

Survey on Decompositions of Multiple Coverings

János Pach, Dömötör Pálvölgyi, Géza Tóth

February 12, 2012

Abstract

The study of multiple coverings was initiated by Davenport and L. Fejes Tóth more than 50 years ago. In 1980 and 1986, the first named author published the first papers about decomposability of multiple coverings. It was discovered much later that, besides its theoretical interest, this area has practical applications to sensor networks. Now there is a lot of activity in this field with several breakthrough results, although, many basic questions are still unsolved. In this survey, we outline the most important results, methods, and questions.

1 Cover-decomposability and the sensor cover problem

Let $\mathcal{P} = \{ P_i \mid i \in I \}$ be a collection of sets in \mathbb{R}^d . We say that \mathcal{P} is an *m-fold covering* if every point of \mathbb{R}^d is contained in at least m members of \mathcal{P} . The largest such m is called the *thickness* of the covering. A 1-fold covering is simply called a *covering*. To formulate the central question of this survey succinctly, we need a definition.

Definition 1.1. *A planar set P is said to be cover-decomposable if there exists a (minimal) constant $m = m(P)$ such that every m -fold covering of the plane with translates of P can be decomposed into two coverings.*

Note that the above term is slightly misleading: we decompose (partition) not the set P , but a collection \mathcal{P} of its translates. Such a partition is sometimes regarded a *coloring* of the members of \mathcal{P} . The problem whether a set P is cover-decomposable is also referred to as the *cover decomposability problem* for P .

The problem of characterizing all cover-decomposable sets in the plane was proposed by Pach [P80] in 1980. He made the following conjecture, which is still unsolved.

Conjecture 1.2. [P80] *Every plane convex set P is cover-decomposable.*

In the present survey, we concentrate on results and proof techniques related to this conjecture. Obviously, in addition to systems of *translates* of a set P , we could study the analogous questions

for systems of *homothets* of P (that is, similar copies in parallel position) or for systems of congruent copies.

In [P86], Conjecture 1.2 was proved for open centrally symmetric convex polygons. More than twenty years later the proof was extended by Tardos and Tóth [TT07] to open triangles and then by Pálvölgyi and Tóth [PT10] to any open convex polygon P . Sections 2 and 3 describe the basic ideas and techniques utilized in these proofs.

Theorem 1.3. [P86] *Every centrally symmetric open convex polygon is cover-decomposable.*

Theorem 1.4. [TT07] *Every open triangle is cover-decomposable.*

Theorem 1.5. [PT10] *Every open convex polygon is cover-decomposable.*

In fact, the proof gives a slightly stronger result: any set, which is the union of finitely many translates of the same open convex polygon is also cover-decomposable. See Section 4 for details.

Given a cover-decomposable set P , one can try to determine the exact value of $m(P)$, that is, the smallest integer m for which every m -fold covering of the plane with translates of P splits into 2 coverings (cf. Definition 1.1). For example, for any open triangle T , we have $12 \geq m(T) \geq 4$ [KP11]. However, in most of the cases, the best known upper and lower bounds are very far from each other.

One can further generalize the cover decomposability problem by asking whether a sufficiently thick multiple covering of the plane can be decomposed into k coverings, for a fixed $k \geq 2$. This question was raised in [P86], and first addressed in detail in [PT07].

Definition 1.6. *Given a set $P \subset \mathbb{R}^2$ and an integer $k \geq 2$, let $m_k(P)$ denote the smallest positive number m with the property that every m -fold covering of the plane with translates of P can be decomposed into k coverings. If such an integer m does not exist, we set $m_k(P) = \infty$.*

We believe that $m_k(P)$ is finite for every cover-decomposable set P , but we cannot verify this conjecture in its full generality. However, the statement is true for all currently known families of cover-decomposable sets. In [P86], it was shown that, for any centrally symmetric convex open polygon P , the parameter $m_k(P)$ exists and is bounded by an exponentially fast growing function of k . In [TT07], a similar result was established for open triangles, and in [PT10] for open convex polygons. However, all these results were improved to the optimal linear bound in a series of papers by Pach and Tóth [PT07], Aloupis et al. [A110], and Gibson and Varadarajan [GV11].

Theorem 1.7. [PT07] *For any open centrally symmetric convex polygon P , we have $m_k(P) = O(k^2)$.*

Theorem 1.8. [A110] *For any open centrally symmetric convex polygon P , we have $m_k(P) = O(k)$.*

Theorem 1.9. [GV11] *For any open convex polygon P , we have $m_k(P) = O(k)$.*

The problem of determining $m_k(P)$ can be reformulated in a slightly different way: for a given m , try to decompose an m -fold covering into as many coverings as possible. This problem, more precisely, a slight generalization of this problem, is called the *sensor cover problem* in sensor network

scheduling. Suppose that we have a finite number of sensors scattered in a region R , each monitoring some part of R , which is called the range of the sensor. Each sensor has a duration for which it can be active and once it is turned on, it has to remain active until this duration is over, after which it will stay inactive. The *load* of a point is the sum of the durations of all ranges that contain it, and the load of the arrangement of sensors is the minimum load of the points of R . A *schedule* for the sensors is a starting time for each sensor that determines when it starts to be active. The goal is to find a schedule to monitor the given area, R , for as long as we can. Clearly, the cover decomposability problem is a special case of the sensor cover problem, when the duration of each sensor is the same (“unit” time).

Buchsbaum et al. [B07] and Gibson and Varadarajan [GV11] proved their result in this more general context. It was shown in [GV11] that for every open convex polygon P , there is a constant $c(P)$ such that for any instance of the sensor cover problem with load $c(P)k$, where the range of each sensor is a translate of P , there is a polynomial time computable schedule such that every point is monitored for k units of time.

Conjecture 1.2 cannot be extended to all (not necessarily convex) polygons.

Theorem 1.10. [PTT05] *No concave quadrilateral is cover-decomposable.*

In Section 5, following [PTT05] and [P10], we describe a large class of concave polygons that are not cover-decomposable.

The definition of cover-decomposability can be extended to higher dimensions in a natural way. It is interesting to note that most of the ideas presented in this survey fail to generalize to higher dimensions. The main reason for this is that statement analogous to Conjecture 1.2 is false in higher dimensions.

Theorem 1.11. [MP86] *For $d \geq 3$, the unit ball in \mathbb{R}^d is not cover-decomposable.*

Theorem 1.12. [P10] *For $d \geq 3$, no convex polytope is cover-decomposable.*

However, there is a notable exception in 3-dimensions, albeit unbounded: the octant $\{(x, y, z) : x, y, z > 0\}$.

Theorem 1.13. [KP11] *The open octant in 3-dimensional space is cover-decomposable.*

As an easy consequence, we obtain that there is an integer m such that any m -fold covering of the plane with homothets of a fixed triangle can be split into two coverings. The statement holds with $m = 12$. In fact, Conjecture 1.2 can be (and was) formulated in the following more general form.

Conjecture 1.14. *For every plane convex set P , there exists a positive integer $m = m(P)$ such that any m -fold covering of the plane with homothets of P can be split into two coverings.*

The methods developed in the first substantial publication in this topic [P86] were used in all later papers. Therefore, in the next two sections we concentrate on this paper and sketch the proof of Theorem 1.3. In Subsections 3.2 and 3.3, we establish Theorems 1.7 and 1.4. In Section 4, we outline the proofs of Theorem 1.5 and Theorem 1.9 for triangles. Section 5 contains constructions proving (an extension of) Theorem 1.10 and Theorem 1.12. We close this paper with some open problems.

2 Basic Tricks

A family of sets \mathcal{P} is called *locally finite* if every point is contained in only finitely many members of \mathcal{P} . It follows by a standard compactness argument that any m -fold covering of the plane with translates of an open polygon P has a locally finite subfamily that forms an m -fold covering. Therefore, in the sequel we will assume without loss of generality that all coverings that we consider are locally finite.

In the next three subsections, we describe three basic tricks from [P86] that enable us to reduce the cover decomposability problem to a finite combinatorial problem for hypergraphs.

2.1 Dualization

Let $\mathcal{P} = \{ P_i \mid i \in I \}$ be a collection of translates of a finite polygon P in the plane, where I is a finite or infinite set. Let O_i denote the center of gravity of P_i . Obviously, \mathcal{P} is an m -fold covering of the plane if and only if every translate of \bar{P} , the reflection of P through the origin, contains at least m elements of the point set $\mathcal{O} = \{ O_i \mid i \in I \}$. Furthermore, $\mathcal{P} = \{ P_i \mid i \in I \}$ can be decomposed into two coverings if and only if the point set $\mathcal{O} = \{ O_i \mid i \in I \}$ can be colored with two colors such that every translate of \bar{P} contains at least one point of each color.

Clearly, the reflected polygon \bar{P} is cover-decomposable if and only if P is. Therefore, we have the following.

Lemma 2.1. *The polygon P is cover-decomposable if and only if there exists an integer m satisfying the following condition. Any point set S in the plane with the property that every translate of P contains at least m elements of S can be colored with two colors so that every translate of P contains at least one point of each color.*

The same argument applies if we want to decompose a covering into $k > 2$ coverings. Almost all later papers in the subject follow this “dual” approach. In the sequel, we also study this version of the problem.

2.2 Divide and conquer—Reduction to wedges

The second trick from [P86] is to cut the plane and the set S in Lemma 2.1 into small regions so that with respect to each of them every translate of our polygon looks like an infinite “wedge”.

We use the following terminology. Two half-lines (rays) emanating from the same point O divide the plane into two connected pieces, called *wedges*. A *closed wedge* contains its boundary, an *open wedge* does not. The point O is called the *apex* of the wedge. The *angle* of a wedge is the angle between its two boundary half-lines, measured inside the wedge.

Let P be an open or closed polygon of n vertices. Consider a multiple covering of the plane with translates of P . Then, the cover decomposition problem can be reduced to wedges as follows. Divide the plane into small regions, say squares, so that each of them intersects at most two consecutive

sides of any translate of P . Every translate of P can intersect only a bounded number c of squares. If a translate of P contains at least cm points of a set S , then at least m of those will belong to one of the squares. Therefore, to find a coloring of the points of S meeting the requirements in Lemma 2.1, it is sufficient to focus on a fixed subset of $S' \subset S$, consisting of all points of S that lie in a single square. It is sufficient to 2-color the elements of S' so that no translate of P that covers at least m points of S' is monochromatic. Notice that, because of our assumption of local finiteness, each subset S' is finite. Moreover, from the point of view of S' any translate of P “looks like” a half-space or a wedge corresponding to one of the vertices of P . To make this statement more precise, denote by v_1, \dots, v_n the vertices of P in cyclic order, and denote by W_i the wedge bounded by the rays $\overrightarrow{v_i v_{i-1}}$ and $\overrightarrow{v_i v_{i+1}}$ which contains a piece of P in any small neighborhood of v_i . (The indices are taken mod n .) Now any subset of S' that can be cut off from S by a translate of P can also be cut off by a translate of one of W_1, \dots, W_n .

Lemma 2.2. *Suppose that there is a positive integer m such that any finite point set S can be colored with two colors such that every translate of any wedge W_i of P that contains at least m elements of S , contains points of both colors. Then P is cover-decomposable.*

A straightforward generalization of the above argument can be applied when we want to decompose a covering into $k \geq 2$ coverings. Thus, from now on, to prove positive cover-decomposability results we will try to find colorings of finite point sets. However, it will turn out that coloring point sets with respect to wedges may also be very useful in proving negative results.

Observe that we can assume without loss of generality that our point set S is in general position with respect to P , that is, none of the lines determined by two points of S is parallel to a side of P . Indeed, if there is such a line, we can slightly perturb the point set such that any subset of S that can be cut off from S by a translate of P , can also be cut off from the perturbed point set S' .

2.3 Totalitarianism

So far we have only considered coverings of the *whole* plane. At this point it will be convenient to extend our definitions to coverings of subsets of the plane.

Definition 2.3. *A set P is said to be totally-cover-decomposable if there exists a (minimal) constant $m^T = m^T(P)$ such that every m^T -fold covering of any (!) point set in the plane with translates of P can be decomposed into two coverings. More generally, for any fixed $k \geq 2$, let $m_k^T(P)$ denote the smallest number m^T with the property that every m^T -fold covering of any planar point set with translates of P can be decomposed into k coverings.*

This notion was formally introduced only in [P10], but, in view of Lemma 2.2, all proofs in earlier papers also work for this stronger version of decomposability. To avoid confusion with this notion, sometimes we will call cover-decomposable sets *plane-cover-decomposable*. By definition, every totally-cover-decomposable set is also plane-cover-decomposable. On the other hand, there exist sets (perhaps even open polygons) that are plane-cover-decomposable but not totally-cover-decomposable.

For example, the disjoint union of a concave quadrilateral and a far enough half-plane is such a set. Using the notion of total decomposability, we obtain the following stronger version of Lemma 2.2.

Lemma 2.4. *A polygon P is totally-cover-decomposable if and only if there exists a positive integer m^T with the property that any finite point set S in the plane can be colored with two colors such that every translate of any wedge of P that contains at least m^T points of S , contains points of both colors.*

Note that if we want to show that a set P is not plane-cover-decomposable, then, using Lemma 2.4 with suitably chosen sets S , we can first show that it is not totally-cover-decomposable, and then we can add more points to S and apply Lemma 2.1. Of course, we have to be careful not to add any points to the translates that guarantee that P is not totally-cover-decomposable. This is the path followed in [MP86, PTT05] (and also in [P10], but there the point set S cannot always be extended). These constructions will be discussed in detail in Section 5.

3 Boundary methods

Let W be a wedge and s be a point in the plane. The translate of W with its apex at s is denoted by $W(s)$. More generally, given a *convex* wedge (whose angle is at most π) W , and points s_1, s_2, \dots, s_k , let $W(s_1, s_2, \dots, s_k)$ denote the *minimal* translate of W (for containment) whose closure contains s_1, s_2, \dots, s_k .

Following [P86], next we will define the boundary of a finite point set with respect to a collection of wedges. We establish and explore some basic combinatorial and geometric properties of the boundary, which will be the heart of the proofs of Theorems 1.3, 1.7, and 1.4. The details of these three proofs from [P86], [PT07], and [TT07], respectively, will be sketched in the next three subsections.

3.1 Decomposition into two parts

In this subsection, we outline the proof of Theorem 1.3 in the special case when P is an axis-parallel square. This square has an *upper-left*, a *lower-left*, an *upper-right*, and a *lower-right* vertex. For each vertex v of the square, there is a corresponding convex wedge, whose apex is at v and whose boundary half-lines contain the sides of the square incident to v . Denote these wedges by W_{ul} , W_{ll} , W_{ur} , and W_{lr} , respectively. We refer to these four wedges as *P -wedges*.

Let S be a finite point set. By Lemma 2.2 it is sufficient to prove the following.

Lemma 3.1. *The set S can be colored with two colors such that any translate of a P -wedge which contains at least five points of S , contains points of both colors.*

At this point, we introduce the notion of the *boundary* of S with respect to the wedges of P . This notion will be similar to that of the boundary of the convex hull. A point s of S belongs to the boundary of the convex hull of S if there is a half-plane which contains s on its boundary, but none of the points of S in its interior. Similarly, a point s of S belongs to the boundary with respect to wedge W if $W(s)$ contains none of the points of S .

Definition 3.2. Let W be an open wedge. The W -boundary of S , that is, the boundary of S with respect to W is defined as $Bd^W(S) = \{s \in S : W(s) \cap S = \emptyset\}$. Two vertices, s and t , of the W -boundary are called neighbors if $W(s, t) \cap S = \emptyset$.

Obviously, one can define a natural ordering on the W -boundary points of S , according to which two vertices are consecutive if and only if they are neighbors. Observe that any translate of W intersects the W -boundary in an interval with respect to this ordering. The boundary of S with respect to the four P -wedges is the union of the W_{ul} -boundary, the W_{ll} -boundary, the W_{ur} -boundary, and the W_{lr} -boundary of S . All points of S that are not boundary vertices with respect to the P -wedges are called interior points.

The W_{lr} -boundary and the W_{ll} -boundary of S meet at the “highest” point of S , that is, at the point of maximum y -coordinate. (Assume, for simplicity that this point is unique). The W_{ll} -boundary and the W_{ul} -boundary meet at the rightmost point of S ; the W_{ul} -boundary and the W_{ur} -boundary meet at the lowest point; and the W_{ur} -boundary and the W_{lr} -boundary meet at the leftmost point. See Figure 1. If it leads to no confusion, the translates of W_{ul} , W_{ll} , W_{ur} , W_{lr} will also be denoted by W_{ul} , W_{ll} , W_{ur} , W_{lr} .




Figure 1: The boundary of a point set.

If we link together the natural orderings of the boundary vertices of S corresponding to W_{ll} , W_{lr} , W_{ur} , and W_{ul} , in this cyclic order, then we obtain a counterclockwise cyclic enumeration of all boundary vertices. The main difference between the boundary of S with respect to P and the boundary of the convex hull of S is that in the cyclic enumeration of the boundary vertices some vertices may occur *twice*. These vertices are called *singular*, and all other vertices *regular*.

It is not hard to show, however, that no boundary vertex can appear *three* times in the cyclic enumeration. Moreover, all singular vertices must have the same type. In our case, all of them belong to both a W_{ul} and a W_{lr} , or all of them belong to a W_{ur} and a W_{ll} . This property generalizes to the case when P is any centrally symmetric convex polygon: all singular boundary vertices must belong to a pair of opposite P -wedges of the same type.

The most important observation is the following.

Observation 3.3. *If the intersection of S with a translate of some P -wedge, say, W_{ll} , is non-empty, then this set can be obtained as the union of three subsets:*

- (i) an interval of consecutive elements in the cyclic enumeration of all vertices of the boundary of S , which contains at least one point from the W_{ll} -boundary;*
- (ii) an interval of consecutive elements in the cyclic enumeration of all vertices of the boundary of S , which contains at least one point from the W_{ur} -boundary;*
- (iii) a set of interior points.*

Note that while the subset in (i) contains at least one element, those in (ii) and (iii) may be empty. Analogous statements hold for the other three wedges, and also for other symmetric polygons.

A first naive attempt to find a suitable coloring of S is to color all boundary vertices blue and all interior vertices red. Unfortunately, it is possible that there is a P -wedge that contains lots of boundary vertices and no interior vertex, so this coloring is not necessarily good.

Another naive attempt is to color the boundary vertices alternately red and blue. Apart from the obvious problems that the size of the boundary may be odd and that the singular vertices are repeated in the cyclic order, there is a more serious difficulty with this approach: the translate of a wedge may contain just one boundary vertex and lots of interior vertices. Consequently, we have to be careful when we color the interior vertices, which may lead to further complications.

It turns out that a “mixture” of the above two naive approaches will work.

Definition 3.4. *A boundary vertex $s \in S$ is called m -rich if there is a translate W of a P -wedge, such that s is the only W -boundary vertex in W , but W contains at least m points of S .¹*

This definition is used in different proofs with a different constant m , but when it leads to no confusion, we simply write “rich” instead of “ m -rich.” In this proof, “rich” means “5-rich,” thus a boundary vertex s is rich if there is a wedge that intersects the W -boundary in s and contains at least four other points.²

Our general coloring rule will be the following.

- (1) Rich boundary vertices are blue.

¹In [P86] and [PT07] a slightly different definition was used: there s was required to be the only vertex from the whole boundary (and not only from the W -boundary) in the translate of W . For symmetric polygons, both definitions work, but, for example, for triangles only the latter one does.

²Instead of $m = 5$, we could also choose $m = 4$ to define rich points in this proof. Only the last line of the argument would require a little more attention.

- (2) There are no two red neighbors along the boundary.
- (3) Color as many points red as possible, that is, let the set of red points $R \subset S$ be *maximal* under condition (1) and (2).

Note that from (3) we can deduce

- (4) Interior points are red.

A coloring that satisfies these conditions is called a *proper coloring*. The same point set may have many proper colorings. For centrally symmetric polygons, any proper coloring will be suitable for our purposes. In [P86], an explicit proper coloring is described.

Now we are ready to sketch the proof of Lemma 3.1. Suppose that S is colored properly and W is a translate of a P -wedge such that it contains at least five points of S . We can assume without loss of generality that W contains exactly five points of S . By Observation 3.3, W intersects the W -boundary of S in an interval.

First, we find a blue point in W . If the above interval contains just one point then this point is rich, as the wedge contains at least five points, and rich points are blue according to (1). If the interval contains at least two points, then one of them should be blue, according to (2).

Now we show that W also has at least one red element. If W contains any interior point, then we are done, according to (4). Thus, we can assume by Observation 3.3 that $W \cap S$ is the union of two intervals and all points in W are blue. Since W has five points, at least one of them, say, x , is not the endpoint of any of the intervals. If x is not rich, then, according to (3), x or one of its neighbors is red. So, x must be rich. But then there is a translate W' of a P -wedge, W , or $-W$, which contains only x as a boundary vertex, and contains five points. Using that S is centrally symmetric, it can be shown that $S \cap W'$ is a proper subset of $S \cap W$, a contradiction, since both contain exactly five points. This concludes the proof of Lemma 3.1.

If we consider wedges with more points, we can guarantee more red points in them.

Lemma 3.5. *In a proper coloring of S , any translate of a P -wedge which contains at least $5i$ points of S contains at least one blue point and at least i red points ($i \geq 1$).*

The proof is very similar to the proof of Lemma 3.1. The difference is that now we color $5i$ -rich points red and we have to be a little more careful when counting red points, especially because of the possible singular points. If we delete the blue points (giving them color 1) and then recolor red points recursively by Lemma 3.5, we obtain an upper bound on $m_k(P)$, exponential in k . An analogous statement holds for any centrally symmetric open convex polygon, therefore, we have

Lemma 3.6. *For any centrally symmetric open convex polygon P , there is a constant c_P such that any c_P^k -fold covering of the plane with translates of P can be decomposed into k coverings.*

3.2 Decomposition to $\Omega(\sqrt{m})$ parts for symmetric polygons

Here we sketch the proof of Theorem 1.7, which is a modification of the argument described in the previous subsection. We continue to assume for simplicity that P is an axis-parallel square. Let $k \geq 2$.

We will color the point set S with k colors such that any P -wedge that contains at least $m = 18k^2$ points has at least one point of each color. Recursively, we define k boundary layers and denote them by B_1, B_2, \dots, B_k . Let B_1 denote the boundary of S , and let $S_2 = S \setminus B_1$. For any $i < k$, if the set $S_i \subset S$ has already been defined, let B_i be the boundary of S_i and let $S_{i+1} = S \setminus B_i$. The coloring of the boundary layer B_i will be “responsible” for color i . Color i takes the role of blue from the previous proof, while those points that were colored red there will be “uniformly” distributed among the other $k - 1$ colors.

Slightly more precisely, a vertex $v \in B_i$ is called *rich* if there is a translate of a P -wedge that intersects S_i in at least $18k^2 - 18ki$ points, and v is the only element of B_i in it. We color all rich vertices of B_i with color i , and color first the remaining singular, then the remaining regular points periodically: $1, i, 2, i, 3, i, \dots, k, i, 1, i, \dots$. The main observation is that, if a P -wedge intersects B_i (for any i) in at least $18k$ points, then it contains a long interval that contains a point of each color. Otherwise, it has to intersect each of the boundary layers B_i ($1 \leq i \leq k$), but then for each i , its intersection with B_i contains a rich point of color i .

3.3 Triangles

The main difficulty with non-symmetric polygons is that Observation 3.3 does not hold here: the intersection with a translate of a P -wedge is not necessarily the union of two boundary intervals and some interior points. See Figure 2. In the case of triangles, Tardos and Tóth [TT07] managed to overcome this difficulty by defining a variant of proper colorings. In this subsection, we sketch their proof of Theorem 1.4. For other polygons, a different approach was needed (see later).

Suppose that P is a triangle with vertices A, B, C . There are three P -wedges, W_A, W_B , and W_C . We define the boundary just like before. It consists of three parts, the *A-boundary*, *B-boundary*, and *C-boundary*. Each of them forms an interval in the cyclic enumeration of the boundary vertices. Here comes the first difficulty: there may exist a singular boundary vertex which appears *three* times in the cyclic enumeration of boundary vertices, once in each boundary. It is easy to see that there exists at most *one* such vertex, and we can get rid of it by decomposing our point set S into at most four subsets such that in each of them all singular boundary points belong to the same pair of boundaries, just like in the case of centrally symmetric polygons. For simplicity of the explanation, assume that S has no singular boundary vertex.

Again, we call a boundary vertex s *rich* if there is a translate W of a P -wedge, such that s is the only W -boundary vertex in W , but W contains at least *five* elements of S .

Our coloring will satisfy the following four conditions.

- (1) Every rich boundary vertex is blue.
- (2) There are no two red neighbors.
- (3) Color as many points red as possible, that is, let the set of red points $R \subset S$ be *maximal* under condition (1) and (2).
- (4) All interior points are red.




Figure 2: Triangle P and the structure of the boundary.

We describe explicitly how to find the set of red points using a greedy algorithm. Consider the linear order on the set of all lines of the plane parallel to the side BC , so that the line through A is *smaller* than the line BC . We define a partial order $<_A$ on our point set as follows. Let $x <_A y$ if the line through x is smaller than the line through y . We have $A <_A B$ and $A <_A C$. Analogously, define the partial order $<_B$ with respect to the side AC such that $B <_B C$ and $B <_B A$, and the partial order $<_C$ with respect to the side AB such that $C <_C A$ and $C <_C B$.

First, color all rich boundary vertices blue. Now take all A -boundary vertices of S and consider them in *increasing* order with respect to $<_A$. If we get to a point that is not colored, we color it red and its neighbors blue. Note that these neighbors may have already been colored blue (because they are rich, or because of an earlier red neighbor), but they were definitely not colored red, as any neighbor of any red point is immediately colored blue. Continue this procedure, until all points of the A -boundary are colored. Color the B -boundary and the C -boundary in a similar fashion, using the other two partial orders.

Suppose that W is a translate of a P -wedge covering *at least* five points of S . We can assume without loss of generality that W covers exactly five points of S . Assume that W is a translate of W_A . The other two cases can be treated similarly. To find a blue point in W , we proceed just like in the previous section; this works for any proper coloring. We know that W intersects the A -boundary

of S in an interval. If this interval contains just one point, then it must be rich and hence blue. If the interval contains at least two points, then one of them must be blue.

It remains to show that W also contains at least one red point. If W contains any interior point, then we are done. Therefore, we assume that all five points in W are boundary vertices. Since there are five points in W , one of them, say, x , is (i) not the first or last A -boundary vertex in W ; (ii) not the $<_A$ -minimal B -boundary point in W ; and (iii) not the $<_A$ -minimal C -boundary point in W .

Suppose that x is rich. Then there is a translate W' of a P -wedge, which contains only x as a boundary vertex, and contains five points. It can be shown by some simple geometric arguments that $S \cap W'$ is a proper subset of $S \cap W$, a contradiction, since both sets contain five points. So, x cannot be rich. But then why would it be blue? The only possible reason is that during the coloring process, one of its neighbors on the boundary, say, y , was colored red earlier. It can be shown that then $y \in W$, which implies that there is a red point in W .

The same idea works if there are some singular boundary vertices, but all of them belong to the A -boundary and the B -boundary, say. The only difference is that in this case we have to synchronize the coloring processes on the A -boundary and on the B -boundary, so that we arrive at the common vertices at the same time. This concludes the proof of Theorem 1.4. The original proof gave that every 43-fold covering with translates of a triangle splits into two coverings, but B. Ács [A10] showed that the statement also holds for every 19-fold covering. Recently it was further improved to 12-fold coverings, by Keszegh and Pálvölgyi [KP11].

By a slightly more careful argument, we can establish

Lemma 3.7. *The points of S can be colored with red and blue such that any translate of a P -wedge which contains at least $5i + 3$ of the points, contains a blue point and at least i red points ($i \geq 1$).*

If we apply Lemma 3.7 recursively, we obtain an bound on $m_k(P)$, exponential in k .

Lemma 3.8. *For any open triangle P , every $\frac{7 \cdot 5^k - 15}{20}$ -fold covering of the plane with translates of P can be decomposed into k coverings.*

4 Path Decomposition and Level Curves

In this section, we present three generalizations of the boundary method that can be used to establish cover-decomposability results.

4.1 Classification of wedges

Pálvölgyi and Tóth [PT10] developed some new ideas to establish Theorem 1.5 which states that all open convex polygons are cover-decomposable. In the previous section, we colored a point set with respect to P -wedges, for some fixed polygon P . Here we color point sets with respect to an *arbitrary* set of wedges.

Definition. A collection of wedges $\mathcal{W} = \{ W_i \mid i \in I \}$ is said to be *non-conflicting* or, simply, *NC* if there is a constant m with the following property. Any finite set of points S in the plane can be colored with two colors so that any translate of a wedge $W \in \mathcal{W}$ that covers at least m points of S contains points of both colors.

It turns out that a single wedge is always NC. One can also characterize all pairs of wedges that are NC. Pálvölgyi and Tóth proved that a set of wedges is NC if and only if each pair is NC. It follows directly from this characterization that for any convex polygon P , the set of P -wedges is NC.

Lemma 4.1. *A single wedge is NC.*

An important tool in the proof of Lemma 4.1 and in the proof of the following lemmas is the *path decomposition*, which is a generalization of the concept of the boundary. To illustrate this technique, we present a proof of Lemma 4.1.




Figure 3: A concave wedge and three points, any two of which can be cut off by a translate.

Proof of Lemma 4.1. Let S be a finite point set and let W be a wedge. We prove that the NC property holds with $m = 3$, that is, we show that S can be colored with two colors such that any translate of W that contains at least 3 points of S , contains a point of both colors. Suppose first that the angle of W is at least π . Then W is the union of two half-planes, A and B . Take the translate of A (resp. B) that contains exactly two points of S , say, A_1 and A_2 (resp. B_1 and B_2). There might be coincidences between A_1, A_2 and B_1, B_2 , but still, we can color the set $\{A_1, A_2, B_1, B_2\}$ such that A_1 and A_2 (resp., B_1 and B_2) are of different colors. Now, if a translate of W contains three points, it contains either A_1 and A_2 , or B_1 and B_2 , and we are done. Note that three is optimal in this statement; see Figure 3.

Suppose now that the angle of W is less than π . We show that in this case the NC property holds with $m = 2$. We can assume that the positive x -axis is in W ; this can be achieved by an appropriate rotation. For simplicity, also suppose that no line determined by a pair of points of S is parallel to

the sides of W . This can be guaranteed by applying a suitable perturbation of the set S that does not effect which subsets of it can be cut off by a translate of W .

For any fixed y , let $W(2; y)$ be the translate of W which

- (1) contains at most two points of S ;
- (2) the y -coordinate of its apex is y ; and
- (3) the x -coordinate of its apex is minimal.




Figure 4: Z replaces X in $W(2; y)$.

For any y , the translate $W(2; y)$ is uniquely determined. Examine, how $W(2; y)$ varies as y runs over the real numbers. If y is very small (smaller than the y -coordinate of the points of S), then $W(2; y)$ contains two points, say X and Y , and one more, Z , on its boundary. As we increase y , the apex of $W(2; y)$ changes continuously. How can the set $\{X, Y\}$ of the two points in $W(2; y)$ change? For a certain value of y , one of them, say X , moves to the boundary. At this point, Y is inside and two points, X and Z , are on the boundary. If we slightly further increase y , then Z replaces X , that is, Y and Z will be in $W(2; y)$ (see Figure 4). As y increases to infinity, the set $\{Z, Y\}$ could change several times, but each time it changes in the above described manner. Define a directed graph whose vertices are the points of S , and there is an edge from u to v if v replaced u during the procedure. We get two paths, P_1 and P_2 . The pair (P_1, P_2) is called the *path decomposition of S with respect to W , of order two* (see Figure 5).

Color the vertices of P_1 red, the vertices of P_2 blue. Observe that each translate of W that contains at least two points, contains at least one vertex of both P_1 and P_2 . This completes the proof. \square

The *path decomposition of S with respect to W , of order m* can be defined very similarly. Let $W(m; y)$ denote the translate of W which

- (1) contains at most m points of S ;
- (2) the y -coordinate of its apex is y ; and
- (3) the x -coordinate of its apex is minimal.




Figure 5: Path decompositions of order two. $P_1 = X_1X_2\dots$, $P_2 = Y_1Y_2\dots$

Suppose that, for a very small value of y , the set $W(m; y)$ contains the points r_1, r_2, \dots, r_m , and at least one more point on its boundary. Just like in the proof above, as we increase the value of y , the set $\{r_1, r_2, \dots, r_m\}$ changes several times. Every time one of the elements of this set is *replaced* by another point. Define a directed graph whose vertices are the points of S , and there is an edge from r to s if r is replaced by s at some point. This graph is the union of m directed paths, $P_1^W, P_2^W, \dots, P_m^W$ (and possibly some isolated vertices), which is called the *order m path decomposition of S with respect to W* . Note that the order 1 path decomposition is just the W -boundary of S , so this notion can be regarded as a generalization of the boundary. In general, in a higher order path decomposition, no path is identical to the boundary. The union of the paths, however, always contains the boundary.

Note that there is a hidden variable in this notation. When we write P_1^W , then it can mean the first path of the path decomposition of *any* order m , so it would be more precise to write $P_1^W(m)$. However, to ease readability, we use the (ambiguous) simpler notation as from the context the value of m will be always clear.

Lemma 4.2. (i) *Any translate of W contains an interval of each of $P_1^W, P_2^W, \dots, P_m^W$.*

(ii) *If a translate of W contains precisely m points of S , then it contains precisely one point from each of $P_1^W, P_2^W, \dots, P_m^W$.*

Now we scrutinize the case when we have *two* wedges, V and W . We distinguish several cases according to their relative position.

Type 1 (Big): One of the wedges has angle at least π .

For the other cases, we can assume without loss of generality that W contains the positive x -axis. Extend the boundary half-lines of W to lines. They divide the plane into four parts: Upper, Lower,

Left, and Right parts, the last of which is W itself. See Figure 6.




Figure 6: Wedge W

Type 2 (Half-plane): One side of V is in the Right part and the other one is in the Left one. That is, the union of the wedges cover a half-plane. See Figure 7.




Figure 7: Type 2 (Half-plane)

Type 3 (Contain): One of the following three conditions is satisfied:
 (i) one side of V is in the Upper part, the other is in the Lower part;
 (ii) both sides are in the Right part;
 (iii) both sides are in the Left part. See Figure 8.




Figure 8: Type 3 (Contain)

Type 4 (Hard): One side of V is in the Left part and the other side is either in the Upper part or in the Lower one. See Figure 9.




Figure 9: Type 4 (Hard)

Type 5 (Special): One of the following three conditions is satisfied:

- (i) one side of V is in the Right part and the other one is in the Upper or Lower part;
- (ii) both sides of V are in the Upper part;
- (iii) both sides are in the Lower part.

That is, the union of the wedges is in an open half-plane whose boundary contains the origin, but neither of them contains the other. See Figure 10.




Figure 10: Type 5 (Special)

It is not hard to see that there are no other possibilities.

Lemma 4.3. *Let $\mathcal{W} = \{V, W\}$ be a set of two wedges, of Type 1, 2, 3, or 4. Then \mathcal{W} is NC.*

Here we omit the proof. It is different for each type, but in each case the basic idea is similar to that of the proof of Lemma 4.1. In the case of pairs of wedges of Type 4 (Hard), we have to take care of singular points in a somewhat similar way as in the previous section, in the proof for triangles. For pairs of wedges of Type 3 (Contain), we can apply an order 4 path decomposition.

Next, we turn to the case of several wedges.

Lemma 4.4. *A set of wedges $\mathcal{W} = \{W_1, W_2, \dots, W_t\}$ is NC if and only if any pair $\{W_i, W_j\}$ is NC.*

It is obvious that if two wedges are not NC, then \mathcal{W} cannot be NC. The proof in the other direction is more involved. It is based on a tricky application of path decompositions. In fact, it can be shown that if \mathcal{W} is NC, then for any k there is an m_k such that any finite point set can be colored with k colors such that if a translate of a wedge from \mathcal{W} contains at least m_k points, then

it contains all k colors. However, the bound obtained in [PT10] grows very fast, the argument gives only $m_k \leq (8k)^{2^{t-1}}$.

To finish the proof of Theorem 1.5, observe that no two wedges corresponding to the vertices of a convex polygon can form a pair of Type 1 (Big) or of Type 5 (Special).

It is shown in [P10] that if $\mathcal{W} = \{V, W\}$ is a set of two wedges of Type 5 (Special), then \mathcal{W} is *not* NC. Therefore, a set of wedges is NC if and only if none of its pairs is of Type 5 (Special). For the construction and its consequences, see Section 5.

4.2 Level curves and decomposition to $\Omega(m)$ parts for symmetric polygons

The *level curve* method, which can be regarded as another extension of the boundary technique, was invented by Aloupis, Cardinal, Collette, Langerman, Orden, and Ramos [Al10] at about the same time, but independently from the introduction of path decompositions.

Suppose that W is an open wedge and its angle is less than π . The *level curve of depth l* , denoted by $\mathcal{C}(l)$, is defined as the boundary of the union of all translates of W that contain fewer than l points. If W contains the positive x -axis, then we can also define $\mathcal{C}(l)$ as the set of the apices of $W(l-1; y)$.




Figure 11: The level curve $\mathcal{C}(l)$.

Note that this curve consists of straight-line segments that are parallel to the sides of W . See Figure 11. $\mathcal{C}(1)$ passes through all boundary points. If $p \in \mathcal{C}(l)$, then $|W(p) \cap S|$ is $l-1$, and $W(p)$ has one or two points of S on its boundary.

Consider all translates of W whose apices are on $\mathcal{C}(l)$. Call these translates $\mathcal{C}(l)$ - W -wedges. Consider a point of S in a $\mathcal{C}(l)$ - W -wedge. The apices of those $\mathcal{C}(l)$ - W -wedges which contain this point form an interval on $\mathcal{C}(l)$. Therefore, each $\mathcal{C}(l)$ - W -wedge corresponds to a point on $\mathcal{C}(l)$, and

every point of S corresponds to an interval of $\mathcal{C}(l)$. The condition that each $\mathcal{C}(l) - W$ -wedge contains at least $l - 1$ points translates to the condition that each of the points of $\mathcal{C}(l)$ is covered by at least $l - 1$ intervals. Here we want to color the *intervals* in such a way that each *point* is covered by intervals of all colors.

Now we sketch the proof of Theorem 1.8 given in [Al10], based on the level curve method.

Suppose that our symmetric polygon P has $2n$ vertices. Denote the wedges belonging to them by W_0, \dots, W_{2n-1} , in clockwise order. Throughout the proof, all the indices should be considered modulo $2n$. Two wedges, W_i and W_j , are called *antipodal* if $i + n \equiv j$ modulo $2n$, that is, if they belong to two opposite vertices of the polygon. A crucial observation, already used in Subsection 3.1 (more generally, in [P86]), is that any two wedges that are not antipodal cover a half-space.

We want to color the points of the point set S with k colors such that every translate of W_i ($i = 0, \dots, 2n - 1$) that contains at least m'_k points, contains a point of each color. For any fixed l , the level curves $\mathcal{C}_i(l)$ that correspond to wedge W_i may cross each other in a complicated way. However, in the “middle” of S they form a structure similar to the boundary in Subsection 3.1. It turns out that it is enough to consider these parts of the level curves.

More precisely, let $l = 6k + 1$. For every side of P , take two lines parallel to it that cut off $2l + 3$ points from each side of S . Denote the intersection of the n strips formed by these lines by \mathcal{T} . For every i , let $\mathcal{C}'_i(l) = \mathcal{C}_i(l) \cap \mathcal{T}$. Call those translates of W_i whose apices are on $\mathcal{C}'_i(l)$ *witness W_i -wedges*. It is not hard to see that only level curves belonging to antipodal wedges may cross inside \mathcal{T} . Some further analysis shows that, in fact, there can be only at most one such pair. (Note the similarity to singular points in case of symmetric polygons.) This means that the regions cut off from \mathcal{T} by the curves $\mathcal{C}'_i(l)$ are all disjoint, with the possible exception of one pair. Without loss of generality we may assume that this pair is $\mathcal{C}'_i(0)$ and $\mathcal{C}'_i(n)$. It is not difficult to verify that any translate of W_i that contains at least $3l + 5$ points, must contain a witness W_i -wedge. Therefore, it is enough to concentrate on wedges with this property.

One can parameterize these witness wedges by $t \in [0, 2n)$ such that $W(t)$ is a translate of $W_{\lfloor t \rfloor}$. The most important geometric observation is that if $p \in W(\lfloor t \rfloor + x) \cap W(\lfloor t \rfloor + z)$, where $0 \leq x \leq 1$ and $0 \leq z \leq n$, then $p \in W(\lfloor t \rfloor + y)$ for all $x \leq y \leq z$.

If $p \in W(\lfloor t \rfloor + x) \cap W(\lfloor t \rfloor + z)$, where $0 \leq x \leq 1$ and $n \leq z \leq n + 1$, then p is contained in two antipodal wedges, which implies that it is contained in translates of W_0 and W_n , but in no translates of any other wedge W_i . Therefore, every p corresponds to either an interval of the circle $[0, 2n)$ or to two intervals, one of which is a subinterval of $[0, 1]$, and the other a subinterval of $[n, n + 1]$.

We can take care of these two cases separately, as any big wedge contains many points from one of these groups. The sets of the first type (intervals) form a circular interval graph. Using a simple greedy algorithm, we can partition the set of these circular intervals into k parts with the property that any point of the circle that is covered by at least $3k$ intervals will be covered by at least one interval in each part. For sets of the second type (unions of two intervals), we want to color points with respect to a wedge W and its antipodal pair $-W$. The greedy algorithm again gives a good partition of a $3k$ -fold covering into k coverings. Since every witness wedge contains at least $6k$ points,

we are done.

Combining these facts, we obtain that $m'_k \leq 18k + 5$ for any system of wedges derived from a convex centrally symmetric polygon. This has to be multiplied by a constant depending on the shape of the polygon that comes from Lemma 2.2, to derive a bound for the multiple-cover-decomposability function m_k of the polygon.

4.3 Decomposition to $\Omega(m)$ parts for triangles

The case of not necessarily centrally symmetric polygons P was settled in [GV11]. In this subsection, we sketch the proof in the special case when P is a triangle, which already contains most of the key ideas of the general argument.

The first step is the usual dualization and reduction to wedges, therefore, it is enough to prove the following statement.

Lemma 4.5. *Let W_A, W_B, W_C be the wedges of a triangle T , and let $k > 0$. Then any point set S can be colored with k colors such that any translate of W_A, W_B , or W_C which covers at least $14k$ points of S contains at least one point in each color.*

Let S be a point set. Consider the level curve $\mathcal{C}_A = \mathcal{C}_A(14k + 1)$ of W_A of depth $14k + 1$. Again, for the coloring it is enough to consider those translates of W_A whose apices are on \mathcal{C}_A . As we have seen in the previous subsection, these wedges contain $14k$ points of S . Call these translates *witness A-wedges*. The *witness B-wedges* and *witness C-wedges* can be defined analogously.

The most important new idea is that first we *partially* color the points of S so that every witness A -wedge contains at least one point of each color, and all witness B -wedges and witness C -wedges have sufficiently many uncolored points. We proceed by extending this coloring in such a way that every witness B -wedge has a point of each color, and it is still true that every witness C -wedge has enough uncolored points. Finally, we take care of the witness C -wedges.

Lemma 4.6. *One can partially color the points of S with k colors such that*
(i) each witness A -wedge contains all k colors, and
(ii) each witness B -wedge and C -wedge contains at least $6k$ uncolored points.

Proof. We will again use the partial orders $<_A, <_B$, and $<_C$, defined in Subsection 3.3. First, we choose a subset $Q \subset S$ in the following way. Initially, set $Q = \emptyset$. Then, for each witness A -wedge W such that $|Q \cap W| < 2k$, we add the points of $S \cap W$ to Q , one by one, in decreasing order with respect to $<_A$, until $|Q \cap W| = 2k$. Then we proceed with another witness A -wedge. There are infinitely many witness A -wedges, but we have to consider only finitely many, since they can intersect S in only finitely many distinct subsets.

In the way described in the previous subsection, each witness A -wedge corresponds to a point on \mathcal{C}_A , and each point of Q corresponds to an interval. Thus, we obtain a system of intervals on \mathcal{C}_A (or, equivalently, on a line) such that each point is covered at least $2k$ times. Take a *minimal* collection of these intervals that still form a covering. It is easy to see that no point can be covered more than

twice. Color these intervals with the first color, take another minimal cover for the second color, and continue until all colors are used. Since we started with a $2k$ -fold covering and in each step the thickness decreased by at most two, we will be able to use all colors. This corresponds to a coloring of a subset $R \subset Q$. It is clear that each witness A -wedge contains at least one point of each color. Observe, that the intervals that correspond to R do not cover any point more than $2k$ times. That is, each witness A -wedge contains at most $2k$ points of R .

Now we prove (ii). By symmetry, it is enough to show that every witness B -wedge contains at least $6k$ uncolored points. Let W be a witness B -wedge, and let p_1, p_2, \dots be the points of $W \cap S$ in increasing order with respect to $<_B$. If none of them is in Q , then none of them is colored and we are done. Otherwise, let j be the largest number such that $p_j \in Q$. If $j < 8k$, then there are at least $6k$ uncolored points in W . Suppose that $j \geq 8k$. Point p_j was added to Q when we considered a certain witness A -wedge, say, V . Wedges W and V can have two types of intersection, since exactly one of them contains the apex of the other one.




Figure 12: The two types of intersections of W and V .

Case 1: V contains the apex of W . Consider the triangle $Z_1 = \{x \mid x \in W, p_j \not<_B x\}$. (See the left part of Figure 12.) It contains j points of S , but at most $2k$ of them are colored, so W contains at least $6k$ uncolored points.

Case 2: W contains the apex of V . Consider the triangle $Z_2 = \{x \mid x \in V, p_j \not<_A x\}$. (See the right part of Figure 12.) Since we added p_j to Q when we processed wedge V , there can be at most $2k - 1$ points p in V with $p_j <_A p$. Therefore, at least $12k$ points are in Z_2 . Since we colored at most $2k$ of them, there must remain at least $10k$ uncolored points in $Z_2 \subset W$. \square

Now we run the same algorithm for the uncolored points and for the witness B -wedges. A very similar argument shows that there will still be at least $2k$ uncolored points in each witness C -wedge. We run the algorithm once more for the uncolored points and for the witness C -wedges. This concludes the proof.

5 Indecomposable Coverings

In this section, we describe some constructions of coverings with arbitrarily high multiplicity that cannot be decomposed into two coverings. The first such example was given by Mani and Pach [MP86], and it shows that the unit ball is not cover-decomposable. In other words, for any m , there exists a covering of \mathbb{R}^3 with unit balls such that every point is covered by at least m balls, but the covering cannot be decomposed into two coverings. Later in [PTT05], several other constructions were given, all based on the geometric realization of the same m -uniform hypergraph (system of m -element sets) not having Property B.³ The same hypergraph is used in the construction described in Subsection 5.1 below. It was shown by Erdős [E63] that every m -uniform hypergraphs that does not have Property B has at least 2^{m-1} hyperedges, so any indecomposable construction must be exponentially large. As one of the first geometric applications of the Lovász Local Lemma [EL75], Pach showed that if a system of translates of a “nice” geometric set has the property that every point is covered by at least m and at most a subexponential (in m) number of sets, then the system is decomposable into two coverings.

First, we present the construction of [PTT05] showing that no concave quadrilateral is cover-decomposable. In Subsection 5.2 (see also [P10]), we show that concave polygons are not totally-cover-decomposable and polyhedra are not space-cover-decomposable. Finally, we discuss the difference between several variants of cover-decomposability.

5.1 Concave quadrilaterals—Proof of Theorem 1.10

We present the construction in the dual setting. Suppose that the vertices of the quadrilateral, Q , are A, B, C , and D , in this order, the reflex angle being at D . This implies that W_A and W_C are of Type 5 (Special) (see Section 4.1 for the definition). Moreover, they belong to an even more special subclass, which we call *Very Special*: when we translate the wedges so that their apices are in the origin, then they are disjoint and their closures are both contained in the same open half-plane (see the two right examples in Figure 10).

First, for any m , we give a finite set of points and a finite number of translates of Q , each covering precisely m points, such that no matter how we color the points by two colors, at least one of the translates will be monochromatic. In the “primal” setting, this corresponds to a finite system of translates of Q with the property that no matter how we partition this system into two, we can find a point contained in precisely m translates, each of which belongs to the same part. Hence, Q is not totally-cover-decomposable. Finally, we show how this construction can be extended to an m -fold covering of the whole plane, which cannot be split into two coverings.

We use translates of the wedges W_A and W_C to realize the following m -uniform hypergraph \mathcal{H} , also used in [MP86]. The vertices of the hypergraph are the vertices of a rooted perfect m -ary tree of height $m - 1$. There are two types of hyperedges. To each vertex v which is not a leaf of the

³We say that a hypergraph has Property B if the elements of its vertex set can be colored with two colors such that every hyperedge contains points of both colors.

tree, we assign a hyperedge of the *first type*, formed by the children of v . To each leaf v , we assign a hyperedge of the *second type*, formed by the vertices along the path from the root to v . More precisely, the vertices of the hypergraph are sequences of length less than m , consisting of the integers from 1 through m : $V(\mathcal{H}) = [m]^{<m} = \cup_{i=0}^{m-1} [m]^i$. The hyperedges of the first kind consist of m -tuples of sequences of length l , for some l ($1 \leq l < m$), such that removing their last elements, we obtain the same sequence of length $l - 1$. The hyperedges of the second kind consist of all initial segments of a sequence of length $m - 1$, where the empty sequence (corresponding to the root) is considered an initial segment of every sequence. Hence, \mathcal{H} has $\sum_{i=0}^{m-1} m^i$ vertices and $\sum_{l=1}^{m-1} m^{l-1} + m^{m-1}$ hyperedges.




Figure 13: Indecomposable covering with two special wedges of a concave quadrilateral.

The hyperedges of the first kind are realized by translates of W_A , the hyperedges of the second kind by translates of W_C . For simplicity, suppose that W_A is a very thin wedge that contains the positive x -axis and W_C is a very thin wedge that contains the negative x -axis; although the construction would work for any pair of convex wedges that belong to opposite vertices of a concave quadrilateral. All vertices of \mathcal{H} are very close to a vertical line. All vertices of a hyperedge of the first kind are on a horizontal line, for each edge on a different one (see Figure 13). It is easy to see that this is indeed a geometric realization of \mathcal{H} , so the points cannot be colored with two colors such that every translate of W_A and W_C of size m contains points of both colors.

Now we switch back to the primal plane. We have a point set S , and a set \mathcal{Q} of translates of Q . It remains to extend \mathcal{Q} to an m -fold covering of the whole plane. Before doing so, notice that it can be achieved that all points of the set S are on a line ℓ , not parallel to the sides of Q . Add to this m -fold covering all translates of Q that are disjoint from S (see Figure 14). It is clear that the resulting arrangement remains indecomposable. The construction can be easily modified to obtain a “locally finite” covering, using a standard compactness argument. Note that the construction of [P10] is not always extendable this way.




Figure 14: Extending the original 2-fold covering of the four points by the solid quadrilaterals to a 2-fold covering of the whole plane by adding the dotted quadrilaterals.

5.2 General concave polygons and polyhedra

The hypergraph \mathcal{H} can be realized by two wedges that form a Very Special pair. Unfortunately, there are concave polygons that do not have two Very Special wedges (see, e.g., Figure 16). In fact, they might not even have two wedges that form a Type 5 (Special) wedge at all; e.g., in the case of the union of two axis-aligned rectangles. The cover-decomposability of such concave polygons follows from the proof of Theorem 1.5 (see Lemma 2.2, 4.3, 4.4). However, it can be shown that every concave polygon that has two wedges that form a pair of Type 5 (Special) is not totally-cover-decomposable. This includes all “typical” concave polygons, as any polygon that has no parallel sides has a Type 5 (Special) pair of wedges.

To prove indecomposability, we have to realize another hypergraph that does not have property B. This construction has fewer points than \mathcal{H} (about 4^m). It is also more general, in the sense that it can be realized by any pair of Type 5 (Special) wedges. In fact, the following statement holds, which implies that no polygon with a Type 5 (Special) pair of wedges is totally-cover-decomposable [P10].

Lemma 5.1. [P10] *For any pair of special wedges, V and W , and for any pair of positive integers, k and l , there is a point set P of size $\binom{k+l}{k} - 1$ such that for every coloring of P with red and blue, either there is a translate of V containing k red points and no blue points, or there is a translate of W containing l blue points and no red points.*

Proof. We proceed by induction on $k+l$. Denote by $P(k,l)$ a set of points that satisfy the conditions

of the lemma for k and l . If k or l is equal to 1, then the statement is trivially true. In the induction step (see the left side of Figure 15), place a point p in the plane and a suitable small scaled down copy of $P(k-1, l)$ with the property that any translate of V with its apex in the neighborhood of $P(k-1, l)$ contains p , but none of the translates of W with its apex in the neighborhood of $P(k-1, l)$ does. Analogously, place a scaled down copy of $P(k, l-1)$ in such a way that any translate of W with its apex in the neighborhood of $P(k, l-1)$ contains p , but none of the translates of V with its apex in the neighborhood of $P(k, l-1)$ does.

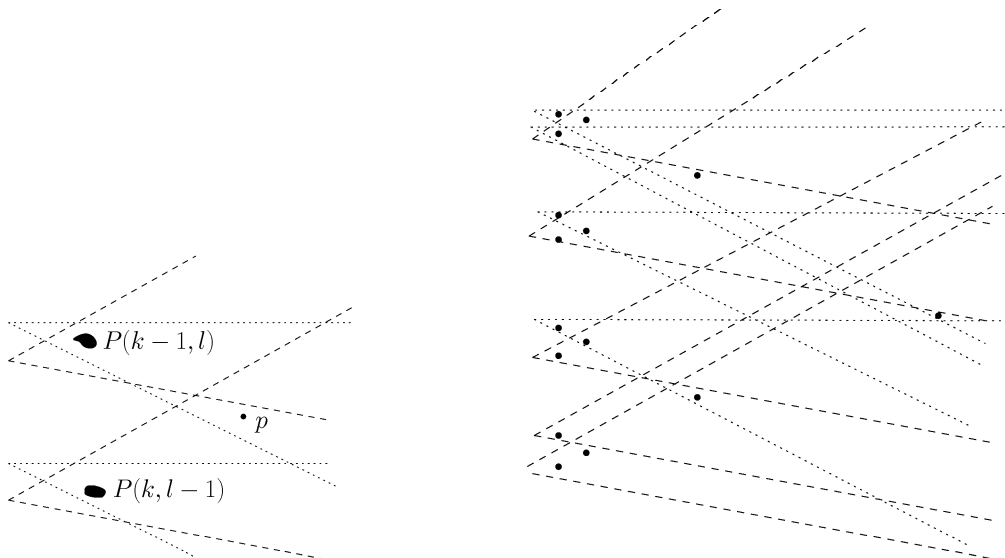


Figure 15: Sketch of one step of the induction and iteration of some steps.

If p is colored red, then either for the first part of the construction, similar to $P(k-1, l)$, there is a translate of V that covers point p together with $k-1$ other red points and no blue ones, or for the part similar to $P(k, l-1)$, there is a translate of W that covers l blue points, no red ones, and it does not contain p . In both cases, we are done. A similar argument works in the case when p is blue. \square

Remark 5.2. *Instead of considering all translates of V and W , in order to find a wedge that meets the requirements of Lemma 5.1, it is sufficient to restrict our attention to a finite set of translates whose apices lie on the same line.*

This construction, combined with Lemma 4.3 and 4.4, gives the following characterization of polygons.

Theorem 5.3. [PT10, P10] *An open polygon P is totally-cover-decomposable if and only if none of the P -wedges form a pair of Type 5 (Special).*

Unfortunately, we still do not have a nice characterization for *plane-cover-decomposability*. The reason is that the above construction cannot always be extended to coverings of the whole plane. As pointed out in Remark 5.2, it is sufficient to consider a finite set of wedges whose apices lie on the same line. However, after dualization the centers of the translates will lie on *two* lines. An example of a polygon which is not totally-cover-decomposable but might be cover-decomposable is depicted in Figure 16. Some special cases when such an extension is always possible, were studied in [P10].




Figure 16: Unknown hexagon: its only special pair of wedges are at A and E.

In higher dimensions, the situation is completely different. According to Theorem 1.12 [P10], for $d \geq 3$, no d -dimensional convex polytope is cover-decomposable.

The proof is based on the observation that for any polytope P , either there is a plane that intersects P in a concave polygon, which always has a special pair of wedges, or there are two parallel planes that intersect P in two polygons such that there is a special pair among their wedges. In both cases, we can take a plane in space and a family of translates of P that realize the above construction in this plane so that the intersection of the plane and the translates of P play the role of the wedges. Then we take the dual of this arrangement. To prove that this construction is extendable to an indecomposable covering of the entire space, observe that the centers of all the translates used in the construction lie in a plane, therefore, we can follow the same argument as for quadrilaterals in the plane.

5.3 Technical difficulties: closed polygons, finite covering

Notice that in all of our positive results (Theorems 1.3-1.9) we considered *open* polygons. This is due to the fact that at the very beginning of Section 2, based on a compactness argument, we restricted our attention to *locally finite* coverings. This does not work for closed polygons. The truth is that at the moment for not locally finite coverings with closed polygons, we cannot prove any positive result. (Our negative results, of course, remain valid for closed polygons as well.) In [P10], we made an attempt to overcome this difficulty. To state the (rather weak) results obtained there,

we need a definition.

Definition 5.4. *A planar set P is said to be finite-cover-decomposable (countable-cover-decomposable) if there exists a constant $m \geq 2$ such that every m -fold covering of any point set with finitely (countably) many translates of P can be decomposed into two coverings.*

By definition, we have: P is totally-cover-decomposable $\Rightarrow P$ is countable-cover-decomposable $\Rightarrow P$ is finite-cover-decomposable. But which of these implications can be reversed? In [P10], it was proved that the first one can be for “nice” sets. The definition of nice includes all closed convex sets and polygons, but is much more general. The proof is based on the hereditary Lindelöf property of the plane.

Unfortunately, we have been unable to prove any such connection between finite-cover-decomposability and countable-cover-decomposability. Hence, the status of closed polygons is still undetermined. We believe, however, that using further geometric observations this problem can be settled.

6 Open questions

The main unsolved problem in the field remains to verify (or refute) Conjecture 1.2 or, more generally Conjecture 1.14.

Problem 6.1. *Is every plane convex set cover-decomposable?*

Concerning coverings with homothetic copies of a set P , the first interesting special cases are when P is a disk or a square.

Problem 6.2. *Does there exist a positive integer m such that every m -fold covering of the plane with open disks of arbitrary radii splits into two coverings?*

Problem 6.3. *Does there exist a positive integer m such that every m -fold covering of the plane with open squares of arbitrary side lengths, whose sides are parallel to the coordinate axes, splits into two coverings?*

As we have seen in the Introduction, the answer to the corresponding question for triangles is affirmative [KP11].

In Subsection 2.3, we defined a notion somewhat stronger than cover-decomposability (see Definition 2.3).

Problem 6.4. *Does there exist a bounded (convex) set P which is cover-decomposable, but not totally-cover-decomposable?*

According to Theorem 1.5, every open convex polygon is cover-decomposable, that is, for every open convex polygon P , there is a positive constant $m(P)$ such that every $m(P)$ -fold covering of the plane with translates of P splits into two coverings. The best known value of $m(P)$ depends on the shape of P .

Problem 6.5. *Is it true that, for any integer $j \geq 3$, there is a positive constant m_j such that every m_j -fold covering of the plane with translates of any convex j -gon P splits into two coverings?*

For open triangles the answer is yes with $m_3 \leq 12$. On the other hand, the same statement is not known for closed triangles, as we do not even know if closed triangles are cover decomposable.

It is possible that for any cover-decomposable set P , there exists a (smallest) positive integer $m = m_3(P)$ with the property that every m -fold covering of the plane with translates of P splits into *three* coverings. More generally, as in the Introduction, let $m_k(P)$ denote the smallest positive integer m such that every m -fold covering of the plane with translates of P splits into k coverings. If such an integer does not exist, let $m_k(P) = \infty$.

Problem 6.6. *Is it true that if $m_2(P) < \infty$, then we also have $m_k(P) < \infty$, for every $k \geq 3$?*

This may be true even in a very general combinatorial setting. Given a finite system of sets \mathcal{F} , a *multiset* of its members (with possible repetition!) is said to form a m -fold covering if every element of the underlying set is contained in at least m members of \mathcal{F} . For any positive integer k , let $m_k(\mathcal{F})$ denote the smallest number $m \geq 1$ such that every m -fold covering with members of \mathcal{F} splits into k coverings. It is easy to see that this number is always finite: for example, we have $m_k(\mathcal{F}) \leq (k-1)|\mathcal{F}| + 1$.

Problem 6.7. *Does there exist a function f such that, for every finite set system \mathcal{F} , we have $m_3(\mathcal{F}) < f(m_2(\mathcal{F}))$?*

It is possible that the answer is yes even with the function $f(x) = O(x)$. As a matter of fact, the relation $m_k(\mathcal{F}) < Ckm_2(\mathcal{F})$ may also hold with an absolute constant $C > 0$.

In spite of substantial progress in this field, our knowledge on decomposability properties of multiple coverings is rather rudimentary. To our surprise, G. Tardos (personal communication) constructed a set system \mathcal{F} , which “almost” refutes Problem 6.7. This set system cannot be decomposed into 3 coverings, although every subsystem of it (with no repetition!) which forms a 2-fold covering splits into 2 coverings.

Finally, we mention another problem for finite set systems that has a strong connection to cover-decomposability.

For a subset $A \subset [n]$, let us denote by a_i the i -th smallest element of A . Given two k -element sets, $A, B \subset [n]$, we write $A \preceq B$ if $a_i \leq b_i$ for every i . A k -uniform hypergraph $\mathcal{H} \subset \binom{[n]}{k}$ is called a *shift-chain* if for any two hyperedges, $A, B \in \mathcal{H}$, we have $A \preceq B$ or $B \preceq A$. (So a shift-chain has at most $k(n-k) + 1$ hyperedges.)

Problem 6.8. *Is it true that if k is sufficiently large, then every k -uniform shift-chain has Property B? In other words, is it true that for every shift-chain $\mathcal{H} \subset \binom{[n]}{k}$, one can color $[n]$ with two colors such that no hyperedge is monochromatic?*

An affirmative answer would be a huge step towards Pach’s conjecture that all planar convex sets are cover-decomposable. To see this, recall the following definition from Section 4.1. For a finite set

of point S in the plane and for a plane convex set P , define $P(k; y)$ as the translate of P which

- (1) contains exactly k points of S ;
- (2) the y -coordinate of its apex is y ; and
- (3) the x -coordinate of its apex is maximal,

if such a translate exists.

If we associate $i \in [n]$ with the element of S with the i -th smallest y -coordinate, then an easy geometric argument shows that $\mathcal{H} = \{P(k; y) \cap S \mid y \in \mathbb{R}\}$ is a shift-chain.

For $k = 2$, there is a trivial counterexample to the above problem: (12),(13),(23). For $k = 3$, a magic counterexample was found by a computer program written by Radoslav Fulek:

(123), (124), (125), (135), (145), (245), (345), (346), (347), (357), (367), (467), (567), (568), (569), (579), (589), (689), (789).

If we allow the hypergraph to be the union of two shift-chains (with the same order), then the construction in Section 5.2 provides a counterexample for any k . Therefore, all arguments using that the average degree is small (like attempts based on Lovász' Local Lemma) would probably fail.

References

- [A10] B. Ács: *Síkfedések szétbonthatósága*, Master Thesis, Eötvös University Budapest, 2010 (in Hungarian).
- [Al10] G. Aloupis, J. Cardinal, S. Collette, S. Langerman, D. Orden, and P. Ramos: Decomposition of multiple coverings into more parts, *Discrete and Computational Geometry* **44** (2010), 706–723. Also in: *Proc. 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 09)*, ACM, New York, 2009, 302–310.
- [B07] A. L. Buchsbaum, A. Efrat, S. Jain, S. Venkatasubramanian, and K. Yi: Restricted strip covering and the sensor cover problem, in: *Proc. 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 07)*, ACM, New York, 2007, 1056–1063.
- [BMP05] P. Brass, J. Pach, and W. Moser: *Research Problems in Discrete Geometry*, Springer, Berlin, 2005.
- [EMS11] M. Elekes, T. Mátrai, and L. Soukup: On splitting infinite-fold covers, *Fundamenta Mathematicae* **212** (2011), 95–127.
- [E63] P. Erdős: On a combinatorial problem, *Nordisk Matematisk Tidskrift* **11** (1963), 5–10.
- [EL75] P. Erdős and L. Lovász: Problems and results on 3-chromatic hypergraphs and some related questions, in: *Infinite and Finite Sets* (to Paul Erdős on his 60th birthday), II. North-Holland, Amsterdam, 1975, 609–627.

- [GV11] M. Gibson and K. Varadarajan: Optimally decomposing coverings with translates of a convex polygon, *Discrete & Computational Geometry* **46** (2011), 313–333. Also in: *Proc. 50th Annual IEEE Symposium on Foundations of Computer Science, (FOCS 09)*, IEEE Computer Soc., Los Alamitos, CA, 2009, 159–168.
- [KP11] B. Keszegh and D. Pálvölgyi: Octants are cover decomposable, *Discrete & Computational Geometry*, DOI 10.1007/s00454-011-9377-1, to appear. Also in: *Proc. 7th Hungarian-Japanese Symposium on Discrete Mathematics and Its Applications, Kyoto, 2011*, 217–226.
- [MP86] P. Mani-Levitska, J. Pach: Decomposition problems for multiple coverings with unit balls, manuscript, 1986.
- [P80] J. Pach: Decomposition of multiple packing and covering, in: *Diskrete Geometrie, 2. Kolloq.* Math. Inst. Univ. Salzburg, 1980, 169–178.
- [P86] J. Pach: Covering the plane with convex polygons, *Discrete & Computational Geometry* **1** (1986), 73–81.
- [PTT05] J. Pach, G. Tardos, and G. Tóth: Indecomposable coverings, *Canadian Mathematical Bulletin* **52** (2009), 451–463. Also in: *The China–Japan Joint Conference on Discrete Geometry, Combinatorics and Graph Theory (CJCDGCGT 2005), Lecture Notes in Computer Science* **4381**, Springer, Berlin, 2007, 135–148.
- [PT07] J. Pach and G. Tóth: Decomposition of multiple coverings into many parts, *Computational Geometry: Theory and Applications* **42** (2009), 127–133. Also in: *Proc. 23rd ACM Symposium on Computational Geometry, (SoCG07)* 2007, 133–137.
- [P10] D. Pálvölgyi: Indecomposable coverings with concave polygons, *Discrete & Computational Geometry* **44**, (2010), 577–588.
- [PT10] D. Pálvölgyi and G. Tóth: Convex polygons are cover-decomposable, *Discrete & Computational Geometry* **43** (2010), 483–496.
- [TT07] G. Tardos and G. Tóth: Multiple coverings of the plane with triangles, *Discrete & Computational Geometry* **38** (2007), 443–450.