

Coloring Points with Respect to Squares

Eyal Ackerman¹ · Balázs Keszegh²  · Mate Vizer²

Received: 1 September 2016 / Revised: 20 May 2017 / Accepted: 30 May 2017 /

Published online: 19 June 2017

© Springer Science+Business Media New York 2017

Abstract We consider the problem of 2-coloring geometric hypergraphs. Specifically, we show that there is a constant m such that any finite set of points in the plane $\mathcal{S} \subset \mathbb{R}^2$ can be 2-colored such that every axis-parallel square that contains at least m points from \mathcal{S} contains points of both colors. Our proof is constructive, that is, it provides a polynomial-time algorithm for obtaining such a 2-coloring. By affine transformations this result immediately applies also when considering 2-coloring points with respect to homothets of a fixed parallelogram.

Keywords Geometric hypergraph coloring · Cover-decomposability · Self-coverability

Mathematics Subject Classification 52C15 (05C15)

Editor in Charge: János Pach

A preliminary version of this paper was presented at the 32nd International Symposium on Computational Geometry (SoCG 2016).

Eyal Ackerman
ackerman@sci.haifa.ac.il

Balázs Keszegh
keszegh@renyi.hu

Mate Vizer
vizermate@gmail.com

¹ Department of Mathematics, Physics, and Computer Science, University of Haifa at Oranim, Tivon 36006, Israel

² Alfréd Rényi Institute of Mathematics, Hungarian Academy of Sciences, Budapest 1053, Hungary

1 Introduction

In this paper we consider the problem of coloring a given set of points in the plane such that every region from a given set of regions contains a point from each color class. To state our results, known results and open problems, we need the following definitions and notations.

A *hypergraph* is a pair $(\mathcal{V}, \mathcal{E})$ where \mathcal{V} is a set and \mathcal{E} is a set of subsets of \mathcal{V} . The elements of \mathcal{V} and \mathcal{E} are called *vertices* and *hyperedges*, respectively. For a hypergraph $H := (\mathcal{V}, \mathcal{E})$, let $H|_m := (\mathcal{V}, \{e \in \mathcal{E} : |e| \geq m\})$. A *proper coloring* of a hypergraph is a coloring of its vertex set such that in every hyperedge not all vertices are assigned the same color. Proper colorability of a hypergraph with two colors is also called *Property B* in the literature. A *polychromatic k -coloring* of a hypergraph is a coloring of its vertex set with k colors such that every hyperedge contains at least one vertex from each of the k colors.

Given a family of regions \mathcal{F} in \mathbb{R}^d (e.g., all disks in the plane), there is a natural way to define two types of finite hypergraphs that are dual to each other. First, for a finite set of points \mathcal{S} , let $H^{\mathcal{F}}(\mathcal{S})$ denote the *primal hypergraph* on the vertex set \mathcal{S} whose hyperedges are all subsets of \mathcal{S} that can be obtained by intersecting \mathcal{S} with a member of \mathcal{F} . We say that a finite subfamily $\mathcal{F}_0 \subseteq \mathcal{F}$ *realizes* $H^{\mathcal{F}}(\mathcal{S})$ if for every hyperedge $S' \subseteq \mathcal{S}$ of $H^{\mathcal{F}}(\mathcal{S})$ there is $F' \in \mathcal{F}_0$ such that $F' \cap \mathcal{S} = S'$. The *dual hypergraph* $H^*(\mathcal{F}_0)$ is defined with respect to a finite multi-subfamily¹ $\mathcal{F}_0 \subseteq \mathcal{F}$. Its vertex set is \mathcal{F}_0 and for each point $p \in \mathbb{R}^d$ it has one hyperedge that consists of exactly those regions in \mathcal{F}_0 that contain p .

The general problems we are interested in are the following.

Problem 1.1 *For a given family of regions \mathcal{F} ,*

- (i) *Is there a constant m such that for any finite set of points \mathcal{S} the hypergraph $H^{\mathcal{F}}(\mathcal{S})|_m$ admits a proper 2-coloring?*
- (ii) *Is there a constant m^* such that for any finite subset $\mathcal{F}_0 \subseteq \mathcal{F}$ the hypergraph $H^*(\mathcal{F}_0)|_{m^*}$ admits a proper 2-coloring?*
- (iii) *Given a constant k , is there a constant m_k such that for any finite set of points \mathcal{S} the hypergraph $H^{\mathcal{F}}(\mathcal{S})|_{m_k}$ admits a polychromatic k -coloring? If so, is $m_k = O(k)$?*
- (iv) *Given a constant k , is there a constant m_k^* such that for any finite subset $\mathcal{F}_0 \subseteq \mathcal{F}$ the hypergraph $H^*(\mathcal{F}_0)|_{m_k^*}$ admits a polychromatic k -coloring? If so, is $m_k^* = O(k)$?*

Examples of families \mathcal{F} for which such coloring problems are studied are translates of convex sets [3, 10, 23, 30, 33, 37], homothets of triangles [6, 7, 14–17], axis-parallel rectangles [8, 9, 26, 28] and half-planes [13, 36]. If \mathcal{F} is the family of disks in the plane, then these hypergraphs generalize Delaunay graphs.

The main motivation for studying proper and polychromatic colorings of such geometric hypergraphs comes from cover-decomposability problems [25] and conflict-free coloring problems [35]. We concentrate on the first connection, as the problems we regard are in direct connection with cover-decomposability problems. We give a

¹ In a multisubfamily we allow taking multiple copies of the same set.

short introduction to this topic here, however we recommend the interested reader to consult the survey paper [25] or the webpage [32] that contains a summary of results about decomposition of multiple coverings and polychromatic colorings.

Multiple coverings and packings were first studied by Davenport and L. Fejes Tóth almost 50 years ago. Since then a wide variety of questions related to coverings and packings has been investigated. In 1986 Pach [23] published the first paper about decomposability problems of multiple coverings. It turned out that this area is rich of deep and exciting questions, and it has important practical applications as well (e.g., in the area of surveillance systems [10, 25]). Following Pach's papers, most of the efforts were concentrated on studying coverings by translates of some given shape. Recently, several researchers started studying cover-decomposability of homothets of a given shape.

A family of planar sets is called an r -fold covering of a region R , if every point of R is contained in at least r members of the family. A 1-fold covering is simply called a covering. A family \mathcal{F} of planar sets is called *cover-decomposable*, if there is an integer ℓ with the property that for any multi-subfamily of \mathcal{F} that forms an ℓ -fold covering of the whole plane can be decomposed into two coverings.² A family \mathcal{F} of planar sets is called *totally-cover-decomposable*, if there is an integer ℓ^T with the property that for any region R , any multi-subfamily of \mathcal{F} that forms an ℓ^T -fold covering of R can be decomposed into two coverings.

One can also ask for a decomposition into more than two coverings. That is, whether there exists an integer ℓ_k (resp., ℓ_k^T) such that any multi-subfamily of \mathcal{F} that forms an ℓ_k -fold covering of the whole plane (resp., an ℓ_k^T -fold covering of R , for any region R), can be decomposed into k coverings of the plane (resp., of R).

If we consider only coverings with finite multi-subfamilies, then we call it the *finite cover-decomposition* problem. It is easy to see that the finite cover-decomposition problem is equivalent to Problems 1.1 (ii) and (iv).

One of the first observations of Pach was that if \mathcal{F} is the family of translates of an open convex set, then finite cover-decomposability implies cover-decomposability. Thus for the family of the translates of an open convex set $\ell \leq m^*$ and $\ell_k \leq m_k^*$ in the notation above. Pach also observed that if \mathcal{F} is the family of translates of some set, then Problems 1.1 (i) and (ii) are equivalent and also Problems 1.1 (iii) and (iv) are equivalent, i.e., $m^* = m$ and $m_k^* = m_k$. Thus, for the family of translates of an open convex set $\ell \leq m^* = m$ and $\ell_k \leq m_k^* = m_k$ and so it is enough to consider the primal hypergraph coloring problem.

Pach conjectured that the family of all translates of any open convex planar set is cover-decomposable [22]. During the years researchers acquired a good understanding of convex planar shapes whose translates are cover-decomposable. On the positive side, Pach's conjecture was verified for every open convex polygon: Pach himself proved it for every open centrally symmetric convex polygon [23], then Tardos and Tóth [37] proved the conjecture for every open triangle, and finally Pálvölgyi and Tóth [33] proved it for every open convex polygon. They also gave a complete characterization

² In fact in early papers that considered only coverings by translates, all mentioned variants of cover-decomposability were defined for sets instead of families, where a set is cover-decomposable if the family of its translates is cover-decomposable in the way we define it in this paper.

of open non-convex polygons whose translates are finite cover-decomposable. For open convex polygons it is even known that $m_k = m_k^* = O(k)$ [3, 10, 30]. However, Pach's conjecture was refuted in [24]. Specifically, it does not hold for disks and for convex shapes with a smooth boundary.

Considering the three dimensional space, it follows from cover-indecomposability of certain non-convex polygons [31] that every bounded polytope is not cover-decomposable. Thus, it is not easy to come up with a cover-decomposable set in the space. An important exception is the *octant*,³ whose translates were proved to be cover-decomposable [14]. The currently best bounds for octants are $5 \leq m \leq 9$ [17] and $m_k = m^* = O(k^{5.09})$ [7, 16, 17]. It is a challenging open problem whether $m_k = m_k^* = O(k)$ in this case.

For a long time no positive results were known about cover-decomposability and geometric hypergraph coloring problems concerning homothets of a given shape. For disks, the answer is negative for all parts of Problem 1.1 [24, 27]. As a first positive result, the cover-decomposability of octants along with a simple reduction implied that both the primal and dual hypergraphs with respect to homothets of a triangle are properly 2-colorable:

Theorem 1.2 ([14, 17]) *For the family \mathcal{F} of all homothets of a given triangle both Problems 1.1 (i) and (iii) have a positive answer with $m = m^* \leq 9$.*

This result was later used to obtain polychromatic colorings of the primal and dual hypergraphs defined by the family of homothets of a fixed triangle. For the dual hypergraph, the best bound comes from the corresponding result about octants and so it is $m_k^* = O(k^{5.09})$. For the primal hypergraph there is a better bound $m_k = O(k^{4.09})$ [15, 17]. An important tool for obtaining these results is the notion of *self-coverability* (see Sect. 2.2), which is also essential for proving our results. The questions whether $m_k = O(k)$ and $m_k^* = O(k)$ for the homothets of a given triangle are still open. The web-page [32] contains an up-to-date collection of results considering all of these problems and related ones.

For polygons other than triangles, somewhat surprisingly, Kovács [20] recently provided a negative answer for Problems 1.1 (ii) and (iv). Namely, he showed that the homothets of any given convex polygon with at least four sides are not cover-decomposable. In other words, there is no constant m^* for which the dual hypergraph consisting of hyperedges of size at least m^* is 2-colorable. Our main contribution is showing that this is not the case when considering 2-coloring the primal hypergraph. Indeed, Problem 1.1 (i) has a positive answer for homothets of any given *parallelogram*.

Theorem 1.3 *There is an absolute constant $m_q \leq 215$ such that the following holds. Given an (open or closed) parallelogram Q and a finite set of points in the plane $S \subset \mathbb{R}^2$, the points of S can be 2-colored in polynomial time, such that any homothet of Q that contains at least m_q points contains points of both colors.*

This is the first example that exhibits such different behavior for coloring the primal and dual hypergraphs with respect to the family of some geometric regions. Fur-

³ An octant is the set of points $\{(x, y, z) : x < a, y < b, z < c\}$ in the space for some a, b and c .

thermore, combined with results about self-coverability, the proof of Theorem 1.3 immediately implies the following generalization to polychromatic k -colorings, thus partially answering also Problem 1.1 (iii) (it remains open whether linearly many points per hyperedge/parallelogram suffice).

Corollary 1.4 *Let Q be a given (open or closed) parallelogram and let S be a set of points in the plane. Then for every integer $k \geq 1$ it is possible to color S with k colors, such that any homothet of Q that contains at least $m_k = \Omega(k^{8.75})$ points from S contains points of all k colors.*

Our proof of Theorem 1.3 also works for homothets of a triangle, i.e., we give a new proof for the primal case of Theorem 1.2 (with a larger constant though):

Theorem 1.5 ([14]) *There is an absolute constant m_t such that the following holds. Given an (open or closed) triangle T and a finite set of points in the plane $S \subset \mathbb{R}^2$, the points of S can be 2-colored in polynomial time, such that any homothet of T that contains at least m_t points contains points of both colors.*

This paper is organized as follows. In Sect. 2 we introduce definitions, notations, tools and some useful lemmas. In Sect. 3 we describe a general 2-coloring algorithm and then apply it for parallelograms and for triangles. Concluding remarks and open problems appear in Sect. 4. A preliminary version of this paper was presented at the 32nd International Symposium on Computational Geometry (SoCG 2016). In the current version some of the proofs are simplified, the constant in Theorem 1.3 is better, and the limitations of our techniques are discussed in more details.

2 Preliminaries

Unless stated otherwise, we restrict ourselves to the two-dimensional Euclidean space \mathbb{R}^2 . For a point $p \in \mathbb{R}^2$ let $(p)_x$ and $(p)_y$ denote the x - and y -coordinate of p , respectively. We denote by ∂S the boundary of a subset $S \subseteq \mathbb{R}^2$ and by $\text{Cl}(S)$ the closure of S . A *homothet* of S is a translated and scaled copy of S . That is, a set $S' = \alpha S + p$ for some number $\alpha > 0$ and a point $p \in \mathbb{R}^2$. We will use the following folklore lemma:

Lemma 2.1 *Let C be a convex and compact set and let C_1 and C_2 be homothets of C . Then if ∂C_1 and ∂C_2 intersect finitely many times, then they intersect in at most two points.*

For a proof of this lemma, see, e.g., [21, Cor. 2.1.2.2].

2.1 Generalized Delaunay Triangulations

For proving Theorems 1.3 and 1.5 we will use the notion of generalized Delaunay triangulations, which are the dual of generalized Voronoi diagrams. In the generalized Delaunay triangulation of a point set S with respect to some compact convex set C , two points of S are connected by a straight-line segment if there is a homothet

of C that contains these two points and does not contain any other point of S . The generalized Delaunay triangulation of S with respect to C is denoted by $\mathcal{DT}(C, S)$. Without causing confusion we also regard such a Delaunay triangulation as an abstract graph with S as its vertex set and the above defined segments correspond to its edges. We say that S is *in general position* with respect to (homothets of) C , if there is no homothet of C whose boundary contains four points from S . If S is in general position with respect to a convex polygon P and no two points of S define a line that is parallel to a line through two vertices of P , then we say that S is in *very general position* with respect to P . The following properties of generalized Delaunay triangulations will be useful.

Lemma 2.2 ([4, 19, 34]) *Let C be a compact convex set and let S be a set of points in general position with respect to C . Then $\mathcal{DT}(C, S)$ is a well-defined connected plane graph whose inner faces are triangles.*

It would be convenient to consider generalized Delaunay triangulations in which the boundary of the outer face is a convex polygon. In such a case we say that $\mathcal{DT}(C, S)$ is *nice*.

Lemma 2.3 *Let P be a closed convex polygon and let S be a set of points in the plane that is in very general position with respect to P . Suppose that P' is a homothet of P and $Z \subseteq S \cap \partial P'$. Then there is a homothet of P , denote it by P'' , such that $P'' \cap S = (P' \cap S) \setminus Z$.*

Proof Since S is in general position, $|\partial P' \cap S| \leq 3$.

If $|Z| = 3$, then there are no other points on $\partial P'$, thus shrinking P' from an inner point gives us the required P'' .

If $|Z| = 2$, then if there is no other point on $\partial P'$, then we can again shrink from an inner point slightly to get the required P'' . Otherwise, there is exactly one point q on $\partial P'$ besides the two points of Z . Now slightly shrink P' from q . As the points are in very general position, the resulting homothet will contain exactly the points $(P' \cap S) \setminus Z$.

Finally, suppose that $|Z| = 1$ and let $Z = \{z\}$. We consider three cases.

Case 1: $\partial P' \cap S = \{z\}$. In this case we slightly shrink P' with respect to some point in its interior and obtain the desired homothet P'' .

Case 2: $|\partial P' \cap S| = 2$. Let $\partial P' \cap S = \{x, z\}$. Since S is in very general position, x and z are on different sides of $\partial P'$. Therefore if we slightly shrink P' with respect to x , the resulting homothet of P contains x , does not contain z and contains all other points in $P' \cap S$.

Case 3: $|\partial P' \cap S| = 3$. Let $\partial P' \cap S = \{x, y, z\}$. In this case first we enlarge P' from z to get a homothet P^+ . Since the points in S are in very general position with respect to P , this can be done so that $\partial P^+ \cap S = \{z\}$ and $P' \cap S = P^+ \cap S$. Then by slightly shrinking P^+ with respect to an interior point we get the desired homothet P'' . \square

For a homothet C' of a compact convex set C we denote by $\mathcal{DT}(C, S)[C']$ the subgraph of $\mathcal{DT}(C, S)$ that is induced by the points of $S \cap C'$. Note that it is not the same as $\mathcal{DT}(C, S \cap C')$, however the following is true.

Lemma 2.4 *Let P be a closed convex polygon, let S be a set of points in very general position with respect to P , and let P' be a homothet of P . Then $\mathcal{DT}(P, S)[P']$ is a connected graph that is contained in P' .*

Proof Let $\mathcal{DT} := \mathcal{DT}(P, S)$. Obviously all the points in $\mathcal{DT}[P']$ are in P' by definition. An edge in $\mathcal{DT}[P']$ must also be in P' , since P' is convex. We prove that $\mathcal{DT}[P']$ is connected by induction on $|\mathcal{S} \cap P'|$. This is true by definition if there are at most two points from \mathcal{S} in P' . Suppose that the claim holds whenever a homothet of P contains $k - 1$ points from \mathcal{S} and let P' be a homothet of P that contains exactly k points from \mathcal{S} with $k \geq 3$. We may assume without loss of generality that P' contains two points, x and z , on its boundary, for otherwise we can continuously shrink P' until such points exist (first from an interior point, then from the point that first appears on the boundary). Now apply Lemma 2.3 twice: once with $Z = \{x\}$ and once with $Z = \{z\}$. Denote the homothets of P that we get by P_x and P_z , respectively. By the induction hypothesis $\mathcal{DT}[P_x]$ and $\mathcal{DT}[P_z]$ are both connected, their intersection contains at least one point (as $k \geq 3$), and their union is contained in $\mathcal{DT}[P']$. Thus, $\mathcal{DT}[P']$ is also connected, as required. \square

Corollary 2.5 *Let P be a closed convex polygon and let S be a set of points in very general position with respect to P . Suppose that P' is a homothet of P and e is an edge of $\mathcal{DT}(P, S)$ that crosses $\partial P'$ twice and thus splits P' into two parts. Then one of these parts does not contain a point from S .*

A rotation of a vertex v in a plane graph G is the clockwise order of its neighbors. For three straight-line edges vx, vy, vz we say that vy is between vx and vz if x, y, z appear in this order in the rotation of v and $\angle xvz < \pi$ ($\angle xvz$ is the angle by which one has to rotate the vector \vec{vx} around v clockwise until its direction coincides with that of \vec{vz}) or if z, y, x appear in this order in the rotation of v and $\angle zvx < \pi$. The following will be useful later on.

Proposition 2.6 *Let C be a compact convex set and let S be a set of points in very general position with respect to C and such that $\mathcal{DT} := \mathcal{DT}(C, S)$ is nice. Let C' be a homothet of C and let v be a vertex in $\mathcal{DT}[C']$. Suppose that x and z are two vertices that are neighbors of v in $\mathcal{DT}[C']$ and $\angle xvz < \pi$ and $xz \notin \mathcal{DT}$. Then there exists an edge $vy \in \mathcal{DT}$ between vx and vz . Moreover, if z immediately follows x in the rotation of v in $\mathcal{DT}[C']$ then $y \notin C'$.*

Proof If x and z are not consecutive in the rotation of v in \mathcal{DT} then by definition there exists a vertex y between them in the rotation of v .

Thus we are done unless x and z are consecutive in the rotation of v in \mathcal{DT} .

Suppose that such an y does not exist, that is, x and z are also consecutive in the rotation of v in \mathcal{DT} . Then the face that is incident to vx and vz and is to the right of \vec{vx} and to the left of \vec{vz} cannot be the outer face since $\angle xvz < \pi$ and \mathcal{DT} is nice. However, since this face is an inner face, then by Lemma 2.2 it must be a triangle and so $xz \in \mathcal{DT}$.

Thus we can conclude that such an y exists and from the definition of the rotation order it follows that $y \notin C'$. \square

Lemma 2.7 For every closed convex polygon P there is a constant $\Delta := \Delta(P)$ such that the following holds. Let \mathcal{S} be a set of points in very general position with respect to P and such that $\mathcal{DT} := \mathcal{DT}(P, \mathcal{S})$ is nice. If P' is a homothet of P such that $\mathcal{DT}[P']$ is a tree, then for every $v \in \mathcal{S} \cap P'$ we have $\deg_{\mathcal{DT}[P']}(v) \leq \Delta$.

Proof Let n be the number of vertices of P and let v_0, v_1, \dots, v_{n-1} be the vertices of P' listed in their clockwise order around P' . Let $v \in \mathcal{S} \cap P'$ be a point and let $N_i := \{u \in \mathcal{S} \cap P' : u \in \Delta v v_i v_{i+1}, vu \in \mathcal{DT}[P']\}$ be the neighbors of v in $\mathcal{DT}[P']$ that are also in the triangle $vv_i v_{i+1}$, for every $i = 0, \dots, n - 1$ (addition is modulo n). Let $\alpha := \alpha(P)$ be the smallest angle formed by three vertices of P (hence, also of P').

Observation 2.8 For every point p' in the interior of P' and every $0 \leq i < j \leq n - 1$ we have $\angle v_i p' v_j \geq \alpha$.

Proposition 2.9 For every $i = 0, \dots, n - 1$ and every $u, u' \in N_i$ we have $\angle uvu' \geq \alpha$.

Proof It is enough to consider the case when u and u' follow each other immediately in the rotation of v . First note that uu' is not an edge in $\mathcal{DT}[P']$, since otherwise there would be a triangle in $\mathcal{DT}[P']$. It follows from Proposition 2.6 that there is a point $z \notin P'$ such that z is a neighbor of v in \mathcal{DT} and is between u and u' in the rotation of v . Thus, there is a homothet of P , denote it by P_z , such that $P_z \cap \mathcal{S} = \{v, z\}$. Let $v'_i v'_{i+1}$ be the side of P_z that corresponds and is parallel to the side $v_i v_{i+1}$ of P' (see Fig. 1). Note that $v'_i v'_{i+1}$ is outside of P' , since $z \notin P'$. Furthermore, since $u, u' \notin P_z$, the side $v'_i v'_{i+1}$ lies inside the wedge whose apex is v and whose boundary consists of the two rays that emanate from v and go through u and u' , respectively. Therefore, $\angle v'_i v v'_{i+1} < \angle uvu'$. It follows from Observation 2.8 that $\angle v'_i v v'_{i+1} \geq \alpha$, thus we have $\angle uvu' \geq \alpha$. □




Fig. 1 An illustration for the proof of Proposition 2.9

To complete the proof of Lemma 2.7 consider the neighbors of v in $\mathcal{DT}[P']$ in their clockwise order around v , and for every set N_i remove the extreme neighbor in this order. It follows from Proposition 2.9 that the angle between any two remaining neighbors of u is at least α . Therefore, $\deg_{\mathcal{DT}[P']}(v) \leq n + \frac{2\pi}{\alpha}$. \square

It follows that if P is an equilateral triangle, then Lemma 2.7 applies with $\Delta(P) \leq 3 + \frac{2\pi}{\pi/3} = 9$. By affine transformations we have:

Corollary 2.10 *Suppose that T is a triangle and S is a set of points in very general position with respect to T and such that $\mathcal{DT} := \mathcal{DT}(T, S)$ is nice. If T' is a homothet of T such that $\mathcal{DT}[T']$ is a tree, then for every point $v \in T' \cap S$ we have $\deg_{\mathcal{DT}[T']}(v) \leq 9$.*

Corollary 2.11 *Let P be a convex polygon and let $\Delta := \Delta(P)$ be the constant from Lemma 2.7. Suppose that S is a set of points in very general position with respect to P and such that $\mathcal{DT} := \mathcal{DT}(P, S)$ is nice. If P' is a homothet of P such that $\mathcal{DT}[P']$ is a tree, then $\mathcal{DT}[P']$ contains a simple path of length at least $2\lfloor \log_{\Delta} |S \cap P'| \rfloor$.*

2.2 Self-coverability of Convex Polygons and Polychromatic k -Coloring

Keszegh and Pálvölgyi introduced in [15] the notion of *self-coverability* and its connection to polychromatic k -coloring. In this section we list the definition and results from their work that we use.

Definition 2.12 ([15]) A collection of closed sets \mathcal{F} in a topological space is *self-coverable* if there exists a *self-coverability function* f such that for any set $F \in \mathcal{F}$ and for any finite point set $S \subset F$, with $|S| = l$ there exists a subcollection $\mathcal{F}' \subset \mathcal{F}$, $|\mathcal{F}'| \leq f(l)$ such that $\bigcup_{F' \in \mathcal{F}'} F' = F$ and no point of S is in the interior of some $F' \in \mathcal{F}'$.

Theorem 2.13 ([15]) *For every convex polygon P there is a constant $c_f := c_f(P)$ such that the family of all homothets of P is self-coverable with $f(l) \leq c_f l$.*

Theorem 2.14 ([15]) *The family of all homothets of a square is self-coverable with $f(l) := 2l + 2$ and this is sharp.*

Theorem 2.15 ([15]) *The family of all homothets of a given triangle is self-coverable with $f(l) := 2l + 1$ and this is sharp.*

Theorem 2.16 ([15, Thm. 2]) *If \mathcal{F} is self-coverable with a monotone self-coverability function $f(l) > l$ and any finite set of points can be colored with two colors such that any member of \mathcal{F} with at least m points contains both colors, then any finite set of points can be colored with k colors such that any member of \mathcal{F} with at least $m_k := m(f(m - 1))^{\lceil \log k \rceil - 1} \leq k^d$ points contains all k colors (where d is a constant that depends only on \mathcal{F}).⁴*

Theorem 1.3 (which we have yet to prove) and Theorems 2.14 and 2.16 immediately imply Corollary 1.4. Indeed, the required assumptions of Theorem 2.16 hold for

⁴ Unless stated otherwise, logarithms in this paper are base 2.

squares with $m \leq 215$ by Theorem 1.3 and $f(l) = 2l + 2$ by Theorem 2.14. We get $m_k = f(m - 1)^{\lceil \log k \rceil - 1} \leq (2m)^{\log k} = k^{\log 430} = O(k^{8.75})$ for squares, and also for parallelograms by affine transformations.

3 A 2-Coloring Algorithm

In this section we prove Theorems 1.3 and 1.5. In fact, we prove a more general result, for which we need the following definitions.

Definition 3.1 (Good paths and good homothets) Let P be an (open or closed) convex polygon, let S be a finite set of points, let $\mathcal{DT} := \mathcal{DT}(P, S)$, and let P' be a homothet of P .

- Let x - y - z be a 2-path in \mathcal{DT} (i.e., a simple path of length two). If $\text{Cl}(P')$ does not contain x and z and y is in the interior of P' , then we say that it *separates* the 2-path x - y - z .
- A 2-path x - y - z is *good*, if there is no homothet of P that separates it such that the edges yx and yz cross the same side of this homothet of P (see Fig. 2 for an example).
- A 3-path x - y - z - w in \mathcal{DT} is *good* if both x - y - z and y - z - w are good 2-paths.
- P' is *good* if it contains a good 3-path or $\mathcal{DT}[P']$ contains a cycle.

Observe that whether a 2-path x - y - z is good depends only on the direction of the vectors \vec{yx} and \vec{yz} .

Definition 3.2 (Universally good polygons) We say that an (open or closed) convex polygon P is *universally good* with a constant $c_g := c_g(P)$ if for any finite set of points S such that S is in very general position with respect to P and $\mathcal{DT}(P, S)$ is nice, every homothet of P that contains at least c_g points from S is good.

Theorem 3.3 Let P be an (open or closed) convex polygon with n vertices such that P is a universally good polygon with a constant $c_g := c_g(P)$, and let $f(l) \leq c_f l$ be a self-coverability function of the family of homothets of $\text{Cl}(P)$ (where $c_f := c_f(P)$ is a constant). Then there is a constant $m := m(P) \leq (c_g - 1)f(n) + n + 1 \leq (c_g - 1)c_f n + n + 1$ such that it is possible to 2-color in polynomial time the points of any given finite set S such that every homothet of P that contains at least m points from S contains points of both colors.




Fig. 2 Considering homothets of an axis-parallel square, x - y - z is a good 2-path whereas x - y - q is not since the square Q' separates it and both xy and qy cross the left side of Q'

We note that for the existence of m in Theorem 3.3 we need only that $f(l)$ exists. However, using that it is upper bounded by $c_f l$ [15] we can get a more explicit bound on m .

Theorems 1.3 and 1.5 immediately follow from Theorems 2.14, 2.15, 3.3, and the following two lemmas.

Lemma 3.4 *Every triangle is a universally good polygon with a constant $c_g \leq 7382$.*

Lemma 3.5 *Every parallelogram is a universally good polygon with a constant $c_g \leq 22$.*

In particular, we get the value in Theorem 1.3 by taking $m = (c_g - 1)f(n) + n + 1 = 215$ with $c_g = 22$, $f(l) = 2l + 2$ and $n = 4$ for squares.

In light of Theorem 3.3, it is enough to prove that a convex polygon is universally good to conclude that Problem 1.1 (i) has a positive solution with respect to homothets of that polygon. However, as it turns out, parallelograms and triangles are the only universally good polygons.

Theorem 3.6 *Let P be a convex polygon which is neither a triangle nor a parallelogram. Then P is not universally good.*

We proceed with the proof of Theorem 3.3, then prove that triangles and parallelograms are universally good, and conclude this section with a proof that no other universally good polygons exist.

3.1 Proof of Theorem 3.3

Let P be an (open or closed) convex polygon with n vertices and let $\bar{P} := \text{Cl}(P)$ be the closure of P (thus, \bar{P} is a closed polygon and $\bar{P} = P$ if P is closed). Let us assume also that P is a universally good polygon with a constant $c_g := c_g(P)$, and let $f(l) \leq c_f l$ be a self-coverability function of the family of homothets of \bar{P} . Set $m := (c_g - 1)f(n) + n + 1 \leq (c_g - 1)c_f n + n + 1$. We first argue that it is enough to prove Theorem 3.3 when P is a closed polygon. Indeed, suppose that P is open and let \mathcal{P} be the family of homothets of P . By slightly shrinking every homothet of P in \mathcal{P} with respect to an interior point, we get a family \mathcal{P}' of homothets that realizes $H^{\mathcal{P}}(S)$ such that there is no point $p \in S$ and homothet $P' \in \mathcal{P}'$ with $p \in \partial P'$.

Note that by definition \bar{P} is universally good with the same constant c_g and is self-coverable with the same self-coverability function as P . Let $\bar{\mathcal{P}}' := \{\text{Cl}(P') : P' \in \mathcal{P}'\}$. Since there is no homothet of P in \mathcal{P}' that contains a point of S on its boundary, every hyperedge of $H^{\mathcal{P}}(S)$ appears also in $H^{\bar{\mathcal{P}}'}(S)$. Thus, it is enough to show that \bar{P} satisfies Theorem 3.3.

Suppose therefore that P is a closed convex polygon. Let \mathcal{P} be the family of homothets of P and let $\mathcal{P}_0 \subseteq \mathcal{P}$ be a smallest size subfamily that realizes $H^{\mathcal{P}}(S)$. For convenience we pick \mathcal{P}_0 such that no $P' \in \mathcal{P}_0$ contains a point from S on its boundary (this can be achieved by slightly inflating homothets if necessary).

We may also assume that S is in very general position with respect to P . Indeed, otherwise note that a small perturbation of the points will achieve that while \mathcal{P}_0

will still realize the same hypergraph $H^{\mathcal{P}}(\mathcal{S})$. It will also be convenient to assume that the boundaries of every two polygons in \mathcal{P}_0 do not overlap, and no edge in $\mathcal{DT} := \mathcal{DT}(P, \mathcal{S})$ crosses the boundary of a polygon in \mathcal{P}_0 at one of its vertices. It follows from Lemma 2.1 that \mathcal{P}_0 is a family of *pseudo-disks*.⁵ This implies that $|\mathcal{P}_0| = O(|\mathcal{S}|^3)$ by a result of Buzaglo et al. [5] who proved the following: Suppose that $(\mathcal{V}, \mathcal{E})$ is a hypergraph where \mathcal{V} is a set of points in the plane and for every hyperedge $e \in \mathcal{E}$ there is a region bounded by a simple closed curve that contains the points of e and no other points from \mathcal{V} . If all the regions that correspond to \mathcal{E} define a family of pseudo-disks, then $|\mathcal{E}| = O(|\mathcal{V}|^3)$.

We can also assume that \mathcal{DT} is nice, that is, the boundary of its outer face is a convex polygon: Set $-P := \{(-x, -y) : (x, y) \in P\}$ and let $-P'$ be a homothet of $-P$ that contains in its interior all the polygons in \mathcal{P}_0 . By adding the vertices of $-P'$ to \mathcal{S} (and perturbing again if needed) we obtain a set of points \mathcal{S}' such that $-P'$ is the boundary of the outer face in its generalized Delaunay triangulation with respect to P . Moreover, the hypergraph we get by intersecting homothets of P with \mathcal{S}' is a hypergraph that contains all the hyperedges that we get by intersecting polygons in \mathcal{P}_0 with \mathcal{S} . The latter hyperedges are exactly the hyperedges we get by intersecting homothets of P with \mathcal{S} , since \mathcal{P}_0 realizes $H^{\mathcal{P}}(\mathcal{S})$. Therefore a valid 2-coloring of the new set of points induces a valid 2-coloring of the original set of points.

Recall that \mathcal{DT} is a plane graph, and therefore, by the Four Color Theorem, we can color the points in \mathcal{S} with four colors, say 1, 2, 3, 4, such that there are no adjacent vertices in \mathcal{DT} with the same color. In order to obtain two color classes, we recolor all the vertices of colors 1 or 2 with the color *light red* and all the vertices of colors 3 or 4 with the color *light blue*.

Call a homothet $P' \in \mathcal{P}_0$ *heavy monochromatic* if it contains exactly c_g points from \mathcal{S} and all of them are of the same light color. If all of these points are colored light blue (resp., red), then we call P' a heavy light blue (resp., red) homothet. Obviously, if there are no heavy monochromatic homothets, then we are done since $m > c_g$ and it follows from Lemma 2.3 that a monochromatic homothet with $m > c_g$ points from \mathcal{S} can be shrunk to a monochromatic homothet with exactly c_g points from \mathcal{S} .

Suppose that P' is a heavy monochromatic homothet of P . Observe that $\mathcal{DT}[P']$ is a tree, for otherwise it would contain a cycle which in turn would contain a triangle by Lemma 2.2. That triangle must be 3-colored in the initial 4-coloring, so not all of its points can be light red or light blue, contradicting the monochromaticity of the points in P' .

Since P is universally good, P' contains c_g points and $\mathcal{DT}[P']$ is a tree, it follows that P' contains a good 3-path x - y - z - w . We associate this 3-path with P' . Suppose that P' is a heavy light red homothet of P . Then one of y and z was originally colored 1 and the other was originally colored 2. Recolor the one whose original color was 1 with the color *dark blue*. Similarly, if P' is a heavy light blue homothet of P , then one of y and z was originally colored 3 and the other was originally colored 4. In this case we recolor the one whose original color was 3 with the color *dark red*. Repeat this for every heavy monochromatic homothet, and, finally, in order to obtain a 2-coloring,

⁵ In a family of pseudo-disks the boundaries of every two regions cross at most twice.

merge the color classes light red and dark red into one color class, red; and merge the color classes light blue and dark blue into one color class, blue.

Lemma 3.7 *There is no homothet $P' \in \mathcal{P}_0$ that contains m points from \mathcal{S} all of which are of the same color.*

Proof Suppose for contradiction that P' is a homothet of P that contains m points from \mathcal{S} all of which of the same color. We may assume without loss of generality that all the points in P' are colored red, therefore, before the final recoloring each point in P' was either light red or dark red. We consider two cases based on the number of dark red points in P' . Recall that n is the number of vertices of P .

Case 1: There are at most n dark red points in P' . By Definition 2.12 there is a set \mathcal{P}' of at most $f(n)$ homothets of P whose union is P' such that no dark red point in P' is in the interior of one of these homothets. Using Lemma 2.3 we can change these homothets slightly such that none of them contains a dark red point yet all light red points are still covered by these homothets. Thus the at least $m - n = (c_g - 1)f(n) + 1$ light red points are covered by these at most $f(n)$ homothets. By the pigeonhole principle one of these homothets, denote it by P'' , contains at least $\lceil \frac{(c_g - 1)f(n) + 1}{f(n)} \rceil = c_g$ light red points and no other points. However, in this case it follows from Lemma 2.3 that there is a heavy light red homothet in \mathcal{P}_0 that contains exactly c_g points from $\mathcal{S} \cap P''$. Therefore, the coloring algorithm should have found within this heavy light red homothet a good 3-path and recolored one of its vertices with dark blue and then blue. This contradicts the assumption that all the points in P' are red.

Case 2: There are more than n dark red points in P' . Let y be one of these dark red points. Then there is a good 3-path x - y - z - w within a heavy light blue homothet $P_y \in \mathcal{P}_0$ with whom this 3-path is associated. Furthermore, the original color of y is 3 and therefore the original color of x and z is 4, and thus their final color is blue. It follows that P' separates x - y - z , moreover, since x - y - z is a good 2-path, the edges yx and yz cross different sides of P' . Let s_x be the side of P' that is crossed by yx , and let q_x be the crossing point of yx and s_x . Similarly, let s_z be the side of P' that is crossed by yz , and let q_z be the crossing point of yz and s_z . See Fig. 3. Note that $\partial P'$ and ∂P_y cross each other exactly twice. Indeed, this follows from Lemma 2.1 and the fact that there are points from \mathcal{S} in each of $P' \cap P_y$ (e.g., y), $P_y \setminus P'$ (e.g., x and z) and $P' \setminus P_y$ (since $|P' \cap \mathcal{S}| \geq m > c_g = |P_y \cap \mathcal{S}|$). The points q_x and q_z partition $\partial P'$ into two parts $\partial P'_1$ and $\partial P'_2$. Note that since $q_x, q_z \in P' \cap P_y$, the two crossing points between $\partial P'$ and ∂P_y must lie either in $\partial P'_1$ or in $\partial P'_2$. Assume without loss of generality that both of them lie in $\partial P'_1$. Thus $\partial P'_2 \subset P_y$. Let v be a vertex of P' in $\partial P'_2$ (note that since $s_x \neq s_z$ each of $\partial P'_1$ and $\partial P'_2$ contains a vertex of P'). We associate the vertex v with the dark red point y . We also define R_y to be the region whose boundary consists of the segments yq_x, yq_z , and the part of $\partial P'_2$ whose endpoints are q_x and q_z (call this part $\partial P'_{xz}$). Observe that $R_y \subseteq P' \cap P_y$.

Proposition 3.8 *There is no other point but y in $\mathcal{S} \cap R_y$.*

Proof Suppose that the claim is false and let $y' \in \mathcal{S} \cap R_y$ be another point in R_y . As y' is in P' , it must be red after the final coloring. Also, as it is also in P_y , it must be




Fig. 3 An illustration for the proof of Lemma 3.7

a dark red point (which was light blue before having been recolored to dark red and finally to red). Thus, y' is a dark red point in R_y .

Since x and y' both lie in the heavy light blue homothet P_y , they are connected by a path in $\mathcal{DT}[P_y]$ that alternates between points of colors 3 and 4 (considering the initial 4-coloring). We may assume without loss of generality that y' is the first point in R_y along this path from x to y' : indeed, there are no points of color 4 in R_y , and if there is a point of color 3 before y' , then we can name it y' . Denote by ℓ the path (in \mathcal{DT}) from y to y' that consists of the edge yx and the above-mentioned path from x to y' . Consider the polygon \hat{P} whose boundary consists of ℓ and a straight-line segment yy' (\hat{P} is not a homothet of P). Since y' and y are the only vertices of \hat{P} in R_y , there is no edge of ℓ that crosses yy' . Indeed, if there was such an edge, then it would split P' into two parts such that one contains y and the other contains y' . This would contradict Corollary 2.5. Hence \hat{P} is a simple polygon.

Since every simple polygon has at least three convex vertices, \hat{P} has a convex vertex different from y and y' (thus this vertex is not in R_y). Denote this vertex by b and let a and c be its neighbors along ℓ such that $\angle abc < \pi$. Since the initial colors of a and c are the same, we have $ac \notin \mathcal{DT}$ and so it follows from Proposition 2.6 that there is a neighbor d of b in \mathcal{DT} in between a and c . Let us choose d such that it is the neighbor of b that is closest to a in the rotation of b . Thus, it is connected to both a and b . Since the initial colors of a and b are 3 and 4, the initial color of d was 1 or 2. Note that $\hat{P} \subseteq P_y$ since all of its edges are inside P_y . Thus $d \notin \hat{P}$ and also $d \notin P_y$ since P_y does not contain vertices of color 1 or 2. Now consider the directed edge bd : it starts inside \hat{P} (since d is in between a and c) and so it must cross yy' . Before doing so bd must cross ∂R_y and so it crosses $\partial P'_{xz}$, since it cannot cross yq_z or yq_x . After crossing yy' , the directed edge bd must cross $\partial P'_{xz}$ again, since $d \notin R_y$. But then bd splits P' into two parts such that one contains y and the other contains y' , which is impossible by Corollary 2.5. \square

In a similar way to the one described above, we associate a vertex of P' with every dark red point in P' . Since there are more than n dark red points in P' , there are two of them, denote them by y and y' , that are associated with the same vertex of

P' , denote it by v . Let x - y - z be the good 2-path that corresponds to y , let yq_x and yq_z be the edge-segments of yx and yz , and let R_y be the region as defined above. Similarly, let x' - y' - z' be the good 2-path that corresponds to y' , let $yq_{x'}$ and $yq_{z'}$ be the edge-segments of $y'x'$ and $y'z'$, and let $R_{y'}$ be the region as defined above.

It follows from Proposition 3.8 that $y \notin R_{y'}$ and $y' \notin R_y$. However, ∂R_y and $\partial R_{y'}$ both contain v . This implies that one of the segments yq_x and yq_z crosses one of the segments $y'q_{x'}$ and $y'q_{z'}$, which is impossible since these are segments of edges of a plane graph. Lemma 3.7 is proved. \square

To complete the proof of Theorem 3.3, we need to argue that the described algorithm runs in polynomial time. Indeed, constructing the generalized Delaunay triangulation and then 4-coloring it can be done in polynomial time. Recall that there are at most $O(|S|^3)$ combinatorially different homothets of P . Among them, we need to consider those that contain exactly c_g points, and for each such heavy monochromatic homothet P' we need to find a good 3-path in $\mathcal{DT}[P']$, for the final recoloring step. This takes a constant time for every heavy monochromatic homothet, since c_g is a constant. Therefore, the overall running time is polynomial with respect to the size of S .

3.2 Triangles are Universally Good

In this section we prove Lemma 3.4.

Let T be a triangle, let \mathcal{S} be a set of points in very general position with respect to T , and let $\mathcal{DT} := \mathcal{DT}(T, \mathcal{S})$ be the generalized Delaunay triangulation of \mathcal{S} with respect to T such that \mathcal{DT} is nice (i.e., the boundary of its outer face is a convex polygon). By applying an affine transformation, if needed, we may assume without loss of generality that T is an equilateral triangle. Suppose that T' is a homothet of T that contains at least 7382 points from \mathcal{S} and that $\mathcal{DT}[T']$ is a tree. We will show that T' contains a good 3-path.

By Corollary 2.10 for every point $v \in T' \cap \mathcal{S}$ we have $\deg_{\mathcal{DT}[T']}(v) \leq 9$. Since $\mathcal{DT}[T']$ is a tree with at least $7382 = 1 + 9 + 9^2 + 9^3 + 9^4 + 1$ vertices of maximum degree 9, it contains a simple path of length 9. Let $Z = v_1-v_2-\dots-v_{10}$ be such a path. We will prove that there is $2 \leq i \leq 8$ such that $v_{i-1}-v_i-v_{i+1}$ and $v_i-v_{i+1}-v_{i+2}$ are good 2-paths, and therefore T contains the good 3-path $v_{i-1}-v_i-v_{i+1}-v_{i+2}$. Call a 2-path $v_{i-1}-v_i-v_{i+1}$ (for $2 \leq i \leq 9$) *bad* if it is not good, that is, there is a homothet of T , T_i , such that T_i contains v_i , does not contain v_{i-1} and v_{i+1} , and the edges $v_i v_{i-1}$ and $v_i v_{i+1}$ cross the same side of T_i .

Denote the sides of T by s_1, s_2, s_3 . For $j = 1, 2, 3$, let B_j be the set of bad 2-paths $v_{i-1}-v_i-v_{i+1}$ such that there is a homothet T_i that contains v_i and does not contain v_{i-1} and v_{i+1} , and the edges $v_i v_{i-1}$ and $v_i v_{i+1}$ both cross the side of T_i that is homothetic to s_j . Suppose for contradiction that Z does not contain two consecutive good 2-paths. Then, at least one of the sets B_j contains two bad 2-paths. Assume without loss of generality that B_1 contains two bad 2-paths $v_{i-1}-v_i-v_{i+1}$ and $v_{k-1}-v_k-v_{k+1}$ such that $i < k$. We may further assume that s_1 is horizontal and that T lies above it.

There is a homothet of T that separates $v_{i-1}-v_i-v_{i+1}$ such that $v_{i-1}v_i$ and $v_i v_{i+1}$ both cross its side that is homothetic to s_1 , therefore both v_{i-1} and v_{i+1} lie below v_i . Similarly, both v_{k-1} and v_{k+1} lie below v_k . Let v_r be the lowest point among v_i, \dots, v_k .




Fig. 4 An illustration for the proof of Lemma 3.4

Since v_{i+1} is lower than v_i and v_{k-1} is lower than v_k it follows that $r \neq i, k$ and so v_r is lower than v_{r-1} and v_{r+1} . Suppose without loss of generality that v_{r+1} is to the right of the line through v_r and v_{r-1} . By applying Proposition 2.6 to two vertices that are between v_{r-1} and v_{r+1} and follow each other immediately in the rotation of v_r in $\mathcal{DT}[T']$ (they can coincide with v_{r-1} and/or v_{r+1}), we get that there is at least one neighbor of v_r between v_{r-1} and v_{r+1} that lies outside of T' . Let u be such a neighbor of v_r and let T_u be a homothet of T that contains v_r and u and no other point from \mathcal{S} . Note that u is higher than v_r , thus $v_r u$ crosses either the right or the left side of T' . Suppose without loss of generality that $v_r u$ crosses the right side of T' at a point q_u (refer to Fig. 4). It follows that the right side of T_u is to the right of the right side of T' . Thus, a horizontal ray that begins at v_r and goes to the right will first cross the right side of T' (denote this crossing point by q) and then cross the right side of T_u (note that this ray does not cross the left sides of T_u and T' since $v_r \in T_u \cap T'$). Now consider the triangle $\Delta q_u q v_r$. All of its vertices are in $T_u \cap T'$, therefore $\Delta q_u q v_r \in T_u \cap T'$. However, since v_{r+1} follows u in the rotation of v_r , it follows that the edge $v_r v_{r+1}$ lies in $\Delta q_u q v_r$ since it cannot cross none of its sides. This is impossible since v_{r+1} should be outside of T_u and hence outside of $\Delta q_u q v_r$. Lemma 3.4 is proved.

3.3 Parallelograms are Universally Good

In this section we prove Lemma 3.5. Let Q be a parallelogram, let \mathcal{S} be a set of points in very general position with respect to Q , and let $\mathcal{DT} := \mathcal{DT}(Q, \mathcal{S})$ be the generalized Delaunay triangulation of \mathcal{S} with respect to Q such that \mathcal{DT} is nice (i.e., the boundary of its outer face is a convex polygon). By applying an affine transformation, we may assume without loss of generality that Q is an axis-parallel square. Since \mathcal{S} is in very general position, no two points in \mathcal{S} share the same x - or y -coordinate.

Suppose that Q' is a homothet of Q that contains at least 22 points from \mathcal{S} and that $\mathcal{DT}[Q']$ is a tree. We will show that Q' contains a good 3-path.

Let $q \in \mathcal{S}$ be a point. We partition the points of the plane into four open quadrants according to their position with respect to q : NE(q) (North-East), NW(q) (North-West), SE(q) (South-East), and SW(q) (South-West).




Fig. 5 An illustration for the proof of Proposition 3.10

Proposition 3.9 *Let x, y, z be three points in \mathcal{S} such that xy and xz are edges in \mathcal{DT} . Then for every quadrant $Q_d \in \{\text{NW}, \text{NE}, \text{SW}, \text{SE}\}$ if $y \in Q_d(x)$, then $z \notin Q_d(y)$.*

Proof Suppose for contradiction and without loss of generality that $y \in \text{NE}(x)$ and $z \in \text{NE}(y)$. Then the smallest rectangle that contains x and z has x at its bottom-left corner, z at its top-right corner and y in its interior. Therefore, there is no square that contains x and z and does not contain y and so xz cannot be an edge in \mathcal{DT} . \square

Proposition 3.10 *For every point $q \in \mathcal{S} \cap Q'$ there are no two neighbors of q in $\mathcal{DT}[Q']$ that lie in the same quadrant of q .*

Proof Suppose for contradiction that q has two neighbors, x and y , that lie in the same quadrant. In this case we can choose them such that there is no other neighbor of q between them. Assume without loss of generality that $x, y \in \text{NE}(q)$, such that qx forms a smaller angle with the x -axis than qy (refer to Fig. 5) and they follow each other immediately in the rotation of q . It follows from Proposition 2.6 that there is a point $z \notin Q'$ such that z is a neighbor of q in \mathcal{DT} and is between x and y in the rotation of q . By Proposition 3.9 we have $y \notin \text{NE}(x)$. Since qx forms a smaller angle with the x -axis than qy we have $y \notin \text{SE}(x)$. If $y \in \text{SW}(x)$, then $x \in \text{NE}(y)$ which is impossible by Proposition 3.9. Thus, $y \in \text{NW}(x)$. Using the same arguments we get that $z \in \text{NW}(x) \cap \text{SE}(y)$. However, this implies that z is contained in any axis-parallel rectangle that contains x and y and thus $z \in Q'$, a contradiction. \square

Proposition 3.11 *Let x and y be two neighbors of q in $\mathcal{DT}[Q']$. Let $z \notin Q'$ be a neighbor of q in \mathcal{DT} that lies between x and y in the rotation of q and let Q_z be a square that contains q and z and no other point from \mathcal{S} . Then:*

- if $x \in \text{NW}(q)$ and $y \in \text{NE}(q)$, then qz crosses the top side of Q' , x is to the left of Q_z and y is to the right of Q_z ;
- if $x \in \text{NE}(q)$ and $y \in \text{SE}(q)$, then qz crosses the right side of Q' , x is above Q_z and y is below Q_z ;
- if $x \in \text{SE}(q)$ and $y \in \text{SW}(q)$, then qz crosses the bottom side of Q' , x is to the right of Q_z and y is to the left of Q_z ; and
- if $x \in \text{SW}(q)$ and $y \in \text{NW}(q)$, then qz crosses the left side of Q' , x is below Q_z and y is above Q_z .

Proof By symmetry it is enough to consider the first case, that is, $x \in \text{NW}(q)$ and $y \in \text{NE}(q)$. Since z is between x and y in the rotation of q we have $z \notin \text{SW}(x)$ and $z \notin \text{SE}(y)$. By Proposition 3.9, $z \notin \text{NW}(x)$ and $z \notin \text{NE}(y)$. Thus z is to the right of x and to the left of y . It follows that z is above Q' and qz crosses the top side of Q' .

Therefore, the top side of Q_z is above Q' . Thus, qx cannot cross the top side of Q_z so it must cross its left side. This implies that x lies to the left of Q_z . Similarly, qy cannot cross the top side of Q_z so it must cross its right side. This implies that y lies to the right of Q_z . \square

Call a (simple) path in DT x -monotone (resp., y -monotone) if there is no vertical (resp., horizontal) line that intersects the path in more than one point.

Proposition 3.12 *Every path in $DT[Q']$ is x -monotone or y -monotone.*

Proof Suppose for contradiction that there is a path $p := q_1q_2 \cdots q_k$ which is neither x -monotone nor y -monotone. Since p is a polygonal path, it follows that there are two points, q_i and q_j , that are “witnesses” to the non- x - and non- y -monotonicity of p , respectively. That is, both q_{i-1} and q_{i+1} are to the left of q_i or both of them are to its right, and both q_{j-1} and q_{j+1} are above q_j or both of them are below q_j . We choose i and j such that $|i - j|$ is minimized, and assume without loss of generality that $i < j$ (note that it follows from Proposition 3.10 that $i \neq j$). Thus, the sub-path $p' := q_iq_{i+1} \cdots q_{j-1}q_j$ is both x -monotone and y -monotone.

By reflecting about the x - and/or y -axis if needed, we may assume that p' is ascending, that is, for every $l = i, \dots, j - 1$ we have $q_{l+1} \in NE(q_l)$. Then it follows from Proposition 3.10 that $q_{i-1} \in SE(q_i)$ and $q_{j+1} \in SE(q_j)$. By applying Proposition 2.6 to two vertices that are between q_{i-1} and q_{i+1} and follow each other immediately in the rotation of q_i in $DT[Q']$ (they can coincide with q_{i-1} and/or q_{i+1}), we get that there is a point $x \notin Q'$ which is a neighbor of q_i and is between q_{i-1} and q_{i+1} in the rotation of q_i , and it follows from Proposition 3.11 that q_ix crosses the right side of Q' . The same argument implies that there is a point $y \notin Q'$ which is a neighbor of q_j and is between q_{j+1} and q_{j-1} in the rotation of q_j , such that q_jy crosses the bottom side of Q' . However, since q_j is to the right of q_i and above it, the edges q_ix and q_jy must cross, which is impossible. \square

Call a 2-path w - q - z *bad* if it is not good, that is, there is an axis-parallel square Q'' that contains q , does not contain w and z , and qw and qz are edges in DT that cross the same side of Q'' . We say that w - q - z is a *bad left* 2-path if qw and qz cross the left side of Q'' , and define *right*, *top*, and *bottom* bad 2-paths analogously.

Proposition 3.13 *Let w - q - z be a 2-path. Then:*

- w - q - z is a bad left 2-path iff $w \in SW(q)$ and $z \in NW(q)$, or vice versa;
- w - q - z is a bad right 2-path iff $w \in SE(q)$ and $z \in NE(q)$, or vice versa;
- w - q - z is a bad top 2-path iff $w \in NW(q)$ and $z \in NE(q)$, or vice versa; and
- w - q - z is a bad bottom 2-path iff $w \in SW(q)$ and $z \in SE(q)$, or vice versa.

Proof By symmetry it is enough to consider the first claim. If w - q - z is a bad left 2-path, then there is a square Q'' that separates it such that the edges qw and qz cross the left side of Q'' . Therefore, these edges go leftwards from q and so $w, z \in SW(q) \cup NW(q)$. It follows from Proposition 3.10 that $w \in SW(q)$ and $z \in NW(q)$, or vice versa.

For the other direction, assume without loss of generality that $w \in SW(q)$ and $z \in NW(q)$. Let Q'' be the square whose left side is the straight-line segment between $((q)_x - \varepsilon, (w)_y)$ and $((q)_x - \varepsilon, (z)_y)$, for some small $\varepsilon > 0$. Then Q'' separates w - q - z and both qw and qz cross its left side, therefore, w - q - z is a bad left 2-path. \square

Proposition 3.14 *Every path in $\mathcal{DT}[Q']$ contains at most four bad 2-paths.*

Proof Let $p := q_1-q_2-\dots-q_k$ be a simple path in $\mathcal{DT}[Q']$ and suppose for a contradiction that p contains at least five bad 2-paths. By Proposition 3.12 the path p is x -monotone or y -monotone. Assume without loss of generality that p is y -monotone and that it goes upwards, that is, q_{i+1} is above q_i for every $i = 1, 2, \dots, k - 1$. It follows that p does not contain bad top or bad bottom 2-paths, for otherwise it would not be y -monotone. It is not hard to see that bad left and bad right 2-paths must alternate along p , that is, between every two bad left 2-paths there is a bad right 2-path and vice versa.

Consider the first five such bad 2-paths along the path p , and denote them by $q_{i_1-1}-q_{i_1}-q_{i_1+1}$, $q_{i_2-1}-q_{i_2}-q_{i_2+1}$, $q_{i_3-1}-q_{i_3}-q_{i_3+1}$, $q_{i_4-1}-q_{i_4}-q_{i_4+1}$ and $q_{i_5-1}-q_{i_5}-q_{i_5+1}$. By symmetry we may assume without loss of generality that $q_{i_1-1}-q_{i_1}-q_{i_1+1}$ is a bad left 2-path, and therefore $q_{i_3-1}-q_{i_3}-q_{i_3+1}$ and $q_{i_5-1}-q_{i_5}-q_{i_5+1}$ are also bad left 2-paths, whereas the 2-paths $q_{i_2-1}-q_{i_2}-q_{i_2+1}$ and $q_{i_4-1}-q_{i_4}-q_{i_4+1}$ are bad right.

Note that we may assume without loss of generality that q_{i_1} is to the right of q_{i_4} , for otherwise q_{i_5} must be to the right of q_{i_2} and by reflecting about the x -axis and renaming the points we get the desired assumption. By applying Proposition 2.6 to two vertices that are between q_{i_1-1} and q_{i_1+1} and follow each other immediately in the rotation of q_{i_1} in $\mathcal{DT}[Q']$ (they can coincide with q_{i_1-1} and/or q_{i_1+1}), we get that q_{i_1} has a neighbor $z \notin Q'$ between q_{i_1-1} and q_{i_1+1} in the rotation of q_{i_1} . Let Q_z be a square that contains q_{i_1} and z and no other point from \mathcal{S} and let s_z be its side length (refer to Fig. 6). It follows from Proposition 3.11 that q_{i_1-1} lies below Q_z , q_{i_1+1} lies above Q_z , and z lies to the left of Q' . Therefore, $(q_{i_1+1})_y - (q_{i_1-1})_y > s_z$. Similarly, q_{i_4} has a neighbor $w \notin Q'$ between q_{i_4+1} and q_{i_4-1} in the rotation of q_{i_4} . Let Q_w be a square that contains q_{i_4} and w and no other point from \mathcal{S} and let s_w be its side length. Then q_{i_4-1} lies below Q_w , q_{i_4+1} lies above Q_w , and w lies to the right of Q' . Therefore, $(q_{i_4+1})_y - (q_{i_4-1})_y > s_w$.

Note that since q_{i_1} is to the right of q_{i_4} and z and w are to the left and to the right of Q' , respectively, we have $s_z + s_w > ((q_{i_1})_x - (z)_x) + ((w)_x - (q_{i_4})_x) > s_{Q'}$, where $s_{Q'}$ is the side length of Q' . Observe also that since there are at least two other vertices between q_{i_1} and q_{i_4} along p , we have that $q_{i_1+1} \neq q_{i_4-1}$, and thus q_{i_1+1} lies below q_{i_4-1} . This implies that $((q_{i_1+1})_y - (q_{i_1-1})_y) + ((q_{i_4+1})_y - (q_{i_4-1})_y) < s_{Q'}$. Combining the inequalities we get, $s_{Q'} > ((q_{i_1+1})_y - (q_{i_1-1})_y) + ((q_{i_4+1})_y - (q_{i_4-1})_y) > s_z + s_w > ((q_{i_1})_x - (z)_x) + ((w)_x - (q_{i_4})_x) > s_{Q'}$, a contradiction. \square

To complete the proof of Lemma 3.5 we will consider a path of length 11 in $\mathcal{DT}[Q']$. It follows from Proposition 3.10 that for every $q \in \mathcal{S} \cap Q'$ we have $\deg_{\mathcal{DT}[Q']}(q) \leq 4$. This implies that if Q' contains at least $1 + \sum_{i=1}^5 4^i = 1366$ points from \mathcal{S} , then $\mathcal{DT}[Q']$ contains a simple path of length at least 11. However, one can show that already 22 points suffice to guarantee the existence of a path of length 11. To prove this, we will need the following proposition.

Proposition 3.15 *There are at most two points in $\mathcal{S} \cap Q'$ whose degree in $\mathcal{DT}[Q']$ is greater than two. If one of these points has degree four, then no other point has degree greater than two.*




Fig. 6 An illustration for the proof of Proposition 3.14

Proof Suppose first that there is a point q_k with $\deg_{\mathcal{DT}[Q']} (q_k) = 4$ and another point q_0 with $\deg_{\mathcal{DT}[Q']} (q_0) \geq 3$ (see Fig. 7 (a) for an example). Let $p := q_0-q_1-q_2-\dots-q_k$ be the path connecting these points in the tree $\mathcal{DT}[Q']$. We may assume without loss of generality that $q_1 \in \text{SE}(q_0)$ and that q_0 has a neighbor $z_0 \in \text{SW}(q_0)$. In this case the path $z_0-q_0-q_1-q_2-\dots-q_k$ is not y -monotone. Therefore this path must be x -monotone by Proposition 3.12, and thus p is also x -monotone. It follows that $q_{k-1} \in \text{SW}(q_k) \cup \text{NW}(q_k)$. Since by Proposition 3.9 a point cannot have two neighbors in the same quadrant and the degree of q_k is four, it has another neighbor $z_k \neq q_{k-1}$ in $\text{SW}(q_k) \cup \text{NW}(q_k)$. Therefore, the path $z_0-q_0-q_1-\dots-q_k-z_k$ is not monotone, a contradiction.

Now suppose that each point in $\mathcal{DT}[Q']$ has degree at most three, and suppose for contradiction that there are at least three points in $S \cap Q'$ whose degree in $\mathcal{DT}[Q']$ is three. Since $\mathcal{DT}[Q']$ is a tree, there must exist a path $p := q_0-q_1-q_2-\dots-q_k$ ($k \geq 2$) between two points q_0 and q_k with degree three that contains a third point q_j ($0 < j < k$) with degree three. We can assume without loss of generality that $q_1 \in \text{SE}(q_0)$. Since the degree of q_0 is three, it has a neighbor in $\text{NE}(q_0) \cup \text{SW}(q_0)$.

We can also assume that q_0 has a neighbor $z_0 \in \text{SW}(q_0)$ (see Fig. 7 (b) for an example). In this case the path $z_0-q_0-q_1-\dots-q_k$ is not y -monotone, thus it must be x -monotone and so p is also x -monotone. It follows, that if q_i , for $i \in \{j, k\}$, has two neighbors in $\text{SW}(q_i) \cup \text{NW}(q_i)$, then for one of them, denote it by $z_i \neq q_{i-1}$, a path that ends with $q_{i-1}-q_i-z_i$ is not x -monotone and hence the path $z_0-q_0-q_1-\dots-q_{i-1}-q_i-z_i$ is not monotone. Therefore, q_i has two neighbors in $\text{SE}(q_i) \cup \text{NE}(q_i)$, for $i \in \{j, k\}$. It follows that q_j has a neighbor $y_j \neq q_{j+1}$ such that a path that starts with $y_j-q_j-q_{j+1}$ is not x -monotone, and q_k has a neighbor y_k such that a path that ends with $q_{k-1}-q_k-y_k$ is not y -monotone. Therefore the path $y_j-q_j-q_{j+1}-\dots-q_{k-1}-q_k-y_k$ is not monotone, a contradiction. \square

Lemma 3.16 *If Q' contains at least 22 points from S , then $\mathcal{DT}[Q']$ contains a simple path of length at least 11.*




Fig. 7 Illustrations for the proof of Proposition 3.15. (a) The degree of q_0 is at least three and the degree of q_k is four. The path $z_0-q_0-q_1-\dots-q_k-z_k$ is not monotone. (b) The degree of q_0, q_j and q_k is three. We can assume that q_0 has a neighbor $z_0 \in SW(q_0)$. In case none of q_j and q_k has a neighbor in both of their NW and SW quadrants, there is a path $y_j-q_j-\dots-q_k-y_k$ which is not monotone

Proof If $\mathcal{DT}[Q']$ contains a vertex q whose degree is four, then it follows from Proposition 3.15 that by deleting it we decompose the tree $\mathcal{DT}[Q']$ into four paths, having 21 vertices altogether. By the pigeonhole principle either the first two or the second two together have at least 11 vertices. Along with q they form a path with 12 vertices, as required.

If $\mathcal{DT}[Q']$ does not contain a vertex whose degree is four, then by Proposition 3.15 it contains at most two vertices with degree three, and no other vertex has degree greater than two. We can assume that we have exactly two vertices whose degree is three, q and q' , and let p be the path connecting them in $\mathcal{DT}[Q']$. By deleting these two vertices we obtain five paths: a path that consists of p without its two endpoints, two paths p_1, p_2 that are incident to q and two paths p_3, p_4 that are incident to q' . Considering the number of vertices in each of these paths we have $|V(p_1)| + |V(p_2)| + |V(p)| + |V(p_3)| + |V(p_4)| + |V(p)| \geq 22 + 2 = 24$, since $|V(p)| \geq 2$. Therefore, one of the paths formed by p_1, p, p_3 and p_2, p, p_4 must contain at least 12 vertices, as required. \square

Let $p := q_1-q_2-\dots-q_{12}$ be a simple path of length 11 in $\mathcal{DT}[Q']$. By Proposition 3.14 there are at most four bad 2-paths $q_{i-1}-q_i-q_{i+1}$ in p . Therefore, there is $2 \leq i \leq 10$ such that $q_{i-1}-q_i-q_{i+1}$ and $q_i-q_{i+1}-q_{i+2}$ are good 2-paths, and therefore Q' contains a good 3-path $q_{i-1}-q_i-q_{i+1}-q_{i+2}$. Lemma 3.5 is proved. \square

3.4 A Universally Good Polygon is Either a Triangle or a Parallelogram

In this section we prove that triangles and parallelograms are the only universally good polygons. That is, for any other polygon P we can construct a set of points \mathcal{S} such that there is a homothet of P that intersects $\mathcal{DT}(P, \mathcal{S})$ in a long path, while every other vertex of this path can be separated from its neighbors by the same side of a homothet of P (thus there is no good 3-path in P). See Fig. 8. The rest of this section contains the exact description and validity of this construction.

We start with a simple statement that will be used later.




Fig. 8 A construction showing that P is not universally good

Proposition 3.17 *Let P' be a homothet of a convex closed polygon P that is contained in P . If $\partial P \cap \partial P' \neq \emptyset$, then an edge of P' is contained in an edge of P (to which it is homothetic).*

Proof Suppose for contradiction that $\partial P \cap \partial P' \neq \emptyset$ and no edge of P' is contained in an edge of P . Then there is a vertex v' of P' that lies on ∂P such that the edges that are incident to v' lie in the interior of P . It is impossible thus that v' coincides with v , the vertex to which it is homothet, since then these vertices would be incident to edges of different slopes. v' cannot coincide with another vertex of P either, since then there is a direction in which v' is extreme in P' and a different vertex than v is extreme at P . Thus v lies on an edge of P . One of the endpoints of this edge must be v , for otherwise as before there is a direction in which v' is extreme in P' and a different vertex than v is extreme at P . But then, again, it follows that v and v' are incident to edges of different slopes. Therefore, there is an edge of P' which is contained in an edge of P , and since both P' and P are on the same side of the line through this edge, it follows that these edges are homothetic to each other. \square

Suppose that P is a convex closed n -gon which is neither a triangle nor a parallelogram. We will show that P is not universally good. Let uv be an edge of P which is not parallel to any other edge of P if $n = 4$ and an arbitrary edge otherwise. Assume without loss of generality that uv is horizontal, u is left of v , and P lies above uv .

In order to have only a single case to deal with, if P has a unique highest vertex, then we regard this vertex as two vertices joined by a horizontal edge of length 0. This way P always has a top horizontal edge yz (possibly of length 0) such that y is left of z and P lies below yz . Since P is neither a triangle nor a parallelogram it has a vertex $x \notin \{u, v, y, z\}$. Assume without loss of generality that x is on the clockwise polygonal chain from y to v on the boundary of P . Denote by y_0 the point on the clockwise polygonal chain from u to y on the boundary of P that has the same y -coordinate as x and observe that the line-segment y_0x lies in P .

Proposition 3.18 *For every two distinct points a and b that lie on y_0x there is a homothet of P , denote it by P' , such that P' is contained in P and $y_0x \cap P' = ab$.*

Proof We can obtain such a homothet P' as follows. Initially, set $P' := P$. Now shrink P' with respect to a until b lies on $\partial P'$. Then shrink P' with respect to b until a lies on $\partial P'$. Since P' is convex we have $y_0x \cap P' = ab$. \square

For a polygon R and a vertex $r \in R$, we denote by r^- and r^+ the vertices of R that precede and succeed r , respectively, in the clockwise order of the vertices of R . For a homothet P' (resp., P_i) of P and a vertex $r \in P$, we denote by r' (resp., r_i) the vertex of P' (resp., P_i) that is homothetic to r . Since, e.g., $(r_i)^+ = (r^+)_i$, we simply write $r_i^+, r_i^-, r_i'^-$, etc.

Given an integer $k \geq 1$, we describe a way to construct a set of k homothets of P , P_1, P_2, \dots, P_k , such that:

- (1) P contains P_i , for every $i \geq 1$;
- (2o) the edge $x_i^-x_i$ of P_i is contained in the edge x^-x of P , for every odd $i \geq 1$;
- (2e) the edge $x_i x_i^+$ of P_i is contained in the edge xx^+ of P , for every even $i \geq 2$;
- (3o) the edge $u_i v_i$ of P_i is contained in the open segment $y_{i-1}x$, for every odd $i \geq 1$;
and
- (3e) the edge $y_i z_i$ of P_i is contained in the open segment $v_{i-1}x$, for every even $i \geq 2$.

We will use the following proposition for the construction.

Proposition 3.19 *Let u^* be a point in the open line-segment y_0x . Then there is a homothet of P , P' , such that P' is contained in P , $u'v'$ is contained in the open line-segment u^*x , and x'^-x' is contained in x^-x .*

Proof For an arbitrary u' on u^*x we can fix a homothet of P , denote it by P'' , such that its edge $u''v''$ that is homothetic to uv coincides with $u'x$ (i.e., $u'' = u'$ and $v'' = x$). Clearly, by choosing u' close enough to x , the intersection of $\partial P''$ and ∂P is contained in the edge x^-x . Since P is convex, it follows that the vector $\overrightarrow{vv^-}$ forms a smaller angle with the positive x -axis than does the vector $\overrightarrow{xx^-}$. Therefore, the (open) edge $v''v''$ of P'' homothetic to vv^- lies outside of P . We now continuously shrink P'' with respect to u' until it is contained in P . Let P' be the resulting homothet of P (see Fig. 9). Clearly, P' is contained in P . Note that an edge of P' is contained in an edge of P by Proposition 3.17. Also, as P' is in P'' and the intersection of $\partial P''$ and ∂P is contained in x^-x , the only possibility is that x'^-x' is contained in x^-x . Observe also that since we had to shrink P'' , the vertex v' of P' that is homothetic to v lies on the open segment $u'x$. \square

By reflecting P about the x -axis and applying Proposition 3.19 we get:

Corollary 3.20 *Let u^* be a point in the open line-segment y_0x . Then there is a homothet of P , P' , such that P' is contained in P , $y'z'$ is contained in the open line-segment u^*x , and x'^+x' is contained in x^+x .*

By applying Proposition 3.19 with $u^* = y_0$ we get a homothet P_1 that satisfies Properties (1)–(3) such that v_1 is on the open segment y_0x . Suppose that we have




Fig. 9 An illustration for the proof of Proposition 3.19

homothets P_1, \dots, P_i that satisfy Properties (1)–(3). For an even i we apply Proposition 3.19 with $u^* = v_i$ to get P_{i+1} . For an odd i we apply Corollary 3.20 with $u^* = z_i$ to get P_{i+1} . Thus we obtain homothets of P that satisfy Properties (1)–(3). Note that P_i is contained in the closed half-plane that is bounded from below (resp., above) by the line containing y_0x for every odd (resp., even) $i \geq 1$.

Suppose for contradiction that P is universally good with a constant $c := c_g(P)$. We may assume that $c \equiv 0 \pmod 4$ (by increasing c if necessary). We will construct a set of points \mathcal{S} such that $|\mathcal{S} \cap P| = c$, $\mathcal{DT}(P, \mathcal{S})[P]$ is a path, and P does not contain a good 3-path. This will contradict the fact that P is universally good.

We begin with an empty set of points \mathcal{S} and a set of homothets of P , P_1, \dots, P_k , as above, for $k = c$. Next, for every $i = 1, \dots, c - 1$ we add a point p_i to \mathcal{S} as follows: if i is odd, then $p_i = u_i$, that is, it coincides with the vertex of P_i that is homothetic to u , whereas if i is even, then $p_i = y_i$, that is, it coincides with the vertex of P_i that is homothetic to y .

Proposition 3.21 *The points $x_1^-, x_3^-, \dots, x_{c-1}^-$ appear on x^-x in this order. The points $x_2^-, x_4^-, \dots, x_c^-$ appear on x^+x in this order.*

Proof By symmetry it is enough to prove the first statement. Suppose for contradiction that there are two odd indices $i_1 < i_2$ such that $x_{i_1}^-$ is closer than $x_{i_2}^-$ to x . Consider the triangles $\Delta u_{i_1} v_{i_1} x_{i_1}^-$ and $\Delta u_{i_2} v_{i_2} x_{i_2}^-$. Since $u_{i_2} v_{i_2}$ is to the right of $u_{i_1} v_{i_1}$ on y_0x , it follows that their boundaries cross at four points (see Fig. 10). However, these triangles are homothetic, and thus such crossing is impossible. \square

It follows from Proposition 3.21 that x_i^- (resp., x_i^+) is not contained in P_{i+2j} for every $i \geq 1$ and every j for which P_{i+2j} exists.

Let $\varepsilon > 0$ be some small positive constant that is much smaller than the smallest distance between any pair of distinct points at vertices of the above mentioned polygons (that is, P and P_1, \dots, P_c). For every odd i we fix a point q_i outside of P at distance εi from x_i^- in the direction of the normal to x^-x and add this point to \mathcal{S} . Similarly, for every even i we fix a point q_i outside of P at distance εi from x_i^+ in the direction of the normal to xx^+ and add this point to \mathcal{S} . Then, for every i , by slightly inflating P_i with respect to some inner point, we obtain a homothet of P , denote it by P'_i , such that $P'_i \cap \mathcal{S} = \{p_i, q_i\}$.




Fig. 10 An illustration for the proof of Proposition 3.21

Let $\mathcal{DT} := \mathcal{DT}(P, \mathcal{S})$ be the generalized Delaunay triangulation of \mathcal{S} with respect to P . It follows from the construction that (p_i, q_i) is an edge in \mathcal{DT} , for every $i = 1, \dots, c$. By Proposition 3.18 it follows that \mathcal{DT} also contains the edge (p_i, p_{i+1}) , for every $i = 1, \dots, c - 1$.

Next we apply a small perturbation of the points in \mathcal{S} and slightly scale and translate the polygons P_i , such that for every $i = 1, \dots, c$ we have:

1. P_i still contains the same (perturbed) points p_i and q_i and no other (perturbed) point from \mathcal{S} ;
2. if i is odd, then the point p_i lies slightly below y_0x ; and
3. if i is even, then the point p_i lies slightly above y_0x .

Thus, the above-mentioned edges of \mathcal{DT} of type (p_i, q_i) and (p_i, p_{i+1}) are still edges of \mathcal{DT} . See Fig. 8 for an example of the construction at this point.

To complete the construction we add some points to \mathcal{S} , as in the proof of Theorem 3.3, to obtain a nice set of points with respect to P and perturb the points to obtain a set of points in very general position with respect to P .

Observe that $\mathcal{DT}[P]$ consists of the path $p_1-p_2-\dots-p_c$. Therefore, if P contains a good 3-path, then it must be of the form $p_i-p_{i+1}-p_{i+2}-p_{i+3}$ for some $1 \leq i \leq c - 3$. Thus, it is enough to prove that for every even i the 2-path $p_{i-1}-p_i-p_{i+1}$ is not good. Suppose for contradiction that $p_{i-1}-p_i-p_{i+1}$ is a good 2-path for some even i . Recall that while we allowed $y = z$, we know that $u \neq v$, that is, uv is an edge of positive length. Since i is even, p_i lies slightly above y_0x whereas p_{i-1} and p_{i+1} lie slightly below y_0x . Let P' be a homothet of P such that the endpoints of its edge $u'v'$ that is homothetic to uv are on y_0x , and u' (resp., v') has the same x -coordinate as p_{i-1} (resp., p_{i+1}). It follows that p_i lies inside P' whereas p_{i-1} and p_{i+1} are outside of P' . Moreover, the edges $p_i p_{i-1}$ and $p_i p_{i+1}$ both cross the edge $u'v'$ of P'_i . Therefore, the 2-path $p_{i-1}-p_i-p_{i+1}$ is not a good 2-path, a contradiction.

4 Discussion

The main open problem related to our work is the following.

Problem 4.1 *Is it true that for every convex polygon P there is a constant $m := m(P)$ such that it is possible to 2-color any set of points \mathcal{S} such that every homothet of P that contains at least m points from \mathcal{S} contains points of both colors?*

By Theorem 3.3 it would be enough to show that every convex polygon is universally good. However Theorem 3.6 shows that no other polygon is universally good besides triangles and parallelograms, thus for other classes of convex polygons additional ideas are needed. We remark that in a recent manuscript using the techniques developed in this article Keszegh and Pálvölgyi [18] solved the above problem with three colors.

We conclude with two challenging related open problems. Considering coloring of points with respect to disks, recall that in [24] it is proved that there is no constant m such that any set of points in the plane can be 2-colored such that any (unit) disk that contains at least m points from the given set is non-monochromatic (that is, contains points of both colors). Coloring the points with four colors such that any disk that contains at least two points is non-monochromatic is easy since the (generalized) Delaunay graph is planar. Therefore, it remains an interesting open problem whether there is a constant m such that any set of points in the plane can be 3-colored such that any disk that contains at least m points is non-monochromatic (this problem was posed originally in [12, 13], for more general variants see also [18]).

Perhaps the most interesting problem of coloring geometric hypergraphs is to color a planar set of points \mathcal{S} with the minimum possible number of colors, such that every axis-parallel rectangle that contains at least two points from \mathcal{S} is non-monochromatic. It is known that $\Omega(\log(|\mathcal{S}|)/\log^2 \log(|\mathcal{S}|))$ colors are sometimes needed [9], and it is conjectured that $\text{polylog}(|\mathcal{S}|)$ colors always suffice. The latter holds when considering rectangles that contain at least three points [1], however, for the original question only polynomial upper bounds are known [2, 8, 11, 29].

Acknowledgements E. Ackerman’s research partially supported by ERC Advanced Research Grant No. 267165 (DISCONV). B. Keszegh’s research supported by the National Research, Development and Innovation Office – NKFIH under the Grant PD 108406 and K 116769 and by the János Bolyai Research Scholarship of the Hungarian Academy of Sciences. M. Vizer’s research supported by the National Research, Development and Innovation Office – NKFIH under the Grant SNN 116095. We thank the reviewers for reading our paper and for their valuable remarks. In particular, suggestions of one of reviewers helped to simplify some basic lemmas and improve the constant in our main theorem.

References

1. Ackerman, E., Pinchasi, R.: On coloring points with respect to rectangles. *J. Combin. Theory Ser. A* **120**(4), 811–815 (2013)
2. Ajwani, D., Elbassioni, K., Govindarajan, S., Ray, S.: Conflict-free coloring for rectangle ranges using $O(n^{382})$ colors. *Discrete Comput. Geom.* **48**(1), 39–52 (2012)
3. Aloupis, G., Cardinal, J., Collette, S., Langerman, S., Orden, D., Ramos, P.: Decomposition of multiple coverings into more parts. *Discrete Comput. Geom.* **44**(3), 706–723 (2010)
4. Bose, P., Carmi, P., Collette, S., Smid, M.: On the stretch factor of convex Delaunay graphs. *J. Comput. Geom.* **1**(1), 41–56 (2010)
5. Buzaglo, S., Pinchasi, R., Rote, G.: Topological hypergraphs. In: Pach, J. (ed.) *Thirty Essays on Geometric Graph Theory*, pp. 71–81. Springer, New York (2013)
6. Cardinal, J., Knauer, K., Micek, P., Ueckerdt, T.: Making triangles colorful. *J. Comput. Geom.* **4**(1), 240–246 (2013)

7. Cardinal, J., Knauer, K., Micek, P., Ueckerdt, T.: Making octants colorful and related covering decomposition problems. *SIAM J. Discrete Math.* **28**(4), 1948–1959 (2014)
8. Chan, T.M.: Conflict-free coloring of points with respect to rectangles and approximation algorithms for discrete independent set. In: Proceedings of the 28th Annual Symposium on Computational Geometry (SoCG'12), pp. 293–302. ACM, New York (2012)
9. Chen, X., Pach, J., Szegedy, M., Tardos, G.: Delaunay graphs of point sets in the plane with respect to axis-parallel rectangles. *Random Struct. Algorithms* **34**(1), 11–23 (2009)
10. Gibson, M., Varadarajan, K.: Optimally decomposing coverings with translates of a convex polygon. *Discrete Comput. Geom.* **46**(2), 313–333 (2011)
11. Har-Peled, S., Smorodinsky, Sh.: Conflict-free coloring of points and simple regions in the plane. *Discrete Comput. Geom.* **34**(1), 47–70 (2005)
12. Keszegh, B.: Weak conflict-free colorings of point sets and simple regions. In: Proceedings of the 19th Canadian Conference on Computational Geometry (CCCG'07), pp. 97–100 (2007)
13. Keszegh, B.: Coloring half-planes and bottomless rectangles. *Comput. Geom.* **45**(9), 495–507 (2012)
14. Keszegh, B., Pálvölgyi, D.: Octants are cover-decomposable. *Discrete Comput. Geom.* **47**(3), 598–609 (2012)
15. Keszegh, B., Pálvölgyi, D.: Convex polygons are self-coverable. *Discrete Comput. Geom.* **51**(4), 885–895 (2014)
16. Keszegh, B., Pálvölgyi, D.: Octants are cover-decomposable into many coverings. *Comput. Geom.* **47**(5), 585–588 (2014)
17. Keszegh, B., Pálvölgyi, D.: More on decomposing coverings by octants. *J. Comput. Geom.* **6**(1), 300–315 (2015)
18. Keszegh, B., Pálvölgyi, D.: Proper coloring of geometric hypergraphs. In: Symposium on Computational Geometry 2017 (SoCG'17). [arXiv:1612.02158](https://arxiv.org/abs/1612.02158)
19. Klein, R.: Concrete and Abstract Voronoi Diagrams. Lecture Notes in Computer Science, vol. 400. Springer, Berlin (1989)
20. Kovács, I.: Indecomposable coverings with homothetic polygons. *Discrete Comput. Geom.* **53**(4), 817–824 (2015)
21. Ma, L.: Bisectors and Voronoi diagrams for convex distance functions. PhD Thesis, FernUniversität Hagen, Hagen (2000). <http://www.fernuni.de/ks/download/pub/tr267.pdf>
22. Pach, J.: Decomposition of multiple packing and covering. 2. Kolloquium über Diskrete Geometrie, pp. 169–178. Institut für Mathematik der Universität Salzburg, Salzburg (1980)
23. Pach, J.: Covering the plane with convex polygons. *Discrete Comput. Geom.* **1**(1), 73–81 (1986)
24. Pach, J., Pálvölgyi, D.: Unsplittable coverings in the plane. *Adv. Math.* **302**, 433–457 (2016)
25. Pach, J., Pálvölgyi, D., Tóth, G.: Survey on decomposition of multiple coverings. In: Bárány, I., et al. (eds.) *Geometry—Intuitive, Discrete, and Convex*. Bolyai Society Mathematical Studies, vol. 24, pp. 219–257. Springer, Berlin (2013)
26. Pach, J., Tardos, G.: Coloring axis-parallel rectangles. *J. Combin. Theory Ser. A* **117**(6), 776–782 (2010)
27. Pach, J., Tardos, G., Tóth, G.: Indecomposable coverings. In: Akiyama, J., et al. (eds.) *Discrete Geometry, Combinatorics and Graph Theory*. Lecture Notes in Computer Sciences, vol. 4381, pp. 135–148. Springer, Berlin (2007)
28. Pach, J., Tardos, G., Tóth, G.: Indecomposable coverings. *Can. Math. Bull.* **52**(3), 451–463 (2009)
29. Pach, J., Tóth, G.: Conflict-free colorings. In: Aronov, B., et al. (eds.) *Discrete and Computational Geometry*. Algorithms and Combinatorics, vol. 25, pp. 665–671. Springer, Berlin (2003)
30. Pach, J., Tóth, G.: Decomposition of multiple coverings into many parts. *Comput. Geom.* **42**(2), 127–133 (2009)
31. Pálvölgyi, D.: Indecomposable coverings with concave polygons. *Discrete Comput. Geom.* **44**(3), 577–588 (2010)
32. Pálvölgyi, D.: Cover decomposition page. <http://www.cs.elte.hu/~dom/covdec>
33. Pálvölgyi, D., Tóth, G.: Convex polygons are cover-decomposable. *Discrete Comput. Geom.* **43**(3), 483–496 (2010)
34. Sariöz, D.: Generalized Delaunay graphs with respect to any convex set are plane graphs (2010). [arXiv:1012.4881](https://arxiv.org/abs/1012.4881)
35. Smorodinsky, Sh.: Conflict-free coloring and its applications. In: Bárány, I. (ed.) *Geometry—Intuitive, Discrete, and Convex*. Bolyai Society Mathematical Studies, vol. 24, pp. 331–389. Springer, Berlin (2013)

36. Smorodinsky, Sh., Yuditsky, Y.: Polychromatic coloring for half-planes. *J. Combin. Theory Ser. A* **119**(1), 146–154 (2012)
37. Tardos, G., Tóth, G.: Multiple coverings of the plane with triangles. *Discrete Comput. Geom.* **38**(2), 443–450 (2007)