Sketch of the proof of Theorem 6. The proof of Theorem 4 relied on the fact that if p is a point in the intersection of two unit disks, then it sees the two intersection points x_1, x_2 in an angle that is at least $\frac{\pi}{2}$. For a general convex set C this angle might be arbitrarily small if the set is thin enough.

To avoid this problem, we can apply an affine transformation to C to make its width more even from different directions. An affine transformation does not change the conclusion of




Figure 6: Angle x_1px_2 is at least $\frac{\pi}{6}$.

Theorem 6, so it is indeed enough to argue for the affine copy. One way to achieve this is via the John ellipsoid of the set, i.e., the largest volume ellipsoid contained in the set. Take the affine transformation that changes the John ellipsoid to be the unit circle. It is well known that in this case the set is contained in the circle whose radius is 2 and shares the center with the John ellipsoid.

Just like in the proof of Lemma 17, we can show that the angle x_1px_2 is at least $\frac{\pi}{6}$. Let C_1 and C_2 be two translates of C . C_1 and C_2 have two common tangent lines, let t_1, t_2 be the tangent points to these lines on C_2 . Let t'_1, t'_2 be the projections of t_1, t_2 to the circle of radius 2 around C_2 . Now x_1px_2 is bigger than angle t_1xt_2 . Clearly the distance of t'_1 and t'_2 is at least 2, since they are separated by the common tangent lines. Hence, $t'_1xt'_2$ is at least as big as the inscribed angle corresponding to a chord of length 2, which is $\frac{\pi}{6}$. Therefore, we can divide the plane into 12 regions and in each region the boundaries of the sets behave as pseudolines and from here the proof is the same, giving $12(2k - 2) + 1 = 12k - 23$. \square

We could not rule out the possibility that with a suitable affine transformation the x_1px_2 angle would be always at least $\frac{\pi}{2}$, just like for disks. This is a nice geometric question.

Proof of Theorem 7. Let P be a fixed polygon and let v_1, v_2, \dots, v_n be the vertices of P . Let $\alpha_{ij} = \inf\{\angle v_i v_j \mid i \in \{1, \dots, n\}, j \in P \setminus \overline{v_i v_j}\}$. It is not hard to see that this angle α_{ij} is minimized for a vertex v . This also implies $\alpha_{ij} > 0$. Let $\alpha = \min_{i,j} \alpha_{ij}$.

Divide the plane into regions around o such that each region is a cone whose angle is less than α . A homothet of P that contains o can have at most one vertex in each region, otherwise they would be visible from an angle α from o . Let P' and P'' be homothetic copies of P that contain o . We will show that P' and P'' intersects at most once in each region.

Suppose there is a region R where they intersect twice. The boundary of R intersects at most two sides of each polygon, since they are convex and contain o . There are essentially two ways P' and P'' can intersect (see Figure 7). Either the two intersection points of P' and P'' fall on the same side of one of the polygons, or not.

We will show that these are not possible for two homothets. In the first case let s be the side that contains the two intersections. In the second case let s be one of the sides that ends at the vertex in R that is closer to o . Let n be the normal vector of the side s and let us look at the extremal points of P' and P'' in the direction of n . Since they are homothets of each




Figure 7: The two ways P' and P'' can intersect each other twice in a region.

other, the extremal points should form a side of the polygon in both P' and P'' . But for one of them we get only a vertex as an extremal point. Therefore two homothets meet at most once in each region. The rest of the proof follows the same way as for Theorem 4. \square

Connection to clustering for planar graphs

We say that a graph G is c -colorable with clustering m if each vertex can be assigned one of c colors so that each monochromatic component has at most m vertices. For example, planar graphs are four-colorable with clustering 1 due to the Four Color Theorem. Clustering questions gained popularity recently, see [3, 6, 16, 17]. It has been shown [14] that for each m there exists a planar graph that is not three-colorable with clustering m . We reprove this result using Theorem 1.

Theorem 19 (Kleinberg et al. [14]). *For each $m > 0$ there exists a planar graph that is not three-colorable with clustering m .*

Proof. Take the point set \mathcal{P} that realizes the non-three-colorable hypergraph $\mathcal{H}^3(m)$, and perturb the points so that no four lie on a circle. Consider the Delaunay graph of this point set, that is, we connect two points if they can be separated from the rest with a circle. From Theorem 1 we know that any three-coloring of this graph contains m monochromatic points that can be separated from the rest by a circle and it is well-known that m such points always form a connected component in the Delaunay graph, thus a monochromatic cluster in our case. \square

6 Open questions

One of the most interesting questions left open is about unit disks. Is there an m such that every point set in the plane can be three-colored such that every unit disk containing exactly m points contains at least two colors?

In fact, this conjecture has a strengthening in the dual, cover-decomposition problem. Is there an m such that every m -fold covering of \mathcal{P} by disks can be partitioned into three parts such that any two parts cover \mathcal{P} ?

If there is a counterexample to these questions, then it has to be quite different from ours, as already the two-color version [19] needed the realization of a different hypergraph.

Coloring questions are also interesting in higher dimensions. Suppose we have a set of points and a family of balls in \mathbb{R}^d . How many colors are needed to color the points such that none of the balls that contain at least m points are monochromatic? Theorem 1 can be extended to show that in some cases we need $d+2$ colors. On the other hand, no upper bound is known already for $d = 3$ that would not depend on m .

Acknowledgment

We would like to thank Balázs Keszegh for useful discussions and for reading the draft of this manuscript and also our anonymous reviewers for their valuable suggestions.

References

- [1] Eyal Ackerman, Balázs Keszegh, and Dömötör Pálvölgyi. Coloring hypergraphs defined by stabbed pseudo-disks and *ABAB*-free hypergraphs. *Acta Math. Univ. Comenian. (N.S.)*, 88(3):363–370, 2019.
- [2] Eyal Ackerman, Balázs Keszegh, and Máté Vizer. Coloring points with respect to squares. *Discrete & Computational Geometry*, 58(4):757–784, Dec 2017.
- [3] Noga Alon, Guoli Ding, Bogdan Oporowski, and Dirk Vertigan. Partitioning into graphs with only small components. *J. Combin. Theory Ser. B*, 87(2):231–243, 2003.
- [4] 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 *Algorithms and data structures*, volume 8037 of *Lecture Notes in Comput. Sci.*, pages 73–84. Springer, Heidelberg, 2013.
- [5] Jean Cardinal, Kolja Knauer, Piotr Micek, and Torsten Ueckerdt. Making triangles colorful. *J. Comput. Geom.*, 4(1):240–246, 2013.
- [6] Louis Esperet and Gwenaël Joret. Colouring planar graphs with three colours and no large monochromatic components. *Combin. Probab. Comput.*, 23(4):551–570, 2014.
- [7] Matt Gibson and Kasturi Varadarajan. Optimally decomposing coverings with translates of a convex polygon. *Discrete Comput. Geom.*, 46(2):313–333, 2011.
- [8] Balázs Keszegh. Coloring half-planes and bottomless rectangles. *Computational geometry*, 45(9):495–507, 2012.
- [9] Balázs Keszegh and Dömötör Pálvölgyi. Octants are cover decomposable. *Discrete Comput. Geom.*, 47:598–609, 2012.
- [10] Balázs Keszegh and Dömötör Pálvölgyi. Convex polygons are self-coverable. *Discrete Comput. Geom.*, 51(4):885–895, 2014.
- [11] Balázs Keszegh and Dömötör Pálvölgyi. More on decomposing coverings by octants. *J. Comput. Geom.*, 6(1):300–315, 2019.
- [12] Balázs Keszegh and Dömötör Pálvölgyi. Proper coloring of geometric hypergraphs. In *Symposium on Computational Geometry*, volume 77 of *LIPICs*, pages 47:1–47:15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.
- [13] 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):1–26, 2019.

- [14] Jon M. Kleinberg, Rajeev Motwani, Prabhakar Raghavan, and Suresh Venkatasubramanian. Storage management for evolving databases. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*, pages 353–362, 1997.
- [15] István Kovács. Indecomposable coverings with homothetic polygons. *Discrete & Computational Geometry*, 53(4):817–824, 2015.
- [16] Nathan Linial, Jiří Matoušek, Or Sheffet, and Gábor Tardos. Graph coloring with no large monochromatic components. *Electronic Notes in Discrete Mathematics*, 29:115 – 122, 2007. European Conference on Combinatorics, Graph Theory and Applications.
- [17] Sergey Norin, Alex Scott, Paul D. Seymour, and David R. Wood. Clustered colouring in minor-closed classes. *Combinatorica*, 39(6):1387–1412, 2019.
- [18] János Pach. Covering the plane with convex polygons. *Discrete Comput. Geom.*, 1(1):73–81, December 1986.
- [19] János Pach and Dömötör Pálvölgyi. Unsplittable coverings in the plane. *Advances in Mathematics*, 302:433–457, 2016.
- [20] János Pach, Dömötör Pálvölgyi, and Géza Tóth. Survey on decomposition of multiple coverings. In *Geometry–Intuitive, Discrete, and Convex*, pages 219–257. Springer, 2013.
- [21] János Pach, Gábor Tardos, and Géza Tóth. Indecomposable coverings. In *Discrete geometry, combinatorics and graph theory*, volume 4381 of *Lecture Notes in Comput. Sci.*, pages 135–148. Springer, Berlin, 2007.
- [22] Dömötör Pálvölgyi. Indecomposable coverings with concave polygons. *Discrete Comput. Geom.*, 44(3):577–588, 2010.
- [23] Dömötör Pálvölgyi and Géza Tóth. Convex polygons are cover-decomposable. *Discrete & Computational Geometry*, 43(3):483–496, Apr 2010.
- [24] Shakhar Smorodinsky and Yelena Yuditsky. Polychromatic coloring for half-planes. *J. Combin. Theory Ser. A*, 119(1):146–154, 2012.
- [25] Gábor Tardos and Géza Tóth. Multiple coverings of the plane with triangles. *Discrete & Computational Geometry*, 38:443–450, 2007.