Unsplittable coverings in the plane

János Pach[†]   Dömötör Pálvölgyi[‡]

November 5, 2018

Et tu mi fili, Brute?[*]
(Julius Caesar)

Abstract

A system of sets forms an *m-fold covering* of a set X if every point of X belongs to at least m of its members. A 1-fold covering is called a *covering*. The problem of splitting multiple coverings into several coverings was motivated by classical density estimates for *sphere packings* as well as by the *planar sensor cover problem*. It has been the prevailing conjecture for 35 years (settled in many special cases) that for every plane convex body C, there exists a constant $m = m(C)$ such that every m-fold covering of the plane with translates of C splits into 2 coverings. In the present paper, it is proved that this conjecture is false for the unit disk. The proof can be generalized to construct, for every m, an unsplittable m-fold covering of the plane with translates of any open convex body C which has a smooth boundary with everywhere *positive curvature*. Somewhat surprisingly, *unbounded* open convex sets C do not misbehave, they satisfy the conjecture: every 3-fold covering of any region of the plane by translates of such a set C splits into two coverings. To establish this result, we prove a general coloring theorem for hypergraphs of a special type: *shift-chains*. We also show that there is a constant $c > 0$ such that, for any positive integer m, every m-fold covering of a region with unit disks splits into two coverings, provided that every point is covered by *at most $c2^{m/2}$* sets.

1 Introduction

Let $\mathcal{C}$ be a family of sets in $\mathbb{R}^d$, and let $P \subseteq \mathbb{R}^d$. We say that $\mathcal{C}$ is an *m-fold covering of P* if every point of P belongs to at least m members of $\mathcal{C}$. A 1-fold covering is called a *covering*. Clearly, the union of m coverings is an m-fold covering. We will be mostly interested in the case when P is a large region or the whole space $\mathbb{R}^d$.

Sphere packings and coverings have been studied for centuries, partially because of their applications in crystallography, Diophantine approximation, number theory, and elsewhere. The research in this field has been dominated by density questions of the following type: What is the most "economical" (i.e., least dense) m-fold covering of space by unit balls or by translates of a fixed convex body? It is suggested by many classical results and physical observations that, at least in low-dimensional spaces, the optimal arrangements are typically periodic, and they can be split into several lattice-like coverings [14, 15]. Does a similar phenomenon hold for all sufficiently "thick" multiple coverings, without any assumption on their densities?

About 15 years ago, a similar problem was raised for *large scale ad hoc sensor networks*; see Feige *et al.* [13], Buchsbaum *et al.* [6]. In the – by now rather extensive – literature, it is usually referred to as the *sensor cover problem*. In its simplest version it can be phrased as follows. Suppose that a large region P is monitored by a set of sensors, each having a circular range of unit radius and each powered by a battery of unit lifetime. Suppose that every point of P is within the range of at least m sensors, that is, the family of ranges of the sensors, $\mathcal{C}$, forms an m-fold covering of P. If $\mathcal{C}$ can be split into k coverings $\mathcal{C}_1, \ldots, \mathcal{C}_k$, then the region can be monitored by the sensors for at least k units of time. Indeed, at time i, we can switch on all sensors whose ranges belong to $\mathcal{C}_i$ $(1 \le i \le k)$. We want to maximize k, in order to guarantee the longest possible service. Of course, the first question is the following.

Problem 1 (Pach, 1980 [31])**.** *Is it true that every m-fold covering of the plane with unit disks splits into two coverings, provided that m is sufficiently large?*

In a long unpublished manuscript, Mani and Pach [27] claimed that the answer to this question was in the affirmative with $m \le 33$. Pach [35] warned that this "has never been independently verified." Winkler [42] even conjectured that the statement is true with $m = 4$. For more than 30 years, the prevailing conjecture has been that for any open plane *convex body* (i.e., bounded convex set) C, there exists a positive integer $m = m(C)$ such that every m-fold covering of the plane with translates of C splits into two coverings. This conjecture was proved in [32] for centrally symmetric convex polygons C. It took almost 25 years to generalize this statement to all convex polygons [40, 38]. Moreover, it was proved by Aloupis *et al.* [3] and Gibson and Varadarajan [19] that in these cases, for every integer k, every at least bk-fold covering splits into k coverings, where $b = b(C)$ is a suitable positive constant. See [33, 36, 34], for surveys.

Here we disprove the above conjecture by giving a negative answer to Problem 1.

Theorem 2. *For every positive integer m, there exists an m-fold covering of the plane with open unit disks that cannot be split into 2 coverings.*

Our construction can be generalized as follows.
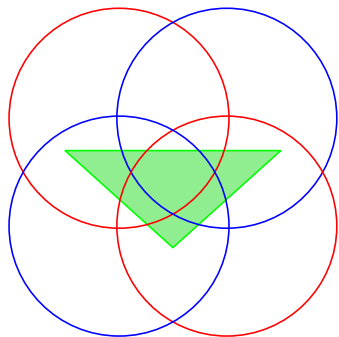
(a) The disks form a 2-fold covering of the green triangle. The colors give a split into 2 coverings. It follows, however, from the proof of Theorem 2, that for any m there is an m-fold covering of a triangle by disks that does not split.
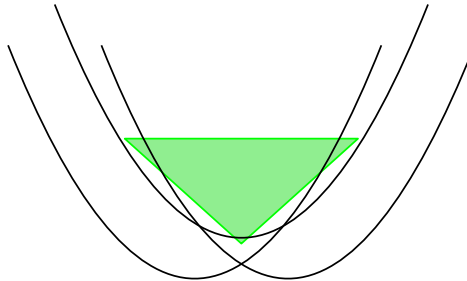
(b) The parabolas form a 2-fold covering of the green triangle, but no matter how we 2-color them, there will be a point not covered by one of the colors. It follows, however, from Theorem 6 and a standard compactness argument, that any 3-fold covering of a closed triangle by the translates of an open parabola splits into 2 coverings.

Figure 1: Two simple examples.

Theorem 3. *Let C be any open plane convex set, which has two parallel supporting lines with positive curvature at their points of tangencies. Then, for every positive integer m, there exists an m-fold covering of the plane with translates of C that cannot be split into 2 coverings.*

As was mentioned above, for every open convex polygon Q, there exists a smallest positive integer $m(Q)$ such that every $m(Q)$-fold covering of the plane with translates of Q splits into 2 coverings. We have that $\sup m(Q) = \infty$, where the sup is taken over all convex polygons Q. Otherwise, we could approximate the unit disk with convex n-gons with n tending to infinity. By compactness, we would conclude that the unit disk C satisfies $m(C) < +\infty$, which contradicts Theorem 2.

Problem 4. *Does there exist, for any $n > 3$, an integer $m(n)$ such that every convex n-gon Q satisfies $m(Q) \le m(n)$?*

For any triangle T, there is an affine transformation of the plane that takes it into an equilateral triangle T_0. Therefore, we have $m(T) = m(T_0)$ and $m(3)$ is finite. For $n = 4$, Problem 4 is open.

In spite of our sobering negative answer to Problem 1 and its analogues in higher dimensions (cp. [27]), there are important classes of multiple coverings such that all of their members are splittable. According to our next, somewhat counter-intuitive result, for example, any m-fold covering of $\mathbb{R}^d$ with unit balls can be split into 2 coverings, provided that no point of the space is covered by too many balls. (We could innocently believe that heavily covered points make it only easier to split an arrangement.)

3

Theorem 5. *For every $d \geq 2$, there exists a positive constant c_d with the following property. For every positive integer m, any m-fold covering of $\mathbb{R}^d$ with unit balls can be split into two coverings, provided that no point of the space belongs to more than $c_d 2^{m/d}$ balls.*

Theorem 5 was one of the first geometric applications of the Lovász local lemma [10], and it was included in [2]. Here, we establish a more general statement (see Theorem 8.2).

One may also believe that *unbounded* convex sets behave even worse than the bounded ones. It turns out, however, that this is not the case.

Theorem 6. *Let C be an unbounded open convex set and let P be a finite set of points in the plane. Then every 3-fold covering of $P \subset \mathbb{R}^2$ with translates of C can be split into two coverings of P.*

In fact, using a standard compactness argument, Theorem 6 also holds if P is any *compact* set in the plane. However, Theorem 6 does not generalize to higher dimensions. Indeed, it follows from the proof of Theorem 2 that, for every positive integer m, there exists a finite family $\mathcal{C}$ of open unit disks in the plane and a finite set $P \subset \mathbb{R}^2$ such that $\mathcal{C}$ is an m-fold covering of P that cannot be split into two coverings. Consider now an unbounded convex cone C' in $\mathbb{R}^3$, whose intersection with the plane $\mathbb{R}^2$ is an open disk. Take a system of translates of C' such that their intersections with the plane coincide with the members of $\mathcal{C}$. These cones form an m-fold covering of P that cannot be split into two coverings.

For interesting technical reasons, the proof of Theorem 6 becomes much easier if we restrict our attention to multiple coverings of the *whole plane*. In fact, in this case, we do not even have to consider *multiple* coverings! Moreover, the statement remains true in higher dimensions.

Proposition 7. *Let C be an unbounded line-free open convex set in $\mathbb{R}^d$. Then every covering of $\mathbb{R}^d$ with translates of C can be split into two, and hence into infinitely many, coverings.*

The reason why we assume here that C is *line-free* (i.e., does not contain a full line) is the following. If C contains a straight line, then it can be obtained as the direct product of a line l and a $(d-1)$-dimensional open convex set C'. Any arrangement $\mathcal{C}$ of translates of C in $\mathbb{R}^d$ is combinatorially equivalent to the $(d-1)$-dimensional arrangement of translates of C', obtained by cutting $\mathcal{C}$ with a hyperplane orthogonal to l. In particular, the problem whether an m-fold covering of $\mathbb{R}^d$ with translates of C can be split into two coverings reduces to the respective question about m-fold coverings of $\mathbb{R}^{d-1}$ with translates of C'.

Proposition 7 is false already in the plane without the assumption that C is open. However, every 2-fold covering of the plane with translates of an unbounded C can be split into two coverings. We omit the proof as it reduces to a simple claim about intervals.

However, in higher dimensions, the similar claim is false.

Theorem 8. *There is a bounded convex set $C' \subset \mathbb{R}^3$ with the following property. One can construct a family of translates of $C = C' \times [0, \infty) \subset \mathbb{R}^4$ which covers every point of $\mathbb{R}^4$ infinitely many times, but which cannot be split into two coverings.*

The construction given in Section 7 is based on an example of Naszódi and Taschuk [30], and explores the fact that the boundary of C' can be rather "erratic." We do not

know whether sufficiently thick coverings of $\mathbb{R}^3$ by translates of an unbounded line-free convex set can be split into two coverings or not.

In the sequel, we will study the equivalent *"dual" form* of the above questions. Consider a family $\mathcal{C} = \{C_i : i \in I\}$ of translates of a set $C \subset \mathbb{R}^d$ that form an m-fold covering of $P \subseteq \mathbb{R}^d$. Suppose without loss of generality that C contains the origin 0. For every $i \in I$, let c_i denote the point of C_i that corresponds to $0 \in C$. In other words, we have $\mathcal{C} = \{C + c_i : i \in I\}$. Assign to each $p \in P$ a translate of $-C$, the reflection of C about the origin, by setting $C_p^* = -C + p$. Observe that

$$p \in C_i \iff c_i \in C_p^*.$$

In particular, the fact that $\mathcal{C}$ forms an m-fold covering of P is equivalent to the following property: Every member of the family $\mathcal{C}^* = \{C_p^* : p \in P\}$ contains at least m elements of $\{c_i : i \in I\}$. Thus, Theorem 2 can be rephrased in the following *dual* form.

Theorem 2'. *For every $m \geq 2$, there is a set of points $P^* = P^*(m)$ in the plane with the property that every open unit disk contains at least m elements of P^*, and no matter how we color the elements of P^* with two colors, there exists a unit disk such that all points in it are of the same color.*

A set system *not* satisfying this condition is said to have *property B* (in honor of Bernstein) or is 2-*colorable* (see [29, 9, 39]). Generalizations of this notion are related to *conflict-free colorings* [12] and have strong connections, e.g., to the theory of ε-*nets*, *geometric set covers* and to *combinatorial game theory* [21, 34, 1, 41, 18].

The rest of this paper is organized as follows. In the next three sections, we prove Theorem 2' in 3 steps. In Section 2, we exhibit a family of non-2-colorable m-uniform hypergraphs $\mathcal{H}(k,l)$. In Section 3, we construct planar "realizations" of these hypergraphs, where the vertices correspond to points and the (hyper)edges to unit disks, preserving the incidence relations. In Section 4, we extend this construction, without violating the colorability condition, so that every disk contains at least m points. In Section 5, we modify these steps in order to establish Theorem 3, a generalization of Theorem 2 to bounded plane convex bodies with a smooth boundary. Sections 6 and 7 contain the proofs of our results related to multiple coverings with *unbounded* convex sets: Theorem 6, Proposition 7, and Theorem 8. The proof of a more general version of Theorem 5, using the Lovász local lemma, can be found in Section 8. Finally, in Section 9 we make some concluding remarks and mention a couple of open problems.

2 A family of non-2-colorable hypergraphs $\mathcal{H}(k,l)$

In this section we define, for any positive integers k and l, an abstract hypergraph $\mathcal{H}(k,l)$ with vertex set $V(k,l)$ and edge set $E(k,l)$. The hypergraphs $\mathcal{H}(k,l)$ are defined recursively. The edge set $E(k,l)$ will be the disjoint union of two sets, $E(k,l) = E_R(k,l) \cup E_B(k,l)$, where the subscripts R and B stand for red and blue. All edges belonging to $E_R(k,l)$ will be of size k, all edges belonging to $E_B(k,l)$ will be of size l. In other words, $\mathcal{H}(k,l)$ is the union of a k-*uniform* and an l-*uniform* hypergraph. If $k = l = m$, we get an m-uniform hypergraph.

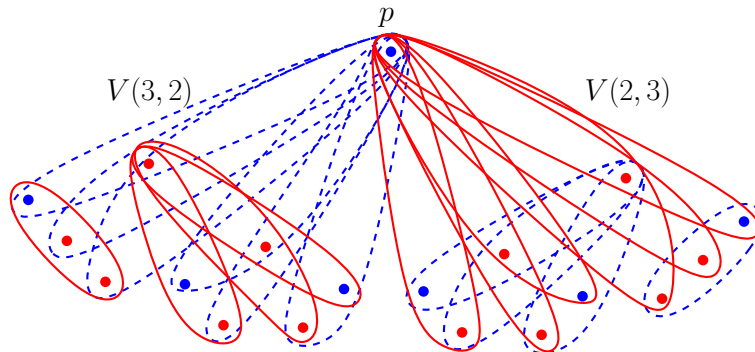**Definition 2.1.** *Let k and l be positive integers.*

Figure 2: The hypergraph $\mathcal{H}(3,3)$ with (arbitrarily) 2-colored vertices. There is a blue (dashed) set with 3 blue vertices or a red (solid) set with 3 red vertices.

1. For $k = 1$, let $V(1,l)$ be an l-element set.
   Set $E_R(1,l) := V(1,l)$ and $E_B(1,l) := \{V(1,l)\}$.

2. For $l = 1$, let $V(k,1)$ be a k-element set.
   Set $E_R(k,1) := \{V(k,1)\}$ and $E_B(k,1) := V(k,1)$.

3. For any $k,l > 1$, we pick a new vertex p, called the root, and let

$$V(k,l) := V(k-1,l) \cup V(k,l-1) \cup \{p\},$$

$$E_R(k,l) := \{e \cup \{p\} : e \in E_R(k-1,l)\} \cup E_R(k,l-1),$$
$$E_B(k,l) := E_B(k-1,l) \cup \{e \cup \{p\} : e \in E_B(k,l-1)\}.$$

By recursion, we obtain that

$$|V(k,l)| = \binom{k+l}{k} - 1,$$

$$|E_R(k,l)| = \binom{k+l-1}{k}, \; |E_B(k,l)| = \binom{k+l-1}{l},$$

$$|E(k,l)| = |E_R(k,l)| + |E_B(k,l)| = \binom{k+l}{k}.$$

Lemma 2.2 ([37]). *For any positive integers k,l, the hypergraph $\mathcal{H}(k,l)$ is not 2-colorable. Moreover, for every coloring of $V(k,l)$ with red and blue, there is an edge in $E_R(k,l)$ such that all of its k vertices are red or an edge in $E_B(k,l)$ such that all of its l vertices are blue.*

For completeness, here we include the proof of Lemma 2.2 from [37]. The induction on two parameters, k and l, is similar to the proof of Ramsey's theorem by Erdős and Szekeres [11].

6

Proof. We will prove that for every coloring of $V(k,l)$ with red and blue, there is an edge in $E_R(k,l)$ such that all of its k vertices are red or an edge in $E_B(k,l)$ such that all of its l vertices are blue.

Suppose first that $k = 1$. If any vertex in $V(1,l)$ is red, then it is a red singleton edge in $\mathcal{H}(1,l)$. If all vertices in $V(1,l)$ are blue, then the (only) edge $V(1,l) \in E_B(1,l)$ contains only blue points. Analogously, the assertion is true if $l = 1$.

Suppose next that $k,l > 1$. Assume without loss of generality that the root p is red. Consider the subhypergraph $\mathcal{H}(k-1,l) \subset \mathcal{H}(k,l)$ induced by the vertices in $V(k-1,l)$. If it has a monochromatic red edge $e \in E_R(k-1,l)$, then $e \cup \{p\} \in E_R(k,l)$ is red. If there is a monochromatic blue edge in $E_B(k-1,l)$, then we are again done, because it is also an edge in $E_B(k,l)$. $\qquad\square$

For other interesting properties of the hypergraphs $\mathcal{H}(k,l)$ related to hereditary discrepancy, see Matoušek [28].

3   Geometric realization of the hypergraphs $\mathcal{H}(k,l)$

The aim of this section is to establish the following weaker version of Theorem 2'.

Theorem 2". *For every $m \geq 2$, there exists a finite point set $P = P(m) \subset \mathbb{R}^2$ and a finite family of unit disks $\mathcal{C} = \mathcal{C}(m)$ with the property that every member of $\mathcal{C}$ contains at least m elements of P, and no matter how we color the elements of P with two colors, there exists a disk in $\mathcal{C}$ such that all points in it are of the same color.*

We *realize* the hypergraph $\mathcal{H}(k,l)$ defined in Section 2 with points and disks. The vertex set $V(k,l)$ is mapped to a point set $P(k,l) \subset \mathbb{R}^2$, and the edge sets, $E_R(k,l)$ and $E_B(k,l)$, to families of open unit disks, $\mathcal{C}_R(k,l)$ and $\mathcal{C}_B(k,l)$, so that a vertex belongs to an edge if and only if the corresponding point is contained in the corresponding disk. The geometric properties of this realization are summarized in the following lemma.

Given two unit disks C, C', let $d(C, C')$ denote the distance between their centers. We fix an orthogonal coordinate system in the plane so that we can talk about the *topmost* and the *bottommost* points of a disk.

Lemma 3.1. *For any positive integers k,l and for any $\varepsilon > 0$, there is a finite point set $P = P(k,l)$ and a finite family of open unit disks $\mathcal{C}(k,l) = \mathcal{C}_R(k,l) \cup \mathcal{C}_B(k,l)$ with the following properties.*

1. *Any disk $C \in \mathcal{C}_R(k,l)$ (resp. $\mathcal{C}_B(k,l)$) contains precisely k (resp. l) points of P.*

2. *For any coloring of P with red and blue, there is a disk in $\mathcal{C}_R(k,l)$ such that all of its points are red or a disk in $\mathcal{C}_B(k,l)$ such that all of its point are blue. In fact, P and $\mathcal{C}(k,l)$ realize the abstract hypergraph $\mathcal{H}(k,l)$ in the above sense.*

3. *For the coordinates (x,y) of any point from P, we have $-\varepsilon < x < \varepsilon$ and $-\varepsilon^2 < y < \varepsilon^2$.*

4. *For the coordinates (x,y) of the center of any disk from $\mathcal{C}_R(k,l)$, we have $-\varepsilon < x < \varepsilon$ and $-\varepsilon^2 < y - 1 < \varepsilon^2$.*

5. *For the coordinates (x,y) of the center of any disk from $\mathcal{C}_B(k,l)$, we have $-\varepsilon < x < \varepsilon$ and $-\varepsilon^2 < y + 1 < \varepsilon^2$.*

6. *The topmost and the bottommost points of a disk $C \in \mathcal{C}(k,l)$ are not covered by the closure of any other member of $\mathcal{C}(k,l)$.*

Looking at our construction from "far away" the two families $\mathcal{C}_R$ and $\mathcal{C}_B$ look like two touching disks, with all points of P very close to the touching point. The segments connecting the centers of disks from different families are almost vertical with all members of $\mathcal{C}_R$ lying "above" all members of $\mathcal{C}_B$. We prove the lemma by induction. Most conditions are needed for the induction to go through. Condition 6 is an exception: it will be used in Section 4.

Proof. We give a recursive construction. We can assume that $\varepsilon < 1/10$. It is easy to see that, for $k = 1$ or $l = 1$, there exists such a family of unit disks for any $\varepsilon > 0$, see Figure 3(a). The family $\mathcal{C}(2,2)$ is depicted in Figure 3(b), where the main idea of the induction may already be visible.

Suppose that $k, l \geq 2$ and we have already constructed $P(k-1,l)$ and $\mathcal{C}(k-1,l)$, and $P(k,l-1)$ and $\mathcal{C}(k,l-1)$, for some $\varepsilon(k-1,l) < \varepsilon/100$ and $\varepsilon(k,l-1) < \varepsilon/100$, respectively. To obtain $P(k,l)$, we place the root p of $\mathcal{H}(k,l)$ into the origin $(0,0)$, and we shift (translate) $P(k-1,l)$ and $P(k,l-1)$ into new positions such that their roots are at $(-\varepsilon/3, -\varepsilon^2/10)$ and $(\varepsilon/3, \varepsilon^2/10)$, respectively. With a slight abuse of notation, the shifted copies will also be denoted $P(k-1,l)$ and $P(k,l-1)$. See Figure 3(c). In this way, it is guaranteed that for the coordinates (x,y) of any point of P, we have

$$-\varepsilon < -(\varepsilon/3 + \varepsilon(k-1,l) + \varepsilon(k,l-1)) < x < \varepsilon/3 + \varepsilon(k-1,l) + \varepsilon(k,l-1) < \varepsilon$$

and

$$-\varepsilon^2 < -(\varepsilon^2/10 + \varepsilon^2(k-1,l) + \varepsilon^2(k,l-1)) < y < \varepsilon^2/3 + \varepsilon^2(k-1,l) + \varepsilon^2(k,l-1) < \varepsilon^2.$$

Thus, **property 3** of the lemma holds.

The family $\mathcal{C}(k,l)$ is defined as the union of two previously defined families, $\mathcal{C}(k-1,l)$ and $\mathcal{C}(k,l-1)$, translated by the same vectors as $P(k-1,l)$ and, resp. $P(k,l-1)$ were. Again, we use the same symbols to denote the translated copies. To verify **properties 4 and 5**, we only have to repeat the above calculations, with the y-coordinates being shifted 1 higher (resp. 1 lower).

Now we show that our set of points $P(k,l)$ and set of disks $\mathcal{C}(k,l)$ realize the hypergraph $\mathcal{H}(k,l)$ (**properties 1 and 2**). It is easy to see that if $C \in \mathcal{C}_R(k-1,l)$ and $s \in P(k,l-1)$, then $s \notin C$ but $p = (0,0) \in C$. The coordinates of the center of C are $\left(-\varepsilon/3 \pm \varepsilon(k-1,l), 1 - \varepsilon^2/10 \pm \varepsilon^2(k-1,l)\right)$ (where here and in the following, $\pm z$ denotes a number that is between $-z$ and z), so the distance of p from C is at most $(\varepsilon/3 + \varepsilon(k-1,l))^2 + (1 - \varepsilon^2/10 + \varepsilon^2(k-1,l))^2 < 1$. On the other hand, the coordinates of s are $\left(\varepsilon/3 \pm \varepsilon(k,l-1), \varepsilon^2/10 \pm \varepsilon^2(k,l-1)\right)$, thus the square of its distance from the center of C is at least

$$\left(2\varepsilon/3 - \varepsilon(k-1,l) - \varepsilon(k,l-1)\right)^2 + \left(1 - 2\varepsilon^2/10 - \varepsilon^2(k-1,l) - \varepsilon^2(k,l-1)\right)^2 > 1.$$

Analogously, if $C \in \mathcal{C}_B(k,l-1)$ and $s \in P(k-1,l)$, then $s \notin C$ but $p = (0,0) \in C$.

Let $C \in \mathcal{C}_R(k,l-1)$ and $s \in P(k-1,l)$. We prove that $p, s \notin C$. The coordinates of the center of C are $\left(\varepsilon/3 \pm \varepsilon(k,l-1), 1 + \varepsilon^2/10 \pm \varepsilon(k,l-1)\right)$. Therefore, the distance of p from the center of C is at least $(\varepsilon/3 - \varepsilon(k,l-1))^2 + (1 + \varepsilon^2/10 - \varepsilon(k,l-1))^2 > 1$. The calculation for s is similar in the case $C \in \mathcal{C}_R(k-1,l)$. Analogously, we have that if $C \in \mathcal{C}_B(k-1,l)$ and $s \in P(k,l-1)$, then $p, s \notin C$. As the disks in $\mathcal{C}(k,l-1)$ (resp. $\mathcal{C}(k-1,l)$) contain precisely the same points of $P(k,l-1)$ (resp. $P(k-1,l)$, as before the shift, we have obtained a geometric realization of $\mathcal{H}(k,l)$, and properties 1 and 2 hold.

(a) Starting step: $\mathcal{C}(k,1)$.

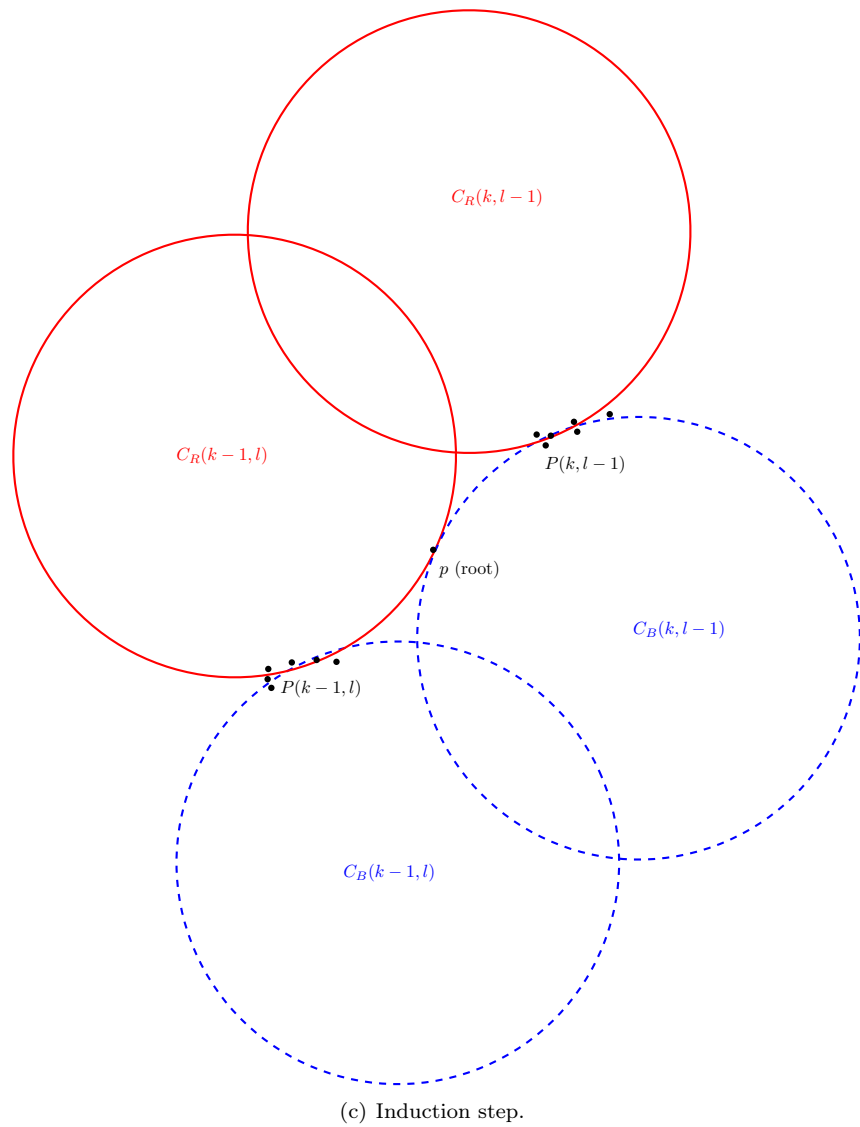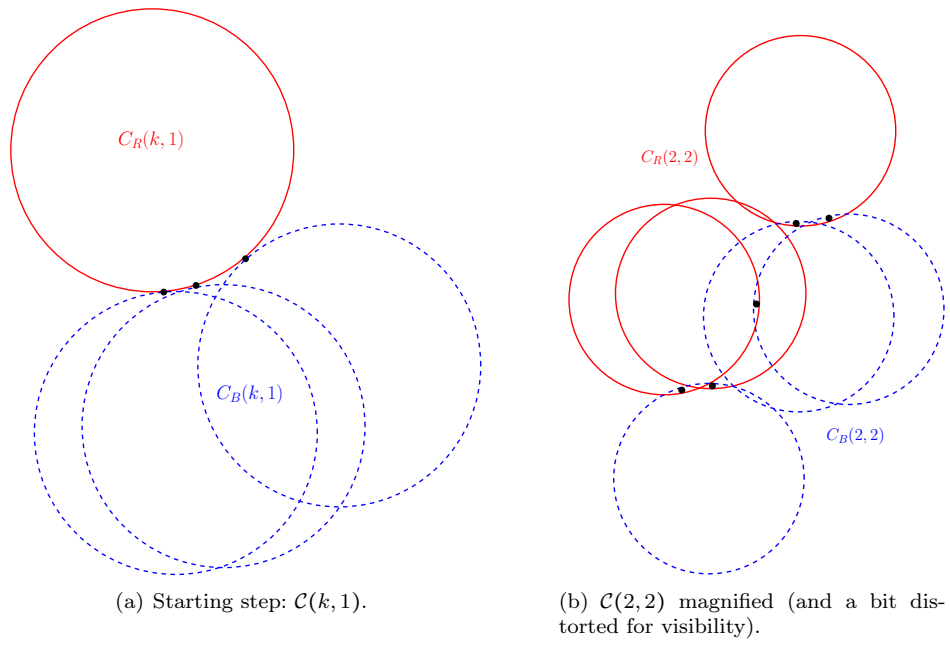(b) $\mathcal{C}(2,2)$ magnified (and a bit distorted for visibility).

(c) Induction step.

Figure 3: The construction.

It remains to prove that the topmost and the bottommost points of a disk $C \in \mathcal{C}(k,l)$ are not covered by any other member of $\mathcal{C}(k,l)$ (**property 6**). Using that our construction and disks are centrally symmetric, it is enough to prove the statement for the topmost points. If $C \in \mathcal{C}_R(k, l-1)$, the coordinates of its topmost point are $\big(\varepsilon/3 \pm \varepsilon(k, l-1), 2 + \varepsilon^2/10 \pm \varepsilon^2(k, l-1)\big)$. If $C \in \mathcal{C}_R(k-1, l)$, the coordinates of its topmost point are $\big(-\varepsilon/3 \pm \varepsilon(k-1, l), 2 - \varepsilon^2/10 \pm \varepsilon^2(k-1, l)\big)$. If $C \in \mathcal{C}_B(k, l-1)$, the coordinates of its topmost point are $\big(\varepsilon/3 \pm \varepsilon(k, l-1), -2 + \varepsilon^2/10 \pm \varepsilon^2(k, l-1)\big)$. If $C \in \mathcal{C}_B(k-1, l)$, the coordinates of its topmost point are $\big(-\varepsilon/3 \pm \varepsilon(k-1, l), -2 - \varepsilon^2/10 \pm \varepsilon^2(k-1, l)\big)$.

If $C \in \mathcal{C}_R(k, l-1)$, by the induction hypothesis, its topmost point cannot be covered by any other disk from $\mathcal{C}(k, l-1)$. Nor can it be covered by any other disk, as the topmost points of all other disks are below it (i.e., have smaller y-coordinates). If $C \in \mathcal{C}_R(k-1, l)$, then the square of the distance of its topmost point from the center of some $C' \in \mathcal{C}_R(k, l-1)$ is at least

$$\big(2\varepsilon/3 - \varepsilon(k, l-1) - \varepsilon(k-1, l)\big)^2 + \big(1 - 2\varepsilon^2/10 - \varepsilon^2(k, l-1) - \varepsilon^2(k-1, l)\big)^2 > 1.$$

If $C \in \mathcal{C}_B(k, l-1)$, then the distance of its topmost point from the center of some $C' \in \mathcal{C}_R(k-1, l)$ is also at least

$$\big(2\varepsilon/3 - \varepsilon(k, l-1) - \varepsilon(k-1, l)\big)^2 + \big(1 - 2\varepsilon^2/10 - \varepsilon^2(k, l-1) - \varepsilon^2(k-1, l)\big)^2 > 1.$$

In all other cases, trivially, the corresponding distances are also larger than 1. This completes the proof of property 6 and hence the lemma. $\qquad \square$

4   Adding points to P — Proof of Theorem 2'

In this section, we extend the proof of Theorem 2" to establish Theorem 2' (which is equivalent to Theorem 2). Note that the only difference between Theorems 2" and 2' is that in the latter it is also required that every unit disk of the plane contains at least m elements of the point set $P^* = P^*(m)$. The set $P = P(m, m)$ constructed in Lemma 3.1, does not satisfy this condition. In order to fix this, we will add all points *not* in $\cup \mathcal{C}(m, m)$ to the set P (or rather a sufficiently dense discrete subset of $\mathbb{R}^2 \smallsetminus \cup \mathcal{C}(m, m)$). In order to show that the resulting set P^* meets the requirements of Theorem 2', all we have to show is the following.

Lemma 4.1. *No (open) unit disk $C \notin \mathcal{C}(k, l)$ is entirely contained in $\cup \mathcal{C}(k, l)$.*

For future purposes, we prove this statement in a slightly more general form. In what follows, we only assume that C is an open convex body with a unique topmost point t and a unique bottommost point b, which divide the boundary of C into two *closed arcs*. They will be referred to as the *left boundary arc* and a *right boundary arc*.

Definition 4.2. *A collection $\mathcal{C}$ of translates of C is said to be* exposed *if the topmost and bottommost points of its members do not belong to the closure of any other member of $\mathcal{C}$.*

By the last condition in Lemma 3.1, the collections of disks $\mathcal{C}(k, l)$ constructed in the previous section are exposed. We prove the following generalization of Lemma 4.1.

Lemma 4.3. *Let $\mathcal{C}$ be a finite exposed collection of translates of an open convex body C with unique topmost and bottommost points. If $C \notin \mathcal{C}$, then $C \nsubseteq \cup\mathcal{C}$.*

For the proof, we need a simple observation.

Claim 4.4. *If the right boundary arcs of two translates of C intersect, then the closure of one of the translates must contain the topmost or bottommost point of the other.*

Proof. Let C_1 and C_2 be the two translates, and let γ_i denote the closed convex curve formed by the right boundary arc of C_i and the straight-line segment connecting its two endpoints (the topmost and the bottommost points of C_i). The curves γ_1 and γ_2 are translates of each other, and since they intersect, they must *cross* twice. (At a *crossing*, one curve comes from the exterior of the other, then it shares an arc with it, which may be a single point, and enters the interior.) It cannot happen that both crossings occur between the right boundary arcs, because they are convex and translates of each other. Therefore, one of the two crossings involves the straight-line segment of one the curves, say, γ_1. But since the condition is that the right boundary arcs intersect, one of the two endpoints of this straight-line segment, either the topmost or the bottommost point of C_1, lies in the closure of C_2 □

Proof of Lemma 4.3. Suppose, for contradiction, that $C \subseteq \cup\mathcal{C}$. By removing some members of $\mathcal{C}$ if necessary, we can assume that $\mathcal{C}$ is a *minimal* collection of translates that covers C. Then C must have a point which belongs to (at least) three translates, $C_1, C_2, C_3 \in \mathcal{C}$. None of the topmost and bottommost points of these translates can be covered by C, otherwise, it would also be covered by another member of $\mathcal{C}$, contradicting the assumption that $\mathcal{C}$ is exposed.

Thus, C intersects either the left or the right boundary arc of every C_i. Without loss of generality, suppose that C intersects the right boundary arcs of C_1 and C_2. These right boundary arcs must intersect inside C, otherwise $C_1 \cap C \subseteq C_2 \cap C$ or $C_2 \cap C \subseteq C_1 \cap C$, and $\mathcal{C}$ would not be minimal. Therefore, we can apply Claim 4.4 to conclude that one of them must contain the topmost or bottommost point of the other. □

Remark 4.5. In the construction described in Lemma 3.1, every disk in $\mathcal{C}(m,m)$ contains at most $|P(m,m)| < 2^{2m}$ points. At the last stage, we added many new points to P. We can keep the maximum number of points of P lying in a unit disk bounded from above by a function $f(m)$. What is the best upper bound? The bound given by our construction depends on $\varepsilon(m,m) \leq 100^{-2m}\varepsilon(1,1)$.

5 Other convex bodies – Proof of Theorem 3

Throughout this section, C denotes an open plane convex body which has two parallel supporting lines with positive curvature at the two points of tangencies. To prove Theorem 3, by duality, it is sufficient to establish the analogue of Theorem 2', where the role of unit disks is played by translates of C.

Theorem 3'. *For every $m \geq 2$, there is a set of points $P^* = P^*(m)$ in the plane with the property that every translate of C contains at least m elements of P^*, and no matter how we color the elements of P^* with two colors, there exists a translate of C such that all points in it are of the same color.*

As in the case of disks, after defining the hypergraphs $\mathcal{H}(k,l)$, the proof consists of two steps:

Step 1: We find a geometric realization of $\mathcal{H} = \mathcal{H}(k,l)$ with translates of C, i.e., a finite point set P representing the vertices and a collection $\mathcal{C}$ of translates of C representing the hyperedges of $\mathcal{H}$ such that a point of P lies in a member of $\mathcal{C}$ if and only if the corresponding vertex belongs to the corresponding hyperedge. We show that $\mathcal{C}$ is an exposed family.

Step 2: We show that no translate of C is entirely contained in $\cup\mathcal{C}$, unless $C \in \mathcal{C}$. Thus, we can add all the points not in $\cup\mathcal{C}(k,l)$ to the points of P to ensure that every translate of C contains many points.

In Section 4, we have shown that Step 2 can be completed, provided that $\mathcal{C}$ is exposed (see Lemma 4.3). Therefore, here we concentrate on Step 1.

Without loss of generality, we can assume that C has unique bottommost and topmost points, b and t, resp., at which the curvature is positive. After applying an affine transformation, we can also attain that the line bt is vertical. Let r_b and r_t denote the reciprocals of the curvatures at b and t, respectively. If we place b at the origin, then, for every $\delta > 0$, in a small neighborhood of b, the boundary of C will lie between the parabolas $y = (1-\delta)r_b x^2$ and $y = (1+\delta)r_b x^2$. Analogously, if we place t at the origin, then in a small neighborhood of it, the boundary of C will lie between the parabolas $y = -(1-\delta)r_t x^2$ and $y = -(1+\delta)r_t x^2$. We find a geometric realization using the following lemma.
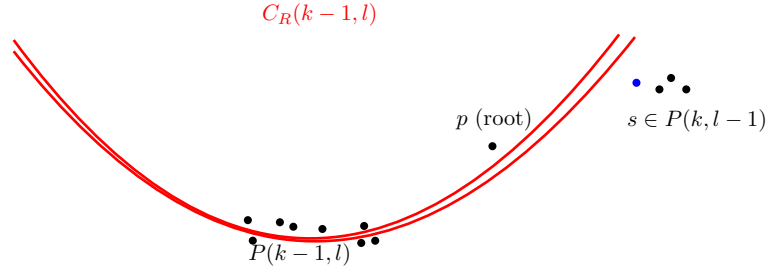

Figure 4: Parabolas enclosing the boundary of C.

Lemma 3.1'. *For any positive integers k,l and for any $\varepsilon > 0$, there is a finite point set $P = P(k,l)$ and a finite family of translates of C, $\mathcal{C}(k,l) = \mathcal{C}_R(k,l) \cup \mathcal{C}_B(k,l)$ with the following properties.*

1. *Any translate from $\mathcal{C}_R(k,l)$ (resp. $\mathcal{C}_B(k,l)$) contains precisely k (resp. l) points of P.*

2. *For any coloring of P with red and blue, there is a translate from $\mathcal{C}_R(k,l)$ such that all of its points are red or a translate from $\mathcal{C}_B(k,l)$ such that all of its point are blue. In fact, P and $\mathcal{C}(k,l)$ realize the abstract hypergraph $\mathcal{H}(k,l)$ in the above sense.*

3. *For the coordinates (x,y) of any point from P, we have $-\varepsilon < x < \varepsilon$ and $-\varepsilon^2 < y < \varepsilon^2$.*

4. *For the coordinates (x, y) of the bottommost point of any translate from $\mathcal{C}_R(k, l)$, we have $-\varepsilon < x < \varepsilon$ and $-\varepsilon^2 < y < \varepsilon^2$.*

5. *For the coordinates (x, y) of the topmost point of any translate from $\mathcal{C}_B(k, l)$, we have $-\varepsilon < x < \varepsilon$ and $-\varepsilon^2 < y < \varepsilon^2$.*

6. *The topmost and the bottommost points of translate from $\mathcal{C}(k, l)$ are not covered by the closure of any other member of $\mathcal{C}(k, l)$.*

Proof. Using an affine transformation, we can suppose that $r_t, r_b < 1$. We fix a δ that is small enough compared to r_t and r_b, and an $\varepsilon = \varepsilon(k, l)$ that is small enough compared to δ, r_t and r_b, but big enough compared to $\varepsilon(k, l-1)$ and $\varepsilon(k-1, l)$. (To keep the presentation simple, we omit the exact required dependencies here.) We will use that the boundary of C in a $2\varepsilon(r_t + r_b)$ neighborhood around t and b is between the (above mentioned) pairs of parabolas, $y = (1 - \delta)r_b x^2$ and $y = (1 + \delta)r_b x^2$, and $y = -(1 - \delta)r_t x^2$ and $y = -(1 + \delta)r_t x^2$.

If $k = 1$ or $l = 1$, the construction is trivial. For $k, l \geq 2$, assume that the point sets $P(k-1, l)$ and $P(k, l-1)$, and the families of translates of C, $\mathcal{C}(k-1, l)$ and $\mathcal{C}(k, l-1)$, have already been defined, and that they satisfy all conditions in the lemma. To obtain $P(k, l)$, we place the root p of $\mathcal{H}(k, l)$ into the origin $(0, 0)$, and we shift $P(k-1, l)$ and $P(k, l-1)$ such that their roots are at $(-r_t \varepsilon, -(1 + 2\delta)r_b r_t^2 \varepsilon^2)$ and $(r_b \varepsilon, (1 + 2\delta)r_t r_b^2 \varepsilon^2)$, respectively. The family of translates $\mathcal{C}(k, l)$ is defined as the union of the families $\mathcal{C}(k-1, l)$ and $\mathcal{C}(k, l-1)$ translated by the same vectors, as $P(k-1, l)$ and $P(k, l-1)$, respectively.

To verify **properties 3, 4, and 5**, we need that $-\varepsilon < -r_t \varepsilon, r_b \varepsilon < \varepsilon$ and $-\varepsilon^2 < -(1 + 2\delta)r_b r_t^2 \varepsilon^2, (1 + 2\delta)r_t r_b^2 \varepsilon^2 < \varepsilon^2$, which hold since $r_t, r_b < 1$ and δ is small. Notice that where we have omitted $\varepsilon(k, l-1)$ and $\varepsilon(k-1, l)$ from these equations to keep the calculations simple. This we can do as the difference of the two sides depends on ε, which we can select to be sufficiently large compared to $\varepsilon(k, l-1)$ and $\varepsilon(k-1, l)$. We will also omit dependencies of $\varepsilon(k, l-1)$ and $\varepsilon(k-1, l)$ later.

To verify **properties 1 and 2**, we have to show that for any $C \in \mathcal{C}_R(k-1, l)$, the origin $p = (0, 0)$ belongs to C, but no point $s \in P(k, l-1)$ does, provided that $\varepsilon > 0$ is sufficiently small. To see this, fix $C \in \mathcal{C}_R(k-1, l)$. The equation of the parabola that touches C from the inside at its bottommost point is approximately $y = (1 + \delta)r_b(x + r_t \varepsilon)^2 - (1 + 2\delta)r_b r_t^2 \varepsilon^2$. If $x = 0$, the value of y is $(1 + \delta)r_b(r_t \varepsilon)^2 - (1 + 2\delta)r_b r_t^2 \varepsilon^2 = -\delta r_b r_t^2 \varepsilon^2$. This is negative, which means that $p = (0, 0)$ lies above the parabola. Thus, we have $p \in C$. Analogously, if $C \in \mathcal{C}_B(k, l-1)$ and $s \in P(k-1, l)$, then $s \notin C$ but $p = (0, 0) \in C$.

On the other hand, the equation of the parabola that touches C at its bottommost point from the outside is approximately $y = (1 - \delta)r_b(x + r_t \varepsilon)^2 - (1 + 2\delta)r_b r_t^2 \varepsilon^2$. If $x = r_b \varepsilon \pm \varepsilon(k, l-1)$ the value of y at x is approximately

$$(1 - \delta)r_b(r_b \varepsilon + r_t \varepsilon)^2 - (1 + 2\delta)r_b r_t^2 \varepsilon^2 = \left((1 - \delta)(r_b^3 + 2r_b^2 r_t) - 3\delta r_b r_t^2\right)\varepsilon^2 \geq \left(r_b^3 + O(\delta)\right)\varepsilon^2.$$

Therefore, $s = (r_b \varepsilon \pm \varepsilon(k, l-1), (1 + 2\delta)r_t r_b^2 \varepsilon^2 \pm \varepsilon^2(k, l-1))$ is below the parabola, if δ is small enough, thus $s \notin C$.

Let $C \in \mathcal{C}_R(k, l-1)$ and $s \in P(k-1, l)$. We prove that $p, s \notin C$. The equation of the parabola that touches C from the outside at its bottommost point is approximately $y = (1 - \delta)r_t(x - r_b \varepsilon)^2 - (1 + 2\delta)r_t r_b^2 \varepsilon^2$. If $x = 0$, the value of y is $(1 + \delta)r_t(-r_b \varepsilon)^2 - (1 + 2\delta)r_t r_b^2 \varepsilon^2 = -\delta r_b r_t^2 \varepsilon^2 < 0$, thus $p \in C$. The calculation for s is similar in the case $C \in \mathcal{C}_R(k-1, l)$. Analogously, we have that if $C \in \mathcal{C}_B(k-1, l)$ and $s \in P(k, l-1)$, then $p, s \notin C$. As the translates in $\mathcal{C}(k, l-1)$ (resp. $\mathcal{C}(k-1, l)$) contain precisely the same

points of $P(k, l-1)$ (resp. $P(k-1, l)$, as before the shift, we have obtained a geometric realization of $\mathcal{H}(k, l)$, and properties 1 and 2 hold.

It remains to prove that the topmost and the bottommost points of a translate $\mathcal{C}(k, l)$ are not covered by any other member of $\mathcal{C}(k, l)$ (**property 6**). Using that our construction is symmetric, it is enough to prove the statement for the topmost points. Recall that the line connecting b and t is vertical and denote their distance, the height of C, by h.

The coordinates of the topmost points of translates from $\mathcal{C}_R(k, l-1)$ are approximately $(r_b \varepsilon + h, (1+2\delta) r_t r_b^2 \varepsilon^2 + h)$. The coordinates of the topmost points of translates from $\mathcal{C}_R(k-1, l)$ are approximately $(-r_t \varepsilon + h, -(1+2\delta) r_b r_t^2 \varepsilon^2 + h)$. The coordinates of the topmost points of translates from $\mathcal{C}_B(k, l-1)$ are approximately $(r_b \varepsilon, (1+2\delta) r_t r_b^2 \varepsilon^2)$. The coordinates of the topmost points of translates from $\mathcal{C}_B(k-1, l)$ are approximately $(-r_t \varepsilon, -(1+2\delta) r_b r_t^2 \varepsilon^2)$.

If $C_1 \in \mathcal{C}_R(k, l-1)$, by the induction hypothesis, its topmost point cannot be covered by any other $C_2 \in \mathcal{C}(k, l-1)$. Nor can it be covered by any other translate, as the topmost points of all other translates are below it (i.e., have smaller y-coordinates). If $C_1 \in \mathcal{C}_R(k-1, l)$, then the vector connecting it to the topmost point of some $C_2 \in \mathcal{C}_R(k, l-1)$ is approximately the same as the vector connecting a point $s \in P(k-1, l)$ to the topmost point of some $C' \in \mathcal{C}_B(k, l-1)$. As we have seen earlier that $s \notin C'$, the same calculation shows that the topmost point of C_1 is not in C_2. If $C_1 \in \mathcal{C}_R(k-1, l)$ and $C_2 \in \mathcal{C}_B(k, l-1)$ or $C_2 \in \mathcal{C}_B(k-1, l)$, then the topmost point of C_2 is lies below the topmost point of C_1. If $C_1 \in \mathcal{C}_B(k, l-1)$ and $C_2 \in \mathcal{C}(k, l-1)$, by induction the topmost point of C_1 is not in C_2. If $C_1 \in \mathcal{C}_B(k, l-1)$ and $C_2 \in \mathcal{C}(k-1, l)$, then the topmost point of C_1 is approximately at the same place as the points of $P(k, l-1)$ which are avoided by C_2, and the same calculation works here. Similarly, if $C_1 \in \mathcal{C}_B(k-1, l)$ and $C_2 \in \mathcal{C}(k-1, l)$, we can use induction, and if $C_1 \in \mathcal{C}_B(k-1, l)$ and $C_2 \in \mathcal{C}(k, l-1)$, we can use that the topmost point of C_1 is approximately at the same place as the points of $P(k-1, l)$ which are avoided by C_2, the same calculation works here. This completes the proof of property 6 and hence the lemma. $\square$

6   Shift-chains – Proof of Theorem 6

Throughout this section, P denotes a fixed set of n points in the plane, no two of which have the same x-coordinate, and C is a fixed open convex set that contains a vertical upward half-line.

Definition 6.1. *For $A \subset [n] = \{1, 2, \ldots, n\}$, denote by a_i the i^{th} smallest element of A. For two equal sized sets, $A, B \subset [n]$, we write $A \preceq B$ if $a_i \le b_i$ for every i.*

An m-uniform hypergraph on the vertex set $[n]$ is called a shift-chain *if its hyperedges are totally ordered by the relation $\preceq$. A shift-chain $\mathcal{H}$ is* special *if for any two hyperedges, $A, B \in \mathcal{H}$ with $A \preceq B$, we have $\max(A \smallsetminus B) < \min(B \smallsetminus A)$.*

For any integer m and real number x, let $C(m; x)$ denote the translate of C which
   a. contains exactly m points of P,
   b. can be obtained from C by translating it through a vector with x-coordinate x,
   c. and has minimum y-coordinate, among all translates satisfying **a** and **b**.
The union of all translates of C through every vector that has x-coordinate x is a vertical strip (or an open half-plane or the whole plane), denoted by $S(x)$. If $S(x)$ contains precisely m points for some x, then in condition **c**, the minimum y-coordinate is $y = -\infty$,

and we set $C(m;x) = S(x)$. If $S(x)$ contains fewer than m points, then $C(m;x)$ is undefined.

Proposition 6.2. *Let $p_1, p_2, \ldots, p_n$ denote the elements of P, listed in the increasing order of their x-coordinates. Then the m-uniform hypergraph consisting of the sets $P(x) = \{i \in [n]; p_i \in C(m;x)\}$, over all $x \in \mathbb{R}$, is a special shift-chain.*

Proof. Notice that if $x < x'$, then the boundary of $C(m;x)$ intersects the boundary of $C(m;x')$ precisely once. Therefore, every element of $(C(m;x) \smallsetminus C(m;x')) \cap P$ is to the left of all elements of $(C(m;x') \smallsetminus C(m;x)) \cap P$. This means that $P(x) \preceq P(x')$. $\square$

In view of the duality described at the end of the introduction, Theorem 6 is an immediate corollary of the following statement.

Theorem 6.3. *For any $m \geq 3$, every m-uniform special shift-chain is 2-colorable. Moreover, such a coloring can be constructed in linear time.*

An example found by Fulek [16] (depicted on Figure 5) shows that Theorem 6.3 is false without assuming that the shift-chain is special.
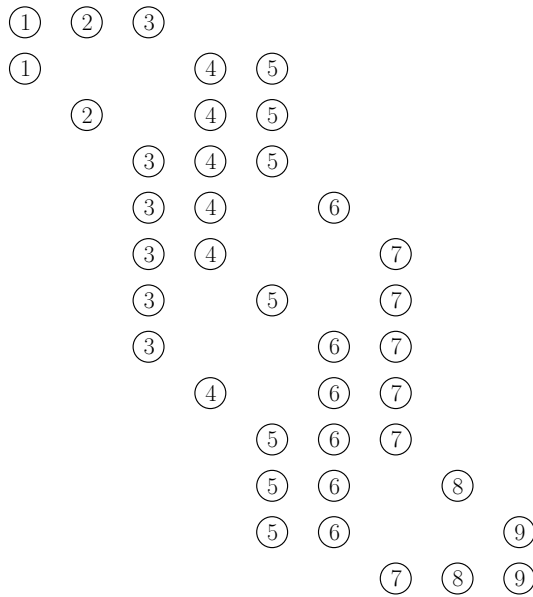

Figure 5: A shift-chain of 13 triples, each of which corresponds to a row. For any 2-coloring of the 9 vertices, one of the triples is monochromatic.

Problem 6.4. *Does there exist an integer $m_0 > 3$ such that for every $m \geq m_0$, every m-uniform shift-chain is 2-colorable?*

If the answer to this question is yes, in some sense this could be regarded as an extension of the Lovász local lemma [10]. For more problems and results related to shift-chains and special shift-chains, consult [36, 25].

Proof of Theorem 6.3. The proof breaks into several simple claims. In the rest of this section, $\mathcal{H}$ denotes a fixed 3-uniform special shift-chain on $[n] = \{1, 2, \ldots, n\}$. For simplicity, a hyperedge (triple) $\{a, b, c\} \in \mathcal{H}$ with $a < b < c$ will be denoted by $\{a < b < c\}$.

Claim 6.5. *If $\{a < b < c\} \in \mathcal{H}$ and $\{a' < b < c'\} \in \mathcal{H}$, then $a' = a$ or $c' = c$.*

Proof. Otherwise, $\{a < b < c\} \smallsetminus \{a' < b < c'\} = \{a < c\}$ and $\{a' < b < c'\} \smallsetminus \{a < b < c\} = \{a' < c'\}$ would not be separated, contradicting our assumption that $\mathcal{H}$ is special. $\qquad\square$

Define a digraph, $D = D(\mathcal{H})$ with vertex set $[n]$ and edge set E, as follows. For any $b < c$, the directed edge $bc \in E$ if and only if there exist $a, a' \in [n]$, $a \neq a'$, such that $\{a < b < c\} \in \mathcal{H}$ and $\{a' < b < c\} \in \mathcal{H}$. Analogously, for any $a < b$, the directed edge $ba \in E$ if and only if there exist $c, c' \in [n]$, $c \neq c'$, such that $\{a < b < c\} \in \mathcal{H}$ and $\{a < b < c'\} \in \mathcal{H}$. According to Claim 6.5, the out-degree of every vertex of D is at most one. Note that an edge may appear in E with both orientations ab and ba.

Claim 6.6. *The directed graph D can be constructed by a linear time algorithm.*

Proof. $\mathcal{H}$, as any 3-uniform shift-chain on n vertices has at most $3n - 8$ hyperedges. Suppose that they are listed in an arbitrary order, and process them one-by-one. Suppose the next triple is $\{a < b < c\}$.

1. If b is a middle vertex for the first time, store it, together with both of its *neighbors*, a and c.

2. If b is a middle vertex for the second time, decide if it was a or c that has been previously stored as one of its neighbors. (By Claim 6.5, we know that one them was.) If it is a, add ba to E, if it is c, add bc.

3. Otherwise, do not add any new edge, and pass to the next triple. $\qquad\square$

Claim 6.7. *For $a < b < c$ (or $c < b < a$) it is not possible that $ac \in E$ and $ba \in E$.*

Proof. Suppose $ac, ba \in E$. By definition, this means that there exist $\{x < a < c\} \in \mathcal{H}$ and $\{a < b < y\} \in \mathcal{H}$, for some x and some $y \neq c$. Obviously, with respect to the ordering of the triples, we have $\{x < a < c\} \prec \{a < b < y\}$. The sets $\{x < a < c\} \smallsetminus \{a < b < y\} = \{x < c\}$ and $\{a < b < y\} \smallsetminus \{x < a < c\} = \{b < y\}$ are not separated, because the maximal element of the first set, c, is larger than the minimal element of the second set, b. This contradicts our assumption that $\mathcal{H}$ was a *special* shift-chain. $\qquad\square$

Claim 6.8. *For $a < b < c < d$ it is not possible that $bd \in E$ and $ca \in E$.*

Proof. This would mean that there exist $\{x < b < d\} \in \mathcal{H}$ with $x \neq a$ and $\{a < c < y\} \in \mathcal{H}$ with $y \neq d$. These two triples are disjoint, but not separated, contradicting the assumption that $\mathcal{H}$ is special. $\qquad\square$

If a directed graph T can be obtained from a directed tree oriented toward its root r, by possibly adding one of the edges pr entering the root also with the reverse orientation rp, then it is called a *quasi-tree*. Note that in this case, we can also think of T as a quasi-tree rooted in p.

Claim 6.9. *The graph D is the vertex-disjoint union of quasi-trees.*

Proof. As no vertex of D has out-degree larger than 1, it is enough to show that D has no directed cycle of length larger than 2. Suppose there is such a directed cycle, and denote its smallest and largest elements by a and b, respectively. By Claim 6.7, we have that $ab \notin E$ and $ba \notin E$. Let $ya \in E$ and $ax \in E$ be the incoming edge and the outgoing edge of the cycle at a. Again, by Claim 6.7, we have $a < x < y < b$. There is a directed path from x to b, and along this path there is a first edge uv with $u < y$ and $v > y$. But then the edges $ya, uv \in E$ would contradict Claim 6.8, as $a < u < y < v$. $\qquad\square$

Now we are in a position to find a 2-coloring of $\mathcal{H}$ in linear time. For every $\{a < b < c\} \in \mathcal{H}$, we will guarantee that the color of its middle vertex, b, will differ from the color of a or the color of c.

First, using breadth-first search, we properly 2-color each connected component. Hence, it will be guaranteed that if the out-degree of b is non-zero, then all triples of the form $\{a < b < c\} \in \mathcal{H}$ contain both colors. Then assign to each vertex $x \in [n]$ an edge $ab \in E$ such that $a < x < b$ or $b < x < a$, provided that such an edge exists. This can be done in linear time, but here we omit the details.

For every $\{x < y < z\} \in \mathcal{H}$ that does not yet contain both colors, its middle vertex y has out-degree zero. If there is an edge $bc \in E$ assigned to y such that $b < y < c$, then there are two different hyperedges $\{a < b < c\}, \{a' < b < c\} \in \mathcal{H}$. Either $a \neq x$ or $a' \neq x$, and thus, necessarily, we have $c = z$. We color y with the same color as b (i.e., differently from $c = z$), so that $\{x < y < z\}$ contains both colors. Note that if the in-degree of y is non-zero, then a simple case analysis shows that the only possibility is $zy \in E$. Thus, this color agrees with the color given earlier to y from its connected component.

In a similar manner, if there is an edge $ab \in E$ assigned to y such that $a < y < b$, then necessarily $a = x$, and we can color y with the same color as b (i.e., differently from $a = x$).

Finally, if there are uncolored vertices, color them in increasing order so that when y is colored, if there are $\{x < y < z\} \in \mathcal{H}$, then y gets the opposite color as x. (This step is well defined, because it follows from the fact that the out-degree of y is zero, that there is only one triple $\{x < y < z\}$ with the above property.)

This completes the proof of Theorem 6.3, as it follows that the middle vertex of any triple will have a different color from another vertex of the triple. $\qquad\square$

7   Covering space with unbounded convex sets

Every open, unbounded, line-free convex set C is contained in a half-space with inner normal vector $\vec{v}$ such that for any $c \in C$, the half-line emanating from c and pointing in the direction of $\vec{v}$ lies entirely in C. We can assume without loss of generality that $\vec{v}$ is the unit vector $e_d = (0, 0, \ldots, 0, 1)$, pointing *vertically upwards*, and that C lies in the *upper half-space*.

First, we prove Proposition 7, according to which every covering of $\mathbb{R}^d$ with translates of a set C satisfying the above conditions can be split into two, and hence into infinitely many, coverings. We prove a slight generalization of this statement, in which C is not required to be convex.

Proposition 7'. *Let C be an open set in the upper half-space of $\mathbb{R}^d$, which has the property that, for every $c \in C$, the half-line starting at c and pointing vertically upwards belongs to C. Then every covering of $\mathbb{R}^d$ with translates of C splits into two coverings.*

Proof. For any positive integer i, let B_i denote the *closed* $(d-1)$-dimensional ball of radius i around the origin in the coordinate hyperplane $x_d = 0$. Let $\mathcal{C}$ be a covering of $\mathbb{R}^d$ with translates of C. As the members of $\mathcal{C}$ cover the whole d-dimensional space, they also cover the $(d-1)$-dimensional ball $B_1 \times \{0\}$, orthogonal to the x_d-axis. This set is *compact* and the members of $\mathcal{C}$ are *open*. Therefore, there is a *finite* subfamily $\mathcal{C}_1 \subset \mathcal{C}$ which covers $B_1 \times \{0\}$. Choose a number $z_1 < 0$ such that all members of $\mathcal{C}_1$ lie strictly above the hyperplane $x_d = z_1$, and consider the $(d-1)$-dimensional ball $B_2 \times \{z_1\}$. Select a finite family $\mathcal{C}_2 \subset \mathcal{C}$ that covers this ball and a number $z_2 < z_1$ such that all members of $\mathcal{C}_2$ lie strictly above the hyperplane $x_d = z_2$. Proceeding like this, we can construct an infinite sequence of disjoint finite subfamilies $\mathcal{C}_1, \mathcal{C}_2, \ldots \subset \mathcal{C}$ and a sequence of reals $z_0 := 0 > z_1 > z_2 > \ldots$ tending to $-\infty$ such that $\mathcal{C}_i$ covers the $(d-1)$-dimensional ball $B_i \times \{z_{i-1}\}$.

Let p be any point of $\mathbb{R}^d$ which is at distance r from the d^{th} coordinate axis and whose d^{th} coordinate is p_d. Notice that p lies above some point of every $(d-1)$-dimensional ball $B_i \times \{z_{i-1}\}$ such that $i \geq r$ and $z_{i-1} \leq p_d$. Consequently, p is covered by the corresponding families $\mathcal{C}_i$. Hence, $\mathcal{C}_1 \cup \mathcal{C}_3 \cup \mathcal{C}_5 \cup \ldots$ and $\mathcal{C}_2 \cup \mathcal{C}_4 \cup \mathcal{C}_6 \cup \ldots$ are two disjoint subfamilies of $\mathcal{C}$, each of which covers the whole space. $\qquad\square$

Next, we establish Theorem 8, which shows that starting from 4-dimensions, Proposition 7 is false if we drop the assumption that C is an *open* set.

Proof of Theorem 8. We have to prove that there is a convex, bounded (not open) set $C' \subset \mathbb{R}^3$ such that $\mathbb{R}^4$ can be covered by translates of $C = C' \times [0, \infty)$ so that every point of $\mathbb{R}^4$ is covered infinitely many times, but this covering cannot be decomposed into two.

The set C' will be the convex hull of $\cup_{i=1}^{\infty}(C_i \times \{\frac{1}{i^2}\})$, where each C_i is in $\mathbb{R}^2$, and thus $C_i \times \{\frac{1}{i^2}\}$ lies in the plane determined by the equation $z = \frac{1}{i^2}$ of $\mathbb{R}^3$. Each C_i is the union of an open disk, defined by the inequality $x^2 + (y - \frac{1}{i})^2 < 1$, and a part of its boundary defined as follows. A point belongs to the boundary of C_i if and only if it can be represented as $(x, \sqrt{1-x^2} + \frac{1}{i})$, where $x \in [0,1]$ and the i^{th} digit of x after the "decimal" point in binary form is 1. For each i, denote the set of such x's by C_i^*. Therefore, C_i^* is the disjoint union of 2^{i-1} closed intervals.

Note that C' is neither closed, nor open. Clearly, C' is a bounded set, as it is contained in the box $[-1, 1] \times [-1, 2] \times [0, 1]$. Observe that for every i, the point $(0, \frac{1}{i}, \frac{1}{i^2})$, the center of the disk $C_i \times \{\frac{1}{i^2}\}$, lies the plane $x = 0$, on the parabola $z = y^2$. Hence, for each i and for every point $p \in \mathbb{R}^3$ whose third coordinate is $\frac{1}{i^2}$ and first coordinate is non-negative, p belongs to the boundary of C' if and only if it is of the form $(x, \sqrt{1-x^2} + z, z^2)$ with $x \in C_i^*$ and $z = \frac{1}{i}$. To see this, it is enough to notice that no point of this form can be obtained as a convex combination of other points in C'.

Now we describe an infinite-fold covering $\mathcal{C}$ of $\mathbb{R}^4$ with translates of C that cannot be decomposed into two coverings. Let $X = \{(x, \sqrt{1-x^2}, 0, -w) \mid x \in [0,1], w \in [0, \infty)\}$. For every point $x \notin X$, select an arbitrary translate of C that covers x and does not intersect X. (It is easy to see that such a translate always exists.) Let $\mathcal{C}$ consist of all these translates, and for every i $(i = 1, 2, \ldots)$, the translate of C through the vector $(0, -\frac{1}{i}, -\frac{1}{i^2}, -i)$, denoted by $\hat{C}_i$.

Notice that the $\hat{C}_i$ covers $(x, \sqrt{1-x^2}, 0, -w) \in X$ if and only if $x \in C_i^*$ and $w \leq i$. This implies that every point of X is covered by infinitely many members of $\mathcal{C}$, because every number has a representation with infinitely many digits that are 1.

It remains to show that $\mathcal{C}$ cannot be split into two coverings. This is a direct consequence of the following statement: For any $I \subset \mathbb{N}$ for which $\mathbb{N} \smallsetminus I$ is infinite, there is a point $(x, \sqrt{1-x^2}, 0, 0) \in X$ that is not covered by $\cup_{i \in I} \hat{C}_i$.

To prove this statement when I is *infinite*, define the i^{th} digit of x as 1 if and only if $i \notin I$. Since this is only one binary representation of x, we have $x \notin \cup_{i \in I} \hat{C}_i^*$ and $(x, \sqrt{1-x^2}, 0, 0) \notin \cup_{i \in I} \hat{C}_i$. If I is *finite*, it can be extended to an infinite set such that $\mathbb{N} \smallsetminus I$ remains infinite. Thus, this case can be reduced to the case when I is infinite. $\square$

8 Bounded coverings

We prove Theorem 5 in a somewhat more general form. For the proof we need the following consequence of the Lovász local lemma.

Lemma 8.1 (Erdős-Lovász [10])**.** *Let $k, m \geq 2$ be integers. If every edge of a hypergraph has at least m vertices and every edge intersects at most $k^{m-1}/4(k-1)^m$ other edges, then its vertices can be colored with k colors so that every edge contains at least one vertex of each color.*

Let $\mathcal{C}$ be a class of subsets of $\mathbb{R}^d$. Given n members $C_1, \ldots, C_n$ of $\mathcal{C}$, assign to each point $x \in \mathbb{R}^d$ a *characteristic vector* $c(x) = (c_1(x), \ldots, c_n(x))$, where $c_i(x) = 1$ if $x \in C_i$ and $c_i(x) = 0$ otherwise. The number of distinct characteristic vectors shows how many "pieces" $C_1, \ldots, C_n$ cut the space into. The *dual shatter function* of $\mathcal{C}$, denoted by $\pi_{\mathcal{C}}^*(n)$, is the maximum of this quantity over all n-tuples $C_1, \ldots, C_n \in \mathcal{C}$. For example, when $\mathcal{C}$ is the family of open *balls* in $\mathbb{R}^d$, it is well known that

$$\pi_{\mathcal{C}}^*(n) \leq \binom{n-1}{d} + \sum_{i=0}^{d} \binom{n}{i} \leq n^d, \tag{1}$$

provided that $2 \leq d \leq n$.

Theorem 8.2. *Let $\mathcal{C}$ be a class of open sets in $\mathbb{R}^d$ with diameter at most D and volume at least v. Let $\pi(n) = \pi_{\mathcal{C}}^*(n)$ denote the dual shatter function of $\mathcal{C}$, and let B^d denote the unit ball in $\mathbb{R}^d$. Then, for every positive integer m, any m-fold covering of $\mathbb{R}^d$ with members of $\mathcal{C}$ splits into two coverings, provided that no point of the space is covered more than $\frac{v}{(2D)^d Vol B^d} \pi^{-1}(2^{m-3})$ times, where $Vol B^d$ is the volume of B^d.*

Proof. Given an m-fold covering of $\mathbb{R}^d$ in which no point is covered more than M times, define a hypergraph $\mathcal{H} = (V, E)$, as follows. Let V consist of all members of $\mathcal{C}$ that participate in the covering. To each point $x \in \mathbb{R}^d$, assign a (hyper)edge $e(x)$: the set of all members of the covering that contain x. (Every edge is counted only once.) Since every point x is covered by at least m members of $\mathcal{C}$, every edge $e(x) \in E$ consists of at least m points.

Consider two edges $e(x), e(y) \in E$ with $e(x) \cap e(y) \neq \emptyset$. Then there is a member of $\mathcal{C}$ that contains both x and y, so that y must lie in the ball $B(x, D)$ of radius D around x. Hence, all members of the covering that contain y lie in the ball $B(x, 2D)$ of radius $2D$ around x. Since the volume of each of these members is at least v, and no point of $B(x, 2D)$ is covered more than M times, we obtain that $B(x, D)$ can be intersected by at most $M Vol B(x, 2D)/v = M(2D)^d Vol B^d/v$ members of the covering. By the definition

19

of the dual shatter functions, those members of the covering that intersect $B(x, D)$ cut $B(x, D)$ into at most $\pi(M(2D)^d VolB^d/v)$ pieces, each of which corresponds to an edge of $\mathcal{H}$. Therefore, for the maximum number N of edges of $\mathcal{H}$ that can intersect the same edge $e(x) \in E$, we have

$$N \le \pi(M(2D)^d VolB^d/v).$$

According to Lemma 8.1 (for $k = 2$), in order to show that the covering can be split into two, i.e., the hypergraph $\mathcal{H}$ is 2-colorable, it is sufficient to assume that $N \le 2^{m-3}$. Comparing this with the previous inequality, the result follows. $\qquad\square$

In the special case where $\mathcal{C}$ is the class of unit balls in $\mathbb{R}^d$, we have $v = VolB^d$, $D = 2$, and, in view of (1), $\pi^{-1}(z) \ge z^{1/d}$. Thus, we obtain Theorem 5 with $c_d = 2^{-2d-3/d}$.

If we want to decompose an m-fold covering into $k > 2$ coverings, then the above argument shows that it is sufficient to assume that

$$\pi(M(2D)^d VolB^d/v) \le k^{m-1}/4(k-1)^m.$$

In case of unit balls, this holds for $M \le c_{k,d}(1 + \frac{1}{k-1})^{m/d}$ with $c_{k,d} = k^{-1/d}4^{-d-1/d}$.

Two sets are *homothets* of each other if one can be obtained from the other by a dilation with positive coefficient followed by a translation. It is easy to see [20] that for $d = 2$, the dual shatter function of the class $\mathcal{C}$ consisting of all homothets of a fixed convex set C is at most $n^2 - n + 2 \le n^2$, for every $n \ge 2$. In this case, Theorem 8.2 immediately implies

Corollary 8.3. *Every m-fold covering $\mathcal{C}$ of the plane with homothets of a fixed convex set can be decomposed into two coverings, provided that no point of the plane belongs to more than $2^{(m-11)/2}$ members of $\mathcal{C}$.*

Naszódi and Taschuk [30] constructed a convex set C in $\mathbb{R}^3$ such that the dual shatter function of the class of all translates of C cannot be bounded from above by any polynomial of n. Therefore, for translates of C, the above approach breaks down. We do not know how to generalize Theorem 5 from balls to arbitrary convex bodies in $\mathbb{R}^d$, for $d \ge 3$.

For some related combinatorial results, see Bollobás et al. [5].

9 Open problems and concluding remarks

Theorem 3 states that, if C is a plane convex body with two antipodal points at which the curvature is positive, then for every m, there exists an m-fold covering of $\mathbb{R}^2$ with translates of C that does not split into two coverings. We also know that this statement is false for any convex polygon. But what happens if C "almost satisfies" the condition concerning the antipodal point pair?

Problem 9.1. *Does there exist an integer m such that every m-fold covering of $\mathbb{R}^2$ with translates of an open semidisk splits into two coverings?*

Another question, which surprisingly is widely open even in a completely abstract setting, is the following.

Problem 9.2. *Suppose that for a body C, there is an integer m such that every m-fold covering of $\mathbb{R}^d$ with translates of C splits into two coverings. Does it follow that for every $k > 2$, there is an integer m_k such that every m_k-fold covering of $\mathbb{R}^d$ with translates of C splits into k coverings? Is it true that (for the smallest such m_k) even $m_k = O_C(k)$?*

According to Theorem 6.3, for any $m \geq 3$, every m-uniform special shift-chain is 2-colorable. Keszegh and the Pálvölgyi [25] recently extended this theorem to show that the vertices of every $(2k-1)$-uniform special shift-chain can be colored by k colors so that every hyperedge contains at least one point of each color.

As was stated in the introduction, for every triangle (in fact, for every convex polygon) C, there is an integer $m(C)$ such that every m-fold covering of the plane with *translates* of C splits into two coverings. Keszegh and the Pálvölgyi [22] extended this theorem to m-fold coverings with *homothets* of a triangle. (Two sets are homothets of each other if one can be obtained from the other by a dilation with positive coefficient followed by a translation.) Using the idea of the proof of our Theorem 2, Kovács [26] has recently showed that the analogous statement is false for homothets of any convex polygon with more than 3 sides. For further results about decomposition of multiple coverings, see [4, 5, 7, 8, 19, 23, 24].

Acknowledgment

References

[1] N. Alon, A non-linear lower bound for planar epsilon-nets, Discrete Comput. Geom. **47** (2012), no. 2, 235–244.

[2] N. Alon and J. H. Spencer, *The probabilistic method*, third edition, Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley, Hoboken, NJ, 2008.

[3] G. Aloupis, J. Cardinal, S. Collette, S. Langerman, D. Orden, and P. Ramos, Decomposition of multiple coverings into more parts, Discrete Comput. Geom. **44** (2010), no. 3, 706–723.

[4] A. Asinowski, J. Cardinal, N. Cohen, S. Collette, T. Hackl, M. Hoffmann, K. Knauer, S. Langerman, M. Lason, P. Micek, G. Rote, and T. Ueckerdt, Coloring hypergraphs induced by dynamic point sets and bottomless rectangles, in *Algorithms and data structures*, 73–84, Lecture Notes in Comput. Sci., 8037, Springer, Heidelberg, 2013.

[5] B. Bollobás, D. Pritchard, T. Rothvoß, and A. Scott, Cover-decomposition and polychromatic numbers, SIAM J. Discrete Math. **27** (2013), no. 1, 240–256.

[6] A. L. Buchsbaum, A. Efrat, S. Jain, S. Venkatasubramanian, and K. Yi, Restricted strip covering and the sensor cover problem, in: *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2007)*, 1056–1063, ACM, New York, 2007.

[7] J. Cardinal, K. Knauer, P. Micek, and T. Ueckerdt, Making triangles colorful, J. Comput. Geom. **4** (2013), no. 1, 240–246.

[8] J. Cardinal, K. Knauer, P. Micek, and T. Ueckerdt, Making octants colorful and related covering decomposition problems, SIAM J. Discrete Math. **28** (2014), no. 4, 1948–1959.

[9] P. Erdős, On a combinatorial problem, Nordisk Mat. Tidskr. **11** (1963), 5–10, 40.

[10] P. Erdős and L. Lovász, Problems and results on 3-chromatic hypergraphs and some related questions, in *Infinite and finite sets (Colloq., Keszthely, 1973; dedicated to P. Erdős on his 60th birthday), Vol. II*, 609–627. Colloq. Math. Soc. János Bolyai, 10, North-Holland, Amsterdam.

[11] P. Erdős and G. Szekeres, A combinatorial problem in geometry, Compositio Math. **2** (1935), 463–470.

[12] G. Even, Z. Lotker, D. Ron, and S. Smorodinsky, Conflict-free colorings of simple geometric regions with applications to frequency assignment in cellular networks, SIAM J. Comput. **33** (2003), no. 1, 94–136.

[13] U. Feige, M. M. Halldórsson, and G. Kortsarz, Approximating the domatic number, SIAM J. Comput. **32** (2002/03), no. 1, 172–195.

[14] G. Fejes Tóth, New results in the theory of packing and covering, in *Convexity and its applications*, 318–359, Birkhäuser, Basel.

[15] G. Fejes Tóth and W. Kuperberg, A survey of recent results in the theory of packing and covering, in *New trends in discrete and computational geometry*, 251–279, Algorithms Combin., 10, Springer, Berlin.

[16] R. Fulek, personal communication, 2010. See also in [36].

[17] R. Fulek, T. Hubai, B. Keszegh, Z. Nagy, T. Rothvoß, and M. Vizer, unpublished.

[18] H. Gebauer, H. Gebauer, Disproof of the neighborhood conjecture with implications to SAT, Combinatorica **32** (2012), no. 5, 573–587.

[19] M. Gibson and K. Varadarajan, Optimally decomposing coverings with translates of a convex polygon, Discrete Comput. Geom. **46** (2011), no. 2, 313–333.

[20] B. Grünbaum, Venn diagrams and independent families of sets, Math. Mag. **48** (1975), 12–23.

[21] D. Haussler and E. Welzl, ϵ-nets and simplex range queries, Discrete Comput. Geom. **2** (1987), no. 2, 127–151.

[22] B. Keszegh and D. Pálvölgyi, Octants are cover-decomposable, Discrete Comput. Geom. **47** (2012), no. 3, 598–609.

[23] B. Keszegh and D. Pálvölgyi, Octants are cover-decomposable into many coverings, Comput. Geom. **47** (2014), no. 5, 585–588.

[24] B. Keszegh and D. Pálvölgyi, Convex polygons are self-coverable, Discrete Comput. Geom. **51** (2014), no. 4, 885–895.

[25] B. Keszegh and D. Pálvölgyi, An abstract approach to polychromatic coloring: shallow hitting sets in ABA-free hypergraphs and pseudohalfplanes, arXiv:1410.0258.

[26] I. Kovács, Indecomposable coverings with homothetic polygons, arXiv:1312.4597.

[27] P. Mani-Levitska and J. Pach, Decomposition problems for multiple coverings with unit balls, manuscript, 1986. Parts of the manuscript are available at
`http://www.math.nyu.edu/~pach/publications/unsplittable.pdf`

[28] J. Matoušek, The determinant bound for discrepancy is almost tight, *Proc. Amer. Math. Soc.* **141** (2013), no. 2, 451–460.

[29] E. W. Miller, On a property of families of sets, C. R. Soc. Sci. Varsovie **30** (1937), 31–38.

[30] M. Naszódi and S. Taschuk, On the transversal number and VC-dimension of families of positive homothets of a convex body, Discrete Math. **310** (2010), no. 1, 77–82.

[31] J. Pach, Decomposition of multiple packing and covering, *Diskrete Geometrie, 2. Kolloq.* Math. Inst. Univ. Salzburg, 1980, 169–178.

[32] J. Pach, Covering the plane with convex polygons, Discrete Comput. Geom. **1** (1986), no. 1, 73–81.

[33] J. Pach, D. Pálvölgyi and G. Tóth, Survey on decomposition of multiple coverings, in *Geometry—intuitive, discrete, and convex*, 219–257, Bolyai Soc. Math. Stud., 24, János Bolyai Math. Soc., Budapest, 2013.

[34] J. Pach, G. Tardos and G. Tóth, Indecomposable coverings, Canad. Math. Bull. **52** (2009), no. 3, 451–463.

[35] J. Pach and G. Tóth, Decomposition of multiple coverings into many parts, Comput. Geom. **42** (2009), no. 2, 127–133.

[36] D. Pálvölgyi, *Decomposition of geometric set systems and graphs*, PhD thesis, EPFL, Lausanne, 2010, arXiv:1009.4641.

[37] D. Pálvölgyi, Indecomposable coverings with concave polygons, Discrete Comput. Geom. **44** (2010), no. 3, 577–588.

[38] D. Pálvölgyi and G. Tóth, Convex polygons are cover-decomposable, Discrete Comput. Geom. **43** (2010), no. 3, 483–496.

[39] J. Radhakrishnan and A. Srinivasan, Improved bounds and algorithms for hypergraph 2-coloring. Random Structures Algorithms 16 (2000), no. 1, 4–32.

[40] G. Tardos and G. Tóth, Multiple coverings of the plane with triangles, Discrete Comput. Geom. **38** (2007), no. 2, 443–450.

[41] K. Varadarajan, Weighted geometric set cover via quasi-uniform sampling, in *STOC'10—Proceedings of the 2010 ACM International Symposium on Theory of Computing*, 641–647, ACM, New York.

[42] P. Winkler, on page 137 of *Mathematical mind-benders*, A K Peters, Wellesley, MA, 2007. See also: P. Winkler, Puzzled: covering the plane, Commun. ACM **52** (2009), no. 11, 112.