



HAL
open science

Rainbow polygons for colored point sets in the plane *

David Flores-Peñaloza, Mikio Kano, Leonardo Martínez-Sandoval, David Orden, Javier Tejel, Csaba D Tóth, Jorge Urrutia, Birgit Vogtenhuber

► To cite this version:

David Flores-Peñaloza, Mikio Kano, Leonardo Martínez-Sandoval, David Orden, Javier Tejel, et al.. Rainbow polygons for colored point sets in the plane *. *Discrete Mathematics*, 2021, 344 (7), pp.112406. 10.1016/j.disc.2021.112406 . hal-03375092

HAL Id: hal-03375092

<https://hal.sorbonne-universite.fr/hal-03375092>

Submitted on 12 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Rainbow polygons for colored point sets in the plane*

David Flores-Peñaloza[†] Mikio Kano[‡] Leonardo Martínez-Sandoval[§]
 David Orden[¶] Javier Tejel^{||} Csaba D. Tóth^{**} Jorge Urrutia^{††}
 Birgit Vogtenhuber^{‡‡}

Abstract

Given a colored point set in the plane, a *perfect rainbow polygon* is a simple polygon that *contains* exactly one point of each color, either in its interior or on its boundary. Let $\text{rb-index}(S)$ denote the smallest size of a perfect rainbow polygon for a colored point set S , and let $\text{rb-index}(k)$ be the maximum of $\text{rb-index}(S)$ over all k -colored point sets in general position; that is, every k -colored point set S has a perfect rainbow polygon with at most $\text{rb-index}(k)$ vertices. In this paper, we determine the values of $\text{rb-index}(k)$ up to $k = 7$, which is the first case where $\text{rb-index}(k) \neq k$, and we prove that for $k \geq 5$,

$$\frac{40\lfloor(k-1)/2\rfloor - 8}{19} \leq \text{rb-index}(k) \leq 10 \left\lfloor \frac{k}{7} \right\rfloor + 11.$$

Furthermore, for a k -colored set of n points in the plane in general position, a perfect rainbow polygon with at most $10 \lfloor \frac{k}{7} \rfloor + 11$ vertices can be computed in $O(n \log n)$ time.

1 Introduction

Given a colored point set in the plane, in this paper we study the problem of finding a simple polygon containing exactly one point of each color. Formally, the problem we consider is the following. Let $k \geq 2$ be an integer and let $\{1, \dots, k\}$ be k distinct colors. For every $1 \leq i \leq k$, let S_i denote a finite set of points of color i in the plane. We always assume that S_i is nonempty and finite for all $i \in \{1, \dots, k\}$, and that $S = \bigcup_{i=1}^k S_i$ is in general position (that is, no three points of S are collinear).

*A preliminary version was presented at the 18th Spanish Meeting on Computational Geometry (2019).

[†]Departamento de Matemáticas, Facultad de Ciencias, Universidad Nacional Autónoma de México, Mexico. Email: dflorespenaloza@ciencias.unam.mx

[‡]Ibaraki University, Hitachi, Ibaraki, Japan. Email: mikio.kano.math@vc.ibaraki.ac.jp

[§]Institut de Mathématiques de Jussieu - Paris Rive Gauche (UMR 7586), Sorbonne Université, Paris, France and Departamento de Matemáticas, Facultad de Ciencias, Universidad Nacional Autónoma de México, Mexico. Email: leontz@im.unam.mx

[¶]Departamento de Física y Matemáticas, Universidad de Alcalá, Alcalá de Henares, Spain. Email: david.orden@uah.es

^{||}Departamento de Métodos Estadísticos, IUMA, Universidad de Zaragoza, Zaragoza, Spain. Email: jtejel@unizar.es

^{**}Department of Mathematics, California State University Northridge, Los Angeles, CA; and Department of Computer Science, Tufts University, Medford, MA, USA. Email: csaba.toth@csun.edu

^{††}Instituto de Matemáticas, Universidad Nacional Autónoma de México, Mexico. Email: urrutia@matem.unam.mx

^{‡‡}Institute of Software Technology, Graz University of Technology, Graz, Austria. Email: bvogt@ist.tugraz.at

For a simple polygon P with m vertices (or a simple m -gon) and a point x in the plane, we say that P *contains* x if x lies in the interior or on the boundary of P . Given a k -colored point set $S = \bigcup_{i=1}^k S_i$ and a simple polygon P in the plane, we call P a *rainbow polygon* for S if P contains at most one point of each color; and P will be called a *perfect rainbow polygon* if it contains exactly one point of each color. The *perfect rainbow polygon problem* for a colored point set S is that of finding a perfect rainbow polygon with the minimum number of vertices.

One can easily check that a perfect rainbow polygon always exists for a colored point set. A way of constructing such a polygon is described below, using the following well-known property for a plane tree: From a tree T embedded in the plane with straight-line edges, a simple polygon can be built by traversing the boundary of the unbounded face of T , placing a copy of a vertex infinitesimally close to that vertex each time it is visited, and connecting the copies according to the traversal order. One can imagine this simple polygon as the “thickening” of the tree. Thus, for a colored point set S , to build a perfect rainbow polygon we can choose one point of each color, form a star connecting one of these points to the rest, and thicken the star; see Figure 1a. Note that the simple polygon obtained in this way can be as close to the star as we wish, so that it contains no other points in S , apart from the points in S that we have chosen.

However, finding a perfect rainbow polygon of minimum size for a given colored point set (where the size of a polygon is the number of its vertices) is in general much more difficult. We believe that this problem is NP-hard. Therefore, we focus on giving combinatorial bounds for the size of minimum perfect rainbow polygons. Let $\text{rb-index}(S)$ denote the *rainbow index* of a colored point set S ; that is, the smallest size of a perfect rainbow polygon for S . We then define the *rainbow index* of k , denoted by $\text{rb-index}(k)$, to be the largest rainbow index among all the k -colored point sets S ; that is,

$$\text{rb-index}(k) = \max \{ \text{rb-index}(S) : S \text{ is a } k\text{-colored point set} \}. \quad (1)$$

In other words, $\text{rb-index}(k)$ is the smallest integer such that, for every k -colored point set S , there exists a perfect rainbow polygon of size at most $\text{rb-index}(k)$.

The two main results in this paper are the following. First, we determine the values of the rainbow index up to $k = 7$, which is the first case where $\text{rb-index}(k) \neq k$, namely

k	3	4	5	6	7
$\text{rb-index}(k)$	3	4	5	6	8

Second, we prove the following lower and upper bounds for the rainbow index

$$\frac{40 \lfloor (k-1)/2 \rfloor - 8}{19} \leq \text{rb-index}(k) \leq 10 \left\lfloor \frac{k}{7} \right\rfloor + 11.$$

Furthermore, for a k -colored set of n points in the plane, a perfect rainbow polygon of size meeting these upper bounds can be computed in $O(n \log n)$ time.

The rainbow index for small values of k is analyzed in Section 3. In Sections 4 and 5, we provide our upper and lower bounds for the rainbow index, respectively. These bounds are based on the analysis of the complexity of noncrossing covering trees for sets of points, under a new measure defined in this paper. This measure and the relationship between perfect rainbow polygons and noncrossing covering trees is given in Section 2.

Related previous work

Starting from the celebrated Ham-Sandwich theorem, a considerable amount of research about discrete geometry on colored point sets (or mass distributions) has been done. For instance,



Figure 1: (a) Thickening a tree to obtain a perfect rainbow polygon. Different colors are represented by different geometric objects. (b) A noncrossing covering tree for the eight black points that can be partitioned into five segments, $s_1 = u_1u_8$, $s_2 = u_2u_3$, $s_3 = u_2u_4$, $s_4 = u_2u_5$, and $s_5 = u_6u_7$; and two forks, u_6 with multiplicity 1 and u_2 with multiplicity 2.

given cg red points and dg blue points in the plane, where c , d , and g are positive integers, the Equitable Subdivision Theorem establishes that there exists a subdivision of the plane into g convex regions such that each region contains precisely c red points and d blue points [10, 34]. It is also known that every d -colored set of points in general position in \mathbb{R}^d can be partitioned into n subsets with disjoint convex hulls such that the set of points and all color classes are partitioned as evenly as possible [12]. For a wide range of geometric partitioning results, the reader is referred to [6, 7, 8, 10, 12, 21, 24, 25, 31, 34] and the references therein.

In addition to geometric partitions, for colored points in the plane some research focuses on geometric structures covering the points in some specific way. For instance, covering the colored points with noncrossing monochromatic matchings [18], noncrossing heterochromatic matchings [26], noncrossing alternating cycles and paths [29], noncrossing alternating spanning trees [11, 26] or noncrossing $K_{1,3}$ stars [1]. In other papers, the main goal is selecting k points with k distinct colors (a rainbow subset) from a k -colored point set such that some geometric properties of the rainbow subset are maximized or minimized. Rainbow subsets with maximum diameter are investigated in [19, 23], with minimum diameter in [20, 33], and rainbow subsets optimizing matchings under several criteria are studied in [9]. In addition, several traditional geometric problems for uncolored point sets become NP-hard for colored point sets. For instance, the following problems are NP-complete [23]: Computing a rainbow subset minimizing (maximizing) the length of its minimum spanning tree, computing a rainbow subset minimizing its convex hull, or computing a rainbow subset maximizing the distance between its closest pair.

Given a 3-colored point set $R \cup B \cup G$ consisting of red, blue, and green points in the plane, a well-known result is that there exists an empty heterochromatic triangle, where the three vertices have distinct colors [15]. In particular, a heterochromatic triangle of minimum area cannot contain any other point from $R \cup B \cup G$ in its interior, hence its interior is empty, and its boundary contains exactly one point of each color. This implies that $\text{rb-index}(3) = 3$. Related work [7] deals with colored lines instead of colored points, showing that in an arrangement of 3-colored lines, there always exists a line segment intersecting exactly one line of each color. Aloupis et al. [4] study the problem of coloring a given point set with k colors so that every axis-aligned strip containing sufficiently many points contains a point from each color class.

2 Covering trees versus perfect rainbow polygons

Given a set of (monochromatic) points, in this section we derive a lower bound for the size of simple polygons that contain the given points and have arbitrarily small area. We also provide a lower bound for the size of a perfect rainbow polygon for some colored point sets.

A noncrossing *covering tree* for a set S of points in the plane is a noncrossing geometric tree (that is, a plane straight-line tree) such that every point of S lies at a vertex or on an edge of the tree; see Figure 1b. Let T be a noncrossing covering tree whose vertices can be collinear. Similarly to [16], we define a *segment* of T as a path of collinear edges in T . Two segments of T may cross at a vertex of degree 4 or higher; we are interested in pairwise noncrossing segments. Any vertex of degree two and incident to two collinear edges can be suppressed; consequently, we may assume that T has no such vertices.

Let \mathcal{M} be a partition of the edges of T into the minimum number of pairwise noncrossing segments. Let $s = s(T)$ denote the number of segments in \mathcal{M} . A *fork* of T (with respect to \mathcal{M}) is a vertex v that lies in the interior of a segment $ab \in \mathcal{M}$ and is an endpoint of another segment in \mathcal{M} . The *multiplicity* of a fork v is 2 if it is the endpoint of two segments that lie on opposite sides of the supporting line of ab ; otherwise its multiplicity is 1. See Figure 1b for an example.

Let $t = t(T)$ denote the sum of multiplicities of all forks in T with respect to \mathcal{M} . We express the number of vertices in a polygon that encloses a noncrossing covering tree T in terms of the parameters s and t . If all edges of T are collinear, then $s = 1$ and T can be enclosed in a triangle. The following lemma addresses the case that $s \geq 2$.

Lemma 1. *Let T be a noncrossing covering tree and \mathcal{M} a partition of the edges into the minimum number of pairwise noncrossing segments. If $s \geq 2$ and $t \geq 0$, then for every $\varepsilon > 0$, there exists a simple polygon P with $2s + t$ vertices such that $\text{area}(P) \leq \varepsilon$ and T lies in P .*

Proof. Let $\delta > 0$ be a sufficiently small constant specified below. For every vertex v of T , let D_v be a disk of radius δ centered at v . We may assume that $\delta > 0$ is so small that the disks D_v , $v \in V(T)$, are pairwise disjoint, and each D_v intersects only the edges of T incident to v . Then the edges of T incident to v partition D_v into $\deg(v)$ sectors. If $\deg(v) \geq 3$, at most one of the sectors subtends a flat angle (that is, an angle equal to π). If $\deg(v) \leq 2$, none of the sectors subtends a flat angle by assumption. Conversely, if one of the sectors subtends a flat angle, then the two incident edges are collinear; they are part of the same segment (by the minimality of \mathcal{M}), and hence v is a fork of multiplicity 1.

In every sector that does not subtend a flat angle, choose a point in D_v on the angle bisector. By connecting these points in counterclockwise order along T , we obtain a simple polygon P that contains T . Note that P lies in the δ -neighborhood of T , so $\text{area}(P)$ is less than the area of the δ -neighborhood of T . The δ -neighborhood of a line segment of length ℓ has area $2\ell\delta + \pi\delta^2$. The δ -neighborhood of T is the union of the δ -neighborhoods of its segments. Consequently, if L is the sum of the lengths of all segments in \mathcal{M} , then the area of the δ -neighborhood of T is bounded from above by $2L\delta + s\pi\delta^2$, which is less than ε if $\delta > 0$ is sufficiently small.

It remains to show that P has $2s + t$ vertices; that is, the total number of sectors whose angle is not flat is precisely $2s + t$. We define a perfect matching between the vertices of P and the set of segment endpoints and forks (with multiplicity) in each disk D_v independently for every vertex v of T . If v is not a fork, then D_v contains $\deg(v)$ vertices of P and $\deg(v)$ segment endpoints. If v is a fork of multiplicity 1, then D_v contains $\deg(v) - 1$ vertices of P and $\deg(v) - 2$ segment endpoints. Finally, if v is a fork of multiplicity 2, then D_v contains $\deg(v)$ vertices of P and $\deg(v) - 2$ segment endpoints. In all cases, there is a one-to-one correspondence between the vertices in P lying in D_v and the segment endpoints and forks (with multiplicity) in D_v . Consequently, the number of vertices in P equals the sum of the multiplicities of all forks plus the number of segment endpoints, which is $2s + t$, as required. \square

Next, we establish a relation between covering trees and points sets, assuming a stronger notion of general position. A point set is in *strong general position* if there is no nontrivial algebraic relation between the coordinates of the points.

Lemma 2. *Let S be a finite set of points in the plane in strong general position and let B be an axis-aligned bounding box of S . Then there exists an $\varepsilon > 0$ such that **for every** simple polygon P with m vertices **such that** $S \subset P$ and $\text{area}(P \cap B) \leq \varepsilon$, **there is** a noncrossing covering tree T of S and a partition of the edges into pairwise noncrossing segments such that $2s + t \leq m$.*

Proof. Let $m \geq 3$ be an integer such that for every $n \in \mathbb{N}$, there exists a simple polygon P_n with precisely m vertices such that $S \subset \text{int}(P_n)$ and $\text{area}(P_n \cap B) \leq \frac{1}{n}$. Let $\overline{\mathbb{R}^2}$ be the compactification of the Euclidean plane \mathbb{R}^2 by adding a circle at infinity, corresponding to the unit (direction) vectors in \mathbb{S}^1 . By compactness, the sequence $(P_n)_{n \geq 1}$ contains a convergent subsequence of polygons in $\overline{\mathbb{R}^2}$. The limit is a (possibly degenerate) polygon \overline{P} with precisely m vertices (some of which may coincide) such that $S \subset \overline{P}$ and $\text{area}(\overline{P} \cap B) = 0$; see Fig. 2(a). Some vertices of \overline{P} may be on the circle at infinity, and so every edge is a single point, a line segment, a ray, or an arc of the circle at infinity. Since S is in strong general position, we may assume that if two rays are parallel, then one contains the other. Hence any edge of \overline{P} along the circle at infinity is an arc with central angle less than π .

Two edges of \overline{P} *overlap* if they intersect in a line segment, a ray, or a circular arc at infinity; and they *cross* if they intersect in a single point that lies in the relative interior of both edges. Since P_n is a simple polygon for all $n \geq 1$, the edges of \overline{P} are noncrossing, although they may overlap. Since $\text{area}(\overline{P} \cap B) = 0$, the union of the edges of \overline{P} is a connected set that contains S . In particular, the union of the m edges of \overline{P} contains a noncrossing covering tree for S in $\overline{\mathbb{R}^2}$.

Let D be a disk whose interior contains S and all vertices of P that are not at infinity. We consider the polygon \overline{P} as a closed curve in $\overline{\mathbb{R}^2}$. Let \mathcal{P} be the collection of maximal **curves** of \overline{P} in the exterior of D . For any two **curves** in \mathcal{P} , the cyclic order of their endpoints along ∂D cannot interleave (i.e., they cannot form an *abab* pattern), since P_n is a simple polygon for all $n \geq 1$. Each **curve** $\gamma \in \mathcal{P}$ contains a unique maximal **curve** along the circle at infinity; denote this **curve** by $\hat{\gamma}$. The **curve** $\hat{\gamma}$ may consist of a single vertex of \overline{P} , or a single edge of \overline{P} (of central angle less than π), or two or more consecutive edges of \overline{P} . Every **curve** $\hat{\gamma}$ is homotopic, with endpoints fixed, to an arc of the circle at infinity (of central angle up to 2π) between the same two vertices of \overline{P} ; let \mathcal{Q} be the set of these arcs. If two arcs in \mathcal{Q} overlap, then one contains the other, otherwise the endpoints of the corresponding **curves** of \overline{P} would interleave in ∂D .

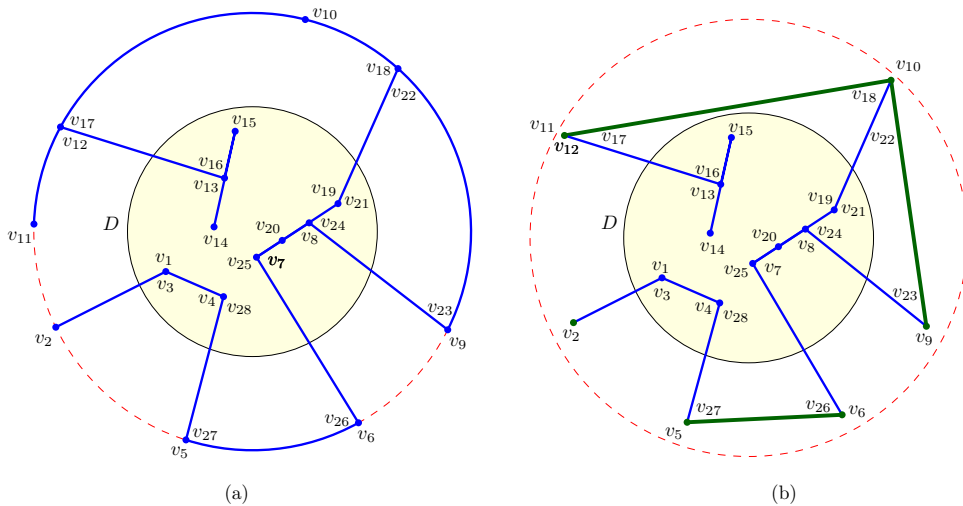


Figure 2: (a) A polygon \overline{P} with 28 vertices and zero area in $\overline{\mathbb{R}^2}$. The dashed circle indicates the circle at infinity. (b) A modified polygon P , where all vertices are in the plane.

We modify \overline{P} to obtain a polygon P in the plane with m vertices and pairwise noncrossing edges such that the union of its edges forms a covering tree for S ; see Fig. 2(b). The transitive

closure of the overlap relation between circular arcs in \mathcal{Q} is an equivalence relation. For each equivalence class of central angle less than π , we replace the arcs with a line segment in \mathbb{R}^2 that intersects the same set of rays. For an equivalence class of central angle at least π , we create two line segments that meet on one of the rays in the arc.

The transitive closure of the overlap relation between the edges of P is an equivalence relation. The union of each equivalence class is a line segment; we call them *segments* for brevity. These segments are pairwise noncrossing (since the edges of P are pairwise noncrossing), and yield a covering of S with a set \mathcal{M} of pairwise nonoverlapping and noncrossing segments.

Analogously to the proof of Lemma 1, at each vertex v of T , there is a one-to-one correspondence between the vertices in P located at v and the segment endpoints and forks (with multiplicity) located at v . This implies $2s + t = m$ with respect to \mathcal{M} . \square

An immediate consequence of Lemma 2 is a lower bound on the size of simple polygons with arbitrarily small area that enclose a point set S .

Corollary 3. *Let S be a finite set of points in the plane in strong general position with an axis-aligned bounding box B , and let T' be a noncrossing covering tree for S minimizing $2s' + t' = m'$. Then there exists an $\varepsilon > 0$ such that if S is contained in a simple polygon P with m vertices and $\text{area}(P \cap B) \leq \varepsilon$, then $m' \leq m$.*

Proof. By Lemma 2, there exists an $\varepsilon > 0$ such that if a simple polygon P with m vertices and $\text{area}(P \cap B) \leq \varepsilon$ contains S , then P also contains a noncrossing covering tree for S . Therefore, by the minimality of T' , necessarily $m' \leq m$. \square

A similar lower bound can be established for perfect rainbow polygons. In particular, for every set S of k points in the plane in strong general position one can build a $(k + 1)$ -colored point set \widehat{S} , such that finding a noncrossing covering tree for S minimizing $2s + t$ is equivalent to finding a minimum perfect rainbow polygon for \widehat{S} .

Theorem 4. *Let S be a set of k points in the plane in strong general position, and let T' be a noncrossing covering tree for S minimizing $2s' + t' = m'$. Then there exists a $(k + 1)$ -colored point set \widehat{S} in strong general position such that every perfect rainbow polygon for \widehat{S} has at least m' vertices.*

Proof. Note that $m' \leq 2k - 2$ since a star centered at one of the points of S is a covering tree for S with $k - 1$ segments and no forks. Let B be an axis-aligned bounding box of S . We may assume, by applying a suitable affine transformation, that B is a unit square. By Lemma 2, there exists an $\varepsilon > 0$ such that if S is contained in a simple polygon P with m vertices and $\text{area}(P \cap B) \leq \varepsilon$, then S admits a noncrossing covering tree and a partition of its edges into segments such that $2s + t \leq m$.

We construct a $(k + 1)$ -colored point set $\widehat{S} = S \cup S_{k+1}$, where S_{k+1} is a point set such that every triangle Δ with $\text{area}(\Delta \cap B) \geq \varepsilon/(2k)$ contains at least two points in S_{k+1} . Each point of S has a unique color and all points in S_{k+1} have the same color. Let $S_{k+1} = (\frac{\varepsilon}{16k} \cdot \mathbb{Z}^2) \cap B$, that is, a section of an integer grid in B . For a triangle Δ with $\text{area}(\Delta \cap B) \geq \varepsilon/(2k)$, the intersection $\Delta \cap B$ is a convex polygon, hence the boundary of $\Delta \cap B$ intersects the interior of at most $4 \cdot (16k/\varepsilon)$ grid cells, and the area of the grid polygon $\text{conv}(\Delta \cap S_{k+1})$ is at least $\varepsilon/(2k) - 4\varepsilon/(16k) = \varepsilon/(4k)$. By Pick's theorem, at least two points of the grid S_{k+1} lie in the interior of $\text{conv}(\Delta \cap S_{k+1})$, hence in the interior of Δ . A random perturbation of S_{k+1} maintains these properties, and the resulting $(k + 1)$ -colored point set \widehat{S} is in strong general position. (This construction yields $|S_{k+1}| = \Theta(k^2(\text{area}(B))/\varepsilon^2)$. A substantially smaller point set S_{k+1} can be constructed with the same properties using classical discrepancy theory [30].)

Now suppose, for the sake of contradiction, that there exists a perfect rainbow polygon P for \widehat{S} with x vertices where $x < m'$. Triangulate P arbitrarily into $x - 2$ triangles. Since each triangle Δ contains at most one point from S_{k+1} , we have $\text{area}(\Delta \cap B) \leq \varepsilon/(2k)$. Summation over triangles yields $\text{area}(P \cap B) \leq (x - 2)\varepsilon/2k \leq \varepsilon$. By the choice of ε , S admits a noncrossing covering tree and a partition of its edges into segments such that $2s + t \leq x$. This contradicts the minimality of T' , which completes the proof. \square

We conjecture that both problems, finding a noncrossing covering tree minimizing $2s + t$ for a given point set and finding a minimum perfect rainbow polygon for a given colored point set, are NP-hard. Many geometric variants of the classical set cover problem are known to be NP-hard. For example covering a finite set of points by the minimum number of lines is APX-hard [13, 28, 32], see also [17, 27]. The minimum-link covering problem (finding a covering path for a set of points with the smallest number of segments) is NP-complete [5]. However, in these problems, the covering objects (lines or edges) may cross. There are few results on covering points with noncrossing segments. It is known, for example, that it is NP-hard to find a maximum noncrossing matching in certain geometric graph [3]. The problem of, given an even number of points, finding a noncrossing matching that minimizes the length of the longest edge is also known to be NP-hard [2].

3 Rainbow indexes of $k = 3, 4, 5, 6, 7$

This section is devoted to determining the rainbow indexes $\text{rb-index}(k)$ up to $k = 7$. The following theorem is the main result of this section, and it summarizes the results proven below.

Theorem 5. *The rainbow indexes of $k = 3, 4, 5, 6, 7$ are the following: $\text{rb-index}(3) = 3$, $\text{rb-index}(4) = 4$, $\text{rb-index}(5) = 5$, $\text{rb-index}(6) = 6$, and $\text{rb-index}(7) = 8$. Furthermore, for every k -colored set of n points, where $3 \leq k \leq 7$, a perfect rainbow polygon of size at most $\text{rb-index}(k)$ can be found in $O(n \log n)$ time.*




Figure 3: Illustrating the proof of Lemma 6. The points of S_1 and S_3 are drawn as red squares and blue circles, respectively.

Our proof for Theorem 5 relies on the following lemma (Lemma 6), which may be of independent interest. Lemma 6 guarantees the existence of a strip containing at least one point of each color, with the additional property that there are at least two color classes that have only one point in the strip.

Before proving the lemma, we introduce some notation. The line segment connecting two points x and y in the plane will be denoted by xy (or yx). Further, a ray emanating from x and passing through y is denoted by \overrightarrow{xy} . Given two parallel lines ℓ_1 and ℓ_2 , defining a strip ST , we

denote by \overline{ST} the closure of the strip; that is, the set of points in the interior of the strip or on the lines ℓ_1 and ℓ_2 .

Lemma 6. *For every k -colored point set $S = \bigcup_{i=1}^k S_i$, where $k \geq 3$, there exist three different colors, i_1 , i_2 , and i_3 , and two parallel lines ℓ_1 and ℓ_2 defining a strip ST , that satisfy the following properties:*

- i) \overline{ST} contains at least one point of S_i for $i = 1, \dots, k$.
- ii) ℓ_1 passes through a point $x \in S_{i_1}$ and a point $y \in S_{i_2}$, such that x and y are the only points with colors i_1 and i_2 in \overline{ST} , respectively.
- iii) ℓ_2 passes through a point $z \in S_{i_3}$.
- iv) If ℓ_2 passes through no other point in S , then z is the only point of color i_3 in \overline{ST} .
- v) If ℓ_2 passes through another point $w \in S$, then either $w \in S_{i_3}$ and z, w are the only points of color i_3 in \overline{ST} , or $w \in S - (S_{i_1} \cup S_{i_2} \cup S_{i_3})$ and z is the only point of color i_3 in \overline{ST} .

Such a strip ST can be computed in $O(n \log n)$ time, where $n = |S|$.

Proof. For $i = 1, \dots, k$, let L_i be the horizontal line passing through the lowest point in S_i . Without loss of generality, assume that $L = L_1$ is the highest line among L_i , $i \in \{1, \dots, k\}$. For $i \in \{2, \dots, k\}$ let S_i^- denote the set of points in S_i that lie strictly below L ; and let U_i be the horizontal line that passes through the highest point in S_i^- . Without loss of generality, assume that $U = U_3$ is the lowest line among all U_i , $i \in \{2, \dots, k\}$ and that $L \cap S_1$ is to the right of $U \cap S_3$; see Figure 3a. Observe that by choosing $i_1 = 1$, $i_3 = 3$, $\ell_1 = L$ and $\ell_2 = U$, all conditions (i) and (iii)–(v) are satisfied.

Let $A = \text{conv}(S_1)$ and $B = \text{conv}(S_3^-)$. We describe a sweepline algorithm in which we maintain a strip \overline{ST} between two parallel lines ℓ_1 and ℓ_2 . Initially, $\ell_1 = L$ and $\ell_2 = U$ are horizontal lines. We also maintain the invariants that

- (I1) ℓ_1 is tangent to A ,
- (I2) ℓ_2 is a tangent to B ,
- (I3) $|\overline{ST} \cap S_i| \geq 1$ for all $\{1, \dots, k\}$, and
- (I4) $ST \cap S_3 = \emptyset$.

During the algorithm ℓ_1 rotates clockwise about the point in $\ell_1 \cap S_1$, and ℓ_2 rotates clockwise about the point $\ell_2 \cap S_3^-$, which are called the *pivot* points of ℓ_1 and ℓ_2 , respectively. Using a fully dynamic convex hull data structure [22], we maintain the convex hull of the points of $S \setminus (A \cup B)$ in ST , above ℓ_1 , and below ℓ_2 , respectively. By computing tangent lines from the two pivot points to the three convex hulls, we can maintain an event queue of when the next point in $S \setminus (A \cup B)$ enters or exits the strip ST (it is deleted from one convex hull and inserted into another) and when the pivot $\ell_1 \cap S_1$ or $\ell_2 \cap S_3^-$ must be updated. We also maintain the number of points in $ST \cap S_i$ for $i = \{2\} \cup \{4, \dots, k\}$, in order to keep track of whether invariant (I3) is satisfied.

Specifically, the rotation stops if the number of points in ST of a color in $\{2\} \cup \{4, \dots, k\}$ drops to zero, or if the line ℓ_1 passes through a point of color 3. Termination is guaranteed, since when $\ell_1 = \ell_2$ is the common tangent of A and B , then $ST = \emptyset$, violating invariant (I3). Each of the lines ℓ_1 and ℓ_2 sweeps through every point in $S \setminus (A \cup B)$ at most once. Since the dynamic convex hull data structures can be updated in $O(\log n)$ amortized time and they support tangent queries in $O(\log n)$ time [22], the sweepline algorithm runs in $O(n \log n)$ total time.

Note that (I1) and (I2) are always satisfied by construction. When the sweep line algorithm terminates, we face two possible scenarios. If (I3) is violated, then there is a color $i_2 \in \{2\} \cup \{4, \dots, k\}$ such that S_{i_2} has exactly one point in $\overline{ST} \cap S_{i_2}$ lying on ℓ_1 or ℓ_2 . In this case our proof is complete with condition (ii) satisfied. Otherwise (I4) is violated; that is, ℓ_1 passes through a point of color 3. In this case, we can translate ℓ_2 towards ℓ_1 until it passes through the last point u in ST of one of the colors $i_2 \in \{2\} \cup \{4, \dots, k\}$; see Figure 3c. Exchanging the roles of colors i_2 and $i_3 = 3$, the lemma follows. \square

Notice that if $k = 3$, then the strip defined by ℓ_1 and ℓ_2 in Lemma 6 is empty, so the triangle $\triangle xyz$ is empty. As a consequence, Lemma 6 provides an alternative proof for $\text{rb-index}(3) = 3$.

In the remainder of this section, we refer to colors 1, 2, 3, 4, 5, 6, and 7 (if they exist) as *red*, *blue*, *green*, *yellow*, *pink*, *orange*, and *black*, respectively (e.g., a 4-colored point set will be red, blue, green, and yellow). Furthermore, when applying Lemma 6, we may assume without loss of generality that the colors i_1 , i_2 , and i_3 are red, blue, and green, respectively, the lines ℓ_1 and ℓ_2 are horizontal, the point x is to the left of point y on ℓ_1 , and if ℓ_2 passes through another point w of S that is not green, then w is yellow. In addition, if p is the intersection point between a ray $\overrightarrow{z\hat{u}}$ and a line ℓ , then p' will denote a point infinitesimally close to p on the ray $\overrightarrow{z\hat{u}}$ towards z ; see Figure 4c.

We can now show that $\text{rb-index}(4) = 4$.

Proposition 7. $\text{rb-index}(4) = 4$.

Proof. We first show that $\text{rb-index}(4) \geq 4$. Consider the 4-colored point set in Figure 4a, where $S_1 = \{x\}$, $S_2 = \{y\}$, $S_3 = \{z\}$, and S_4 consists of two points in the interior of the triangle $\triangle xyz$. Every triangle that contains a point of color 1, 2, and 3 must contain $\triangle xyz$, hence two points of S_4 . It follows that there exists no perfect rainbow triangle.




Figure 4: Illustrating the proof of Proposition 7.

We now show that $\text{rb-index}(4) \leq 4$. Let $S = S_1 \cup S_2 \cup S_3 \cup S_4$ be a point set in the plane whose points are colored red, blue, green, and yellow. By Lemma 6, there is a strip defined by two horizontal lines, ℓ_1 and ℓ_2 , where ℓ_1 passes through a red point x and a blue point y , and ℓ_2 passes through a green point z , such that either there are only yellow points in the interior of the strip, or the strip is empty and ℓ_2 passes through a yellow point w . In the first case, we rotate the horizontal ray emanating from z clockwise until it encounters a yellow point u in the interior of the strip; see Figure 4b. Let p be the intersection point of $\overrightarrow{z\hat{u}}$ and ℓ_1 . By symmetry, we may assume that p is to the left of x or on the segment xy . If p is to the left of x , then $py \cup pz$ is a covering tree T for $\{x, y, z, u\}$, which can be thickened to a perfect rainbow quadrilateral by Lemma 1; see Figure 4b. If p is on xy , then $yxzp'$ is a perfect rainbow quadrilateral; see

Figure 4c. Finally, if the strip is empty and ℓ_2 contains a yellow point w , then $xyzw$ is a perfect rainbow quadrilateral. \square

Before moving to the next proposition, let us prove the following useful lemma that in fact works for monochromatic points.

Lemma 8. *Let $P = \{u, v, w\}$ be a set of three points in the interior of a triangle Δxyz such that $P \cup \{x, y, z\}$ is in general position. Then one can label the points in P by a, b , and c so that a lies on the line segment xr , for $r = \vec{x\bar{a}} \cap \vec{y\bar{b}}$, b lies on the line segment ys , for $s = \vec{y\bar{b}} \cap \vec{z\bar{c}}$, and c lies on the line segment zt , for $t = \vec{z\bar{c}} \cap \vec{x\bar{a}}$.*

Proof. We first show that we can label the three points in P by a, b , and c so that $\Delta xy a$, $\Delta yz b$, and $\Delta zx c$ are interior-disjoint. By symmetry, we may assume that the line passing through u and w intersects edges xy and xz of Δxyz , and u is closer to xy than w . Let $p = \vec{y\bar{u}} \cap \vec{z\bar{w}}$. Since $P \cup \{x, y, z\}$ is in general position, v lies neither on $\vec{y\bar{u}}$ nor on $\vec{z\bar{w}}$; see Figure 5. Three cases arise: If v is in Δyzp , then the triangles Δxyu , Δyzv , and Δzxw are interior-disjoint; see Figure 5a. If v is in Δxyp , then Δxyv , Δyzu , and Δzxw are interior-disjoint; see Figure 5b. Finally, if v is in Δxzp , then we can take the triangles Δxyu , Δyzw , and Δzxv ; see Figure 5c. By labeling u, v, w as a, b, c in the first case, as b, a, c in the second case, and as a, c, b in the third case, this part of the proof follows.

We now show that a lies on the line segment xr , for $r = \vec{x\bar{a}} \cap \vec{y\bar{b}}$, b lies on the line segment ys , for $s = \vec{y\bar{b}} \cap \vec{z\bar{c}}$, and c lies on the line segment zt , for $t = \vec{z\bar{c}} \cap \vec{x\bar{a}}$. If a were not on the line segment xr , then a would lie below $\vec{y\bar{b}}$, and so $\Delta xy a$ and $\Delta yz b$ would intersect, a contradiction; see Figure 5d. The other cases, b lying on ys and c lying on zt , hold by symmetry. \square




Figure 5: Illustrating the proof of Lemma 8.




Figure 6: (a)–(c) Building a new expedient triangle $\Delta r's't'$ from the expedient triangle Δrst , when replacing c by d . (d) A perfect rainbow hexagon for six of the colors.

We say that the triangle $\triangle rst$ described in Lemma 8 is *expedient* with respect to $\{a, b, c\}$. Note that an expedient triangle can be computed in $O(1)$ time. Further, note that the labeling given in Lemma 8 is not unique, and thus expedient triangles are not uniquely determined. Using Lemma 8, one can find perfect rainbow hexagons for six colors in some special colored point sets, as the following lemma shows.

Lemma 9. *Let S be a k -colored set of n points, with $k \geq 6$, such that $\text{conv}(S) = \triangle xyz$. Assume that $S_1 = \{x\}$, $S_2 = \{y\}$, and $S_3 = \{z\}$. Then there is a perfect rainbow hexagon for six of the colors, which can be found in $O(n)$ time.*

Proof. Let a, b , and c be points in the interior of $\triangle xyz$ with three different colors i_1, i_2 , and i_3 , respectively, and consider an expedient triangle $\triangle rst$ with respect to $\{a, b, c\}$, which exists by Lemma 8. If there are points of S in the interior of $\triangle rst$, then we choose one of them, say d , we replace c by d in $\{a, b, c\}$, and we take a new expedient triangle $\triangle r's't'$ with respect to the set $\{a, b, d\}$. Without loss of generality, we may assume that if the color of d is some of i_1, i_2 , and i_3 , then c is the point having the same color as d .

We show how to construct an expedient triangle $\triangle r's't'$ such that $\triangle r's't' \subset \triangle rst$. Assume that a lies on xr , for $r = \overrightarrow{xa} \cap \overrightarrow{yb}$, b lies on ys , for $s = \overrightarrow{yb} \cap \overrightarrow{zc}$, and c lies on zt , for $t = \overrightarrow{zc} \cap \overrightarrow{xa}$; cf. Figures 6a–6c. Consider the ray \overrightarrow{zb} and distinguish cases based on whether d is to the left or to the right of \overrightarrow{zb} . If d is to the right of \overrightarrow{zb} (Figure 6a), then $\triangle r's't'$ is an expedient triangle, where a lies on xr , b lies on ys' , and d lies on zt' . Suppose now that d is to the left of \overrightarrow{zb} . Two subcases arise, depending on whether d is to the left or to the right of the ray \overrightarrow{ya} . If d is to the right of \overrightarrow{ya} (Figure 6b), then $\triangle r's't'$ is an expedient triangle, where a lies on xr' , d lies on ys' , and b lies on zt' . Finally, if d is to the left of \overrightarrow{ya} (Figure 6c), then $\triangle r's't'$ is an expedient triangle, where d lies on xr' , a lies on ys' , and b lies on zt' . In all three cases, $\triangle r's't' \subset \triangle rst$, as required.

Since $\triangle r's't' \subset \triangle rst$, and since d is in the interior of $\triangle rst$ but not in the interior of $\triangle r's't'$, it follows that $\triangle r's't'$ contains fewer points of S than $\triangle rst$. Hence we can repeat this procedure until we find an expedient triangle that is empty of points of S . From this expedient triangle, we can obtain a perfect rainbow hexagon for the six colors involved, by slightly moving the vertices of the expedient triangle towards the vertices of $\triangle xyz$ as depicted in Figure 6d. Furthermore, an *empty* expedient triangle can be computed in $O(n)$ time: We can start with an arbitrary expedient triangle $\triangle rst$. For each point $s \in S$, we can test whether $s \in \triangle rst$ and update it to a smaller triangle $\triangle r's't' \subset \triangle rst$ if necessary in $O(1)$ time. Consequently, a perfect rainbow polygon for six of the colors can also be found in $O(n)$ time. \square




Figure 7: 5-, 6-, and 7-colored points sets whose rainbow indices are 5, 6, and 8, respectively.

We are now ready to prove the following result.

Proposition 10. $\text{rb-index}(5) = 5$, $\text{rb-index}(6) = 6$, and $\text{rb-index}(7) = 8$.

Proof. We first show that $\text{rb-index}(5) \geq 5$, $\text{rb-index}(6) \geq 6$, and $\text{rb-index}(7) \geq 8$, by constructing a 5-colored point set $S(5)$, a 6-colored point set $S(6)$ and a 7-colored point set $S(7)$, such that $\text{rb-index}(S(5)) = 5$, $\text{rb-index}(S(6)) = 6$, and $\text{rb-index}(S(7)) = 8$.

The set $S(5)$ consists of four one-element color classes $S_1 = \{x\}$, $S_2 = \{y\}$, $S_3 = \{z\}$, and $S_4 = \{w\}$, where w is in the interior of $\triangle xyz$. The set S_5 of black points contains $\triangle xyz$ in its convex hull, as described in the proof of Theorem 4; that is, every triangle of area ε or more contains at least two black points; see Figure 7a. The set $S(6)$ is obtained from $S(5)$ by adding a one-element color class $S_6 = \{u\}$, where u is in the interior of $\triangle xyz$; see Figure 7b. The set $S(7)$ is based on two triangles, $\triangle xyz$ and $\triangle uvw$, where $\triangle uvw$ lies in the interior of $\triangle xyz$ and its vertices are very close to the midpoints of the edges of $\triangle xyz$; see Figure 7c. The set $S(7)$ then consists of six one-element color classes, $S_1 = \{x\}$, $S_2 = \{y\}$, $S_3 = \{z\}$, $S_4 = \{u\}$, $S_5 = \{v\}$, and $S_6 = \{w\}$, and the dense class S_7 of black points as described in the proof of Theorem 4. In the three sets, $S(5)$, $S(6)$ and $S(7)$, we assume that x , y , z , u , v , and w (if defined) are in strong general position.

It is easy to see that a noncrossing covering tree for $\{x, y, z, w\}$ in $S(5)$, minimizing $2s + t$, requires either two segments and a fork, or at least three segments (and no fork). Hence, by Theorem 4, the size of a minimum perfect polygon for $S(5)$ is at least 5. Figure 7a illustrates a perfect rainbow pentagon based on a covering tree that uses a segment to cover x and z and another segment to cover y and w . Every noncrossing covering tree for $\{x, y, z, w, u\}$ in $S(6)$ requires at least three segments, so the size of any perfect rainbow polygon for $S(6)$ is at least 6 by Theorem 4. Figure 7b shows a perfect rainbow hexagon based on three segments that cover x and u , y and w , and z , respectively.

Finally, consider a noncrossing covering tree for $\{x, y, z, w, u, v\}$ in $S(7)$. It has at least three segments, by the pigeonhole principle, since no three points are collinear. If it has four or more segments, then the size of the corresponding perfect rainbow polygon for $S(7)$ is at least 8. Otherwise it consists of exactly three segments, and then an analysis of the possible choices shows that at least two forks are always required. Therefore, the size of a minimum perfect rainbow polygon for $S(7)$ is at least 8. Figure 7c illustrates the perfect rainbow octagon for $S(7)$ based on the segments that cover $\{x, u\}$, $\{y, v\}$, and $\{z, w\}$, respectively.

We next show that $\text{rb-index}(5) \leq 5$, $\text{rb-index}(6) \leq 6$, and $\text{rb-index}(7) \leq 8$. Let S be a k -colored point set in the plane, with $k \in \{5, 6, 7\}$. By Lemma 6, there is a strip defined by two horizontal lines ℓ_1 and ℓ_2 , with ℓ_1 passing through a red point x and a blue point y , and ℓ_2 passing through a green point z , such that either the strip contains some points of all other colors, or ℓ_2 passes through a yellow point w , the interior of the strip does not contain any other yellow point, but contains at least one point from the remaining $k - 4$ color classes. We analyse these cases in detail as follows.

The strip ST contains points of all $k - 3$ other colors. Consider the horizontal ray emanating from z to the left and rotate it in clockwise direction, sweeping all the colored points in the strip until we find two consecutive points of S , say u and v , with different colors, say yellow and pink; see Figure 8a. Let p and q be the intersection points of $\vec{z}u$ and $\vec{z}v$ with ℓ_1 , respectively. Assume that p is to the left of x . If q is also to the left of x or on the line segment xy , then $ypzq'$ is a perfect rainbow quadrilateral for five of the colors; see Figure 8a. If q is to the right of y , then $\triangle pqz$ is a perfect rainbow triangle for five of the colors; see Figure 8b. If $k = 5$, we are done. If $k = 6$, then we connect an orange point to z and thicken this edge (dotted line segments in Figures 8a and 8b). In this way, we obtain a perfect rainbow hexagon in the first case and a perfect rainbow pentagon in the second case. If $k = 7$, we repeat this process connecting a black point to z , to obtain either a perfect rainbow octagon or a perfect rainbow heptagon.




Figure 8: Illustrating the proof of Proposition 10.

Suppose now that p is to the right of x . Arguing in an analogous way when rotating the horizontal ray emanating from z to the right counterclockwise, if u' and v' are two consecutive points with different colors, then we may assume that the intersection point p_1 between ℓ_1 and $\overrightarrow{zu'}$ is to the left of y . When this happens, p to the right of x and p_1 to the left of y , it is straightforward to see that $\triangle xyz$ must contain at least one point of each color, and that q is on xy . If $k = 5$, then $yxp'zq'$ is a perfect rainbow pentagon; see Figure 8c. If $k = 6$, a perfect rainbow hexagon exists by Lemma 9; see Figure 6d. If $k = 7$, we can build a perfect rainbow hexagon for six of the colors by Lemma 9, and form a perfect rainbow octagon for S by connecting a black point to z and thicken this edge.

The only yellow point w in \overline{ST} is on ℓ_2 . When $k = 5$, the strip contains only pink points. Thus, we rotate the horizontal ray emanating from y to the left counterclockwise until it encounters a pink point u in the strip; see Figure 8d. Let p be the intersection point of \overrightarrow{yu} and ℓ_2 . By symmetry, there are two cases to consider: either p is to the left of z , or on the segment zw . In the first case, $yxpp'w'$ is a perfect rainbow pentagon; see Figure 8d. In the second case $yxzwp'$ is a perfect rainbow pentagon; see Figure 8e.

When $k = 6$ or $k = 7$, the strip contains points of at least two of the colors and we can argue as before, but now rotating clockwise about w instead of about z , to look for the first two consecutive points u and v with different colors, say pink and orange. If the intersection point p between \overrightarrow{wu} and ℓ_1 is to the left of x , then we can build a perfect rainbow quadrilateral or a perfect rainbow triangle for five of the colors, as shown in Figures 8a and 8b. After that, we connect z (and a black point if $k = 7$) to w to form a perfect rainbow hexagon (or a perfect octagon if $k = 7$) for S ; see Figure 8f.

Finally, if p is to the right of x , then there are points of at least two of the colors to the left of

\overrightarrow{xw} . Therefore, when rotating the horizontal ray emanating from x to the right clockwise until finding two consecutive points u and v of different colors, the intersection point between \overrightarrow{xu} and ℓ_2 will be necessarily to the right of w , and we can carry out symmetric constructions. \square

The following corollary is straightforward from the proofs of Propositions 7 and 10.

Corollary 11. *For $k = 3, 4, 5, 6, 7$, a perfect rainbow polygon with at most $\text{rb-index}(k)$ vertices can be found in $O(n \log n)$ time for any k -colored set S of n points.*

Proof. By Lemma 6, the strip ST used in the proofs of Propositions 7 and 10 can be found in $O(n \log n)$ time. In addition, the cyclic order of the points in S around any of x, y, z or w can be computed in $O(n \log n)$ time. A perfect rainbow hexagon as described in Lemma 9 can be obtained in $O(n)$ time. Therefore, the corollary follows. \square

4 Upper bound for rainbow indexes

We show in this section that for every k -colored point set, there exists a perfect rainbow polygon of size at most $10\lfloor \frac{k}{7} \rfloor + 11$. We begin with an auxiliary lemma showing that any seven (monochromatic) points in a vertical strip can be covered by a noncrossing forest of two trees of order four and two, respectively, such that both trees are fully contained in the strip.

Lemma 12. *Let $S = \{p_1, \dots, p_7\}$ be a point set in general position in the plane, ordered by x -coordinate. Let B be the strip defined by the two vertical lines passing through p_1 and p_7 , respectively. Then, in $O(1)$ time, we can construct two noncrossing trees T_1 and T_2 , of order four and two, respectively, with the following properties:*

- (i) *The union of T_1 and T_2 covers S and is contained in B .*
- (ii) *For $i \in \{1, 2\}$, the tree T_i has a leaf v_i such that the ray emanating from v_i in the direction opposite to the edge incident to v_i goes to the left and does not cross T_i . Moreover, if the extension at v_i hits T_j , $j \neq i$, then the extension at v_j does not hit T_i ; that is, the two trees and the two extensions do not create cycles.*

Proof. Let ℓ be the line passing through p_1 and p_7 . Without loss of generality, we may assume that S contains at least $\lceil 5/2 \rceil = 3$ points below ℓ . Note that p_1 and p_7 are extremal points in S , hence they are vertices of the convex hull $\text{conv}(S)$ of S . Points p_1 and p_7 decompose the boundary of $\text{conv}(S)$ into two convex arcs, an *upper arc* and a *lower arc*. Since S contains at least 3 points below ℓ , the lower arc must have at least 3 vertices (including p_1 and p_7). We distinguish between two cases depending on the number of vertices of the lower arc of $\text{conv}(S)$.

The lower arc of $\text{conv}(S)$ has 3 vertices. Assume that the lower arc of $\text{conv}(S)$ is the path (p_1, p_i, p_7) , where $1 < i < 7$; see Figures 9a–9b. Since S contains at least 3 points below ℓ , at least 2 points of S are in the interior of $\triangle p_1 p_i p_7$. Rotate the ray $\overrightarrow{p_1 p_i}$ counterclockwise until it encounters a point p_a in the interior of $\triangle p_1 p_i p_7$; rotate the ray $\overrightarrow{p_7 p_i}$ clockwise until it encounters a point p_b in the interior of $\triangle p_1 p_i p_7$. We distinguish between two cases depending on whether p_a and p_b are distinct.

In the first case, assume that $p_a \neq p_b$; see Figure 9a. The rays $\overrightarrow{p_1 p_a}$ and $\overrightarrow{p_7 p_b}$ intersect in the interior of $\triangle p_1 p_i p_7$, at some point q . By construction, the remaining two points in $S \setminus \{p_1, p_i, p_7, p_a, p_b\}$ are above the path (p_1, q, p_7) , in a wedge bounded by $\overrightarrow{qp_1}$ and $\overrightarrow{qp_7}$. This wedge is convex, hence it contains the line segment between the two points. Let T_1 be the star centered at q with edges $p_1 q$, $p_i q$, and $p_7 q$; and let T_2 be the line segment spanned by the two points of S above (p_1, q, p_7) .




Figure 9: Finding two trees when the lower arc of $\text{conv}(S)$ has three vertices (9a–9b), and when it has four or more vertices (9c).

In the second case, assume that $p_a = p_b$; see Figure 9b. Let r be the first point along the ray $\overrightarrow{p_i p_a}$ such that the line segment $p_1 r$ or $p_7 r$ contains a point of S in the interior of $\triangle p_1 p_a p_7$. Denote this point by $p_c \in S$. By construction, the remaining two points in $S \setminus \{p_1, p_i, p_7, p_a, p_c\}$ are above the path (p_1, r, p_7) , in a convex wedge bounded by $\overrightarrow{r p_1}$ and $\overrightarrow{r p_7}$. Let T_1 be the star centered at r with edges $p_1 r$, $p_i r$, and $p_7 r$, and let T_2 be the line segment spanned by the two points of S above (p_1, r, p_7) .

In both cases, $T_1 \cup T_2$ covers all seven points in S , is in B , and is noncrossing, as required. Moreover, the second property of the lemma is clearly satisfied by choosing $v_1 = p_1$ and v_2 the leftmost point of T_2 .

The lower arc of $\text{conv}(S)$ has 4 or more vertices. Let p_1, p_i, p_j, p_k be the first four vertices of the lower arc of $\text{conv}(S)$ in counterclockwise order (possibly $p_k = p_7$). Let q be the intersection point of the lines passing through p_1 and p_i , and through p_j and p_k , respectively. Rotate the ray $\overrightarrow{p_k p_j}$ clockwise until it encounters a point in $S \setminus \{p_1, p_i, p_j, p_k\}$, and denote it by p_a . Let T_1 be the union of the path (p_1, q, p_k, p_a) , and connect the two remaining points of S to define T_2 ; see Figure 9c. Note that $T_1 \cup T_2$ covers all seven points in S , is in B , and is noncrossing, as required. Moreover, the second property of the lemma is clearly satisfied by choosing $v_1 = p_1$ and v_2 the leftmost point of T_2 . \square

Using Lemma 12, the following theorem provides a method to find noncrossing covering trees with few segments and forks.

Theorem 13. *Let S be a finite set of $n = 7j + r$ points in the plane in general position, with $j \geq 0$ and $0 \leq r \leq 6$. Then, in $O(n \log n)$ time, we can construct a noncrossing covering tree T consisting of $4\lfloor \frac{n}{7} \rfloor + 1 + \lceil \frac{r}{2} \rceil$ segments and $2\lfloor \frac{n}{7} \rfloor + \lceil \frac{r}{2} \rceil$ forks with multiplicity 1.*

Proof. By rotating the point set if necessary, we may assume that the points in S have distinct x -coordinates. We assume first that $n = 7j$, for some integer $j > 0$. Figure 10 illustrates the method to obtain a noncrossing covering tree with $4\lfloor \frac{n}{7} \rfloor + 1$ segments and $2\lfloor \frac{n}{7} \rfloor$ forks with multiplicity 1. We partition the n points from left to right into j groups G_1, G_2, \dots, G_j of seven points each; see Figure 10a. We apply Lemma 12 to every group G_i to cover the points in G_i by two trees, consisting of 4 segments in total; see Figure 10b. In this way, we obtain a forest F formed by $2j$ trees with $4j$ segments. In addition, by the same lemma, every tree T_i of F contains a special leaf v_i that can be extended to the left without crossing T_i .

We add a long vertical segment P' to the left of the point set such that the extension of any tree T_i of F at v_i crosses P' ; see Figure 10c. For every tree T_i , we extend the edge incident to

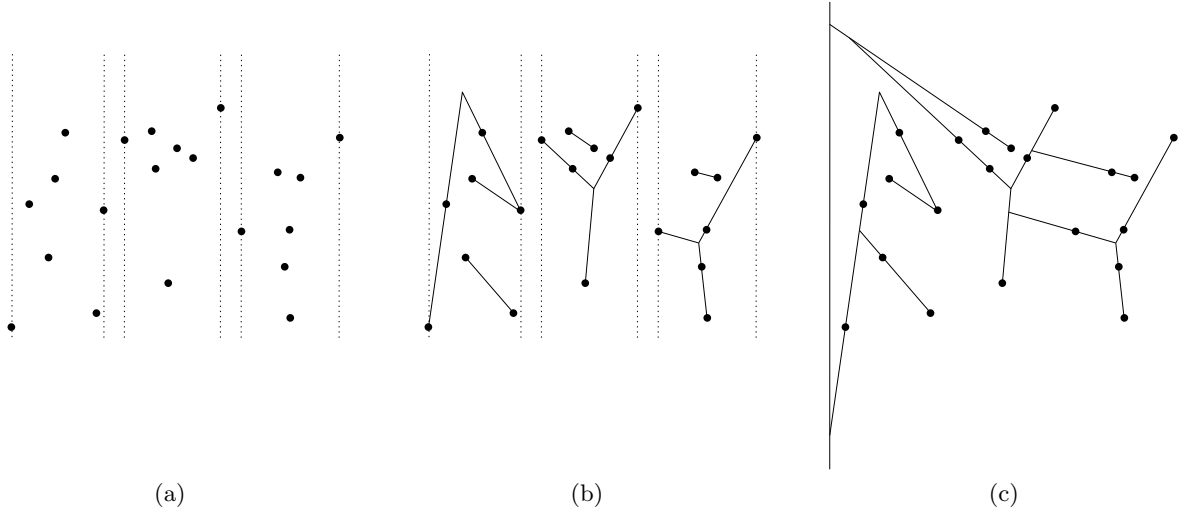


Figure 10: Part (a): Dividing the n points into groups of size 7. Part (b): Applying Lemma 12 to each group. Part (c): Joining all trees to a vertical segment.

its special leaf v_i to the left until the extension hits another tree, another extension, or P' ; see Figure 10c. This is carried out exploring, for example, the special leaves from right to left. Thus, we join the $2j$ trees of F and the segment P' to form a single component. This component is necessarily a noncrossing covering tree T with $4j + 1$ segments and $2j$ forks with multiplicity 1, since all extensions go to the left without creating cycles. Therefore, there exists a noncrossing covering tree T consisting of $4\lfloor \frac{n}{7} \rfloor + 1$ segments and $2\lfloor \frac{n}{7} \rfloor$ forks with multiplicity 1.

Consider the case that $n = 7j + r$, where $1 \leq r \leq 6$. Using the first $7j$ points from left to right, we proceed as before, and we build a noncrossing covering tree T with $4j + 1$ segments and $2j$ forks with multiplicity 1. If $j = 0$, the previous step is not required. The last r points can be covered by connecting the point at position $7j + 1$ to the following one, the point at position $7j + 3$ to the following one, and so on. If the last point cannot be connected to the following one, we assign a small horizontal segment to it. In this way, we are covering the last r points with $\lceil \frac{r}{2} \rceil$ segments. These segments can be joined to T by extending their leftmost points. Therefore, we can obtain a noncrossing covering tree T' consisting of $4\lfloor \frac{n}{7} \rfloor + 1 + \lceil \frac{r}{2} \rceil$ segments and $2\lfloor \frac{n}{7} \rfloor + \lceil \frac{r}{2} \rceil$ forks with multiplicity 1.

It remains to show that the construction above can be implemented in $O(n \log n)$ time. We can sort the points in S in increasing order by x -coordinates in $O(n \log n)$ time, and hence partition S into $O(n)$ groups of size seven. For each group, we can find two trees in $O(1)$ time by Lemma 12. Finally, we can compute the left extensions of the special leaves v_i of all trees T_i by a standard sweepline algorithm [14, Sec. 2.1] as follows: We sweep a vertical line L right to left from the rightmost point to one unit left of the leftmost point in S . In the course of the algorithm, we maintain the intersection of L with the forest F , and the left extensions of all leaves v_i to the right of L . An event queue maintains the time steps when L passes through a vertex of F , when a left extension hits an edge of F (in which case the left extension ends), and when two left extensions meet (in which case one extension ends and other one continues). There are $O(n)$ events, and the event queue can be updated in $O(\log n)$ time for each event. Consequently, the sweepline algorithm runs in $O(n \log n)$ time. In the last step, the sweepline L is to the left of S ; we can let the vertical line segment P' be the convex hull of the intersections of L with all surviving left extensions. \square

Notice that by construction, the minimum number of pairwise noncrossing segments into which T' can be decomposed is precisely $4\lfloor \frac{n}{7} \rfloor + 1 + \lceil \frac{r}{2} \rceil$. As a consequence of this theorem, we can give an upper bound for the size of a perfect rainbow polygon.

Theorem 14. *Let S be a k -colored set of n points in general position. Then a perfect rainbow polygon P of size at most $10\lfloor \frac{k}{7} \rfloor + 11$ can be computed in $O(n \log n)$ time.*

Proof. We choose a point of each color to define a point set S' of cardinality $k = 7j + r$, with $j \geq 0$ and $0 \leq r \leq 6$. By Theorem 13, there is a noncrossing covering tree T' for the point set S' , consisting of $4\lfloor \frac{k}{7} \rfloor + 1 + \lceil \frac{r}{2} \rceil$ segments and $2\lfloor \frac{k}{7} \rfloor + \lceil \frac{r}{2} \rceil$ forks with multiplicity 1, and it can be computed in $O(k \log k)$ time. By Lemma 1, given a noncrossing covering tree T and a partition \mathcal{M} of the edges into the minimum number s of pairwise noncrossing segments, for every $\varepsilon > 0$, there exists a simple polygon P with $2s + t$ vertices such that $\text{area}(P) \leq \varepsilon$ and T lies in P , where t is the sum of the multiplicities of all forks in T . Thus, for every $\varepsilon > 0$, we can construct a simple polygon P' with $2(4\lfloor \frac{k}{7} \rfloor + 1 + \lceil \frac{r}{2} \rceil) + 2\lfloor \frac{k}{7} \rfloor + \lceil \frac{r}{2} \rceil \leq 10\lfloor \frac{k}{7} \rfloor + 11$ vertices such that $\text{area}(P') \leq \varepsilon$ and S' lies in P' . By choosing ε sufficiently small so that P' contains no other point in S except for the points in S' , we can construct a perfect rainbow polygon for S of size at most $10\lfloor \frac{k}{7} \rfloor + 11$.

A suitable $\varepsilon > 0$ can be half of the minimum distance between the covering tree T' and the points in $S \setminus S'$. To find this distance, we can compute the Voronoi diagram for a set of sites, which consists of the $O(k) \leq O(n)$ edges of T' and the $O(n)$ points in $S \setminus S'$ in $O(n \log n)$ time [14, Sec. 7.3]. The Voronoi diagram is formed by $O(n)$ line segments and parabolic arcs; and we can find the closest point in T' (hence in $S \setminus S'$) for each of these arcs in $O(1)$ time. \square

5 Lower bound for rainbow indexes

For every $k \geq 3$, Dumitrescu et al. [16] constructed a set S of $n = 2k$ points in the plane in strong general position (without colors) for which every noncrossing covering path has at least $(5n - 4)/9$ edges. They also showed that every noncrossing covering tree for S has at least $(9n - 4)/17$ edges. Furthermore, every set of $n \geq 5$ points in general position in the plane admits a noncrossing covering tree with at most $\lceil n/2 \rceil$ noncrossing segments, and this bound is the best possible. We recall that a segment is defined as a path of collinear edges.

In this section, we use the point sets constructed in [16] to derive a lower bound for the complexity of a covering tree as defined in Section 2. This bound, in turn, yields a lower bound on the complexity of perfect rainbow polygons for colored point sets built from such sets.




Figure 11: A sketch of the construction given in [16] for $k = 4$. The figure is not to scale.

Construction. We use the point set constructed by Dumitrescu et al. [16]. We review some of its properties here. For every $k \in \mathbb{N}$, they construct a set of $n = 2k$ points, $S = \{a_i, b_i : i = 1, \dots, k\}$. The pairs $\{a_i, b_i\}$ ($i = 1, \dots, k$) are called *twins*. The points a_i ($i = 1, \dots, k$) lie on the parabola $\alpha = \{(x, y) : y = x^2\}$, sorted by increasing x -coordinate. The points b_i ($i = 1, \dots, k$) lie on a convex curve β above α , such that $\text{dist}(a_i, b_i) < \varepsilon$ for a sufficiently small ε , and the lines $a_i b_i$ are almost vertical with monotonically decreasing positive slopes (hence the supporting lines of any two twins intersect below α). For $i = 1, \dots, k$, they also define pairwise disjoint disks $D_i(\varepsilon)$ of radius ε centered at a_i such that $b_i \in D_i(\varepsilon)$, and the supporting lines of segments $a_i a_j$ and $b_i b_j$ meet in $D_i(\varepsilon)$ for every $j, i < j \leq k$. Furthermore, (1) no three points in S are collinear; (2) no two lines determined by the points in S are parallel; and (3) no three lines determined by disjoint pairs of points in S are concurrent. The x -coordinates of a_i ($i = 1, \dots, k$) are chosen such that (4) for any four points c_1, c_2, c_3, c_4 from S , labeled by increasing x -coordinate, the supporting lines of $c_1 c_4$ and $c_2 c_3$ cross to the left of these points. See Figure 11 for a sketch of the construction. Finally, the point set S is perturbed into strong general position maintaining these properties (except that the points a_i ($i = 1, \dots, k$) are no longer on a parabola).

Analysis. Let S be a set of $n = 2k$ points defined in [16] as described above, for some $k > 1$. Let \mathcal{M} be a set of pairwise noncrossing line segments in the plane whose union is connected and contains S . In particular, if T is a noncrossing covering tree for S , then any partition of the edges of T into pairwise noncrossing segments could be taken to be \mathcal{M} .

A segment in \mathcal{M} is called *perfect* if it contains two points in S ; otherwise it is *imperfect*. By perturbing the endpoints of the segments in \mathcal{M} , if necessary, we may assume that every point in S lies in the relative interior of a segment in \mathcal{M} . By the construction of S , no three perfect segments are concurrent, so we can define the set Γ of maximal paths of perfect segments; we call these *perfect paths* or *perfect chains*.

Dumitrescu et al. [16, Lemmata 4–10] proved several properties of a covering *path* for S . Clearly, a covering path has precisely two leaves, while a covering tree may have arbitrarily many leaves. Their results are based on local configurations, however, and hold for any set of noncrossing segments \mathcal{M} where the endpoints of perfect chains play the same role as the endpoints of a covering path. We restate their key results for a set \mathcal{M} of noncrossing covering segments.

Lemma 15. [16, Lemma 7] *Let pq be a perfect segment in \mathcal{M} that contains one point from each of the twins $\{a_i, b_i\}$ and $\{a_j, b_j\}$, where $i < j$. Assume that p is the left endpoint of pq . Let s be the segment in \mathcal{M} containing the other point of the twin $\{a_i, b_i\}$. Then one of the following four cases occurs.*

Case 1: p is the endpoint of a perfect chain;

Case 2: s is imperfect;

Case 3: s is perfect, one of its endpoints v lies in $D_i(\varepsilon)$, and v is the endpoint of a perfect chain;

Case 4: s is perfect and p is the common left endpoint of segments pq and s .

Lemma 16. [16, Lemma 9] *Let pq be a perfect segment in \mathcal{M} that contains a twin $\{a_i, b_i\}$, and let q be the upper (that is, right) endpoint of pq . Then q is the endpoint of a perfect chain.*

Let Γ_x be the set of maximal x -monotone chains of perfect segments in \mathcal{M} .

Lemma 17. *The right endpoints of the chains in Γ_x are distinct.*

Proof. Suppose, for the sake of contradiction, that $\gamma_1, \gamma_2 \in \Gamma_x$ have a common right endpoint q . Let pq and rq , respectively, be the rightmost segments of γ_1 and γ_2 .

If pq contains a twin, then pq has positive slope (by construction), and so q is the upper endpoint of pq . In this case segment rq is imperfect by Lemma 16, contradicting the assumption that rq is in γ_2 . We may assume that neither pq nor rq contains a twin. In this case, their supporting lines intersect to the left of the points in S on pq and pr by property (4), contradicting our assumption that q is the right endpoint of both segments. \square

Corollary 18. *Every chain in Γ consists of at most two chains in Γ_x .*

Denote by s_0, s_1 , and s_2 , respectively, the number of segments in \mathcal{M} that contain 0, 1, and 2 points from S . An adaptation of a charging scheme from [16, Lemma 4] yields the following result, where t is the number of forks (with multiplicity) in \mathcal{M} .

Lemma 19. $s_2 \leq 8s_0 + 9s_1 + 4(t + 1)$.

Proof. Let pq be a perfect segment of \mathcal{M} , and part of a chain $\gamma \in \Gamma$. We charge pq to either an endpoint of γ or some imperfect segment.

We define the charging as follows. If pq contains a twin, then charge pq to the top vertex of pq , which is the endpoint of a perfect chain by Lemma 16. Assume now that pq does not contain a twin, its left endpoint is p , and it contains a point from each of the twins $\{a_i, b_i\}$ and $\{a_j, b_j\}$, with $i < j$. We consider the four cases presented in Lemma 15.

In Case 1, charge pq to p , which is the endpoint of a perfect chain. In Case 2, charge pq to the imperfect segment s containing a point of the twin $\{a_i, b_i\}$. In Case 3, charge pq to the endpoint v of a perfect chain located in $D_i(\varepsilon)$. Now, consider Case 4 of Lemma 15. In this case, pq is the leftmost segment of a maximal x -monotone chain γ_x . We charge pq to the right endpoint of γ_x , which is the endpoint of a perfect chain by Lemma 17. This completes the definition of the charges.

Note that every imperfect segment and every right endpoint of a chain in Γ is charged at most once for perfect segments in Cases 1–3, and every left endpoint of a chain is charged at most twice. By Corollary 18, each endpoint of a perfect chain is charged at most once for perfect segments in Case 4. Overall, every imperfect segment containing one point of S is charged at most once, and every endpoint of a perfect chain is charged at most twice. Consequently,

$$s_2 \leq s_1 + 4|\Gamma|. \quad (2)$$

We bound $|\Gamma|$ from above in terms of s_0, s_1 , and t . Choose an arbitrary root vertex in T , and direct all edges in T towards the root. Every perfect chain has a unique vertex v closest to the root. As all chains in Γ are maximal and as no three perfect segments are concurrent, v must be a fork, the endpoint of an imperfect segment, or the root. There are at most t forks, the $s_0 + s_1$ imperfect segments jointly have at most $2(s_0 + s_1)$ distinct endpoints, and we have one root. This yields $|\Gamma| \leq 2(s_0 + s_1) + t + 1$. Combined with (2), this yields,

$$s_2 \leq s_1 + 4[2(s_0 + s_1) + t + 1] = 8s_0 + 9s_1 + 4(t + 1),$$

as claimed. \square

Lemma 20. *Let S be a set of $n = 2k \geq 4$ points from [16]. Then every covering tree T of S satisfies $2s + t \geq (20n - 8)/19$.*

Proof. The combination of Lemma 19 and $n = s_1 + 2s_2$ yields

$$\begin{aligned}
2s + t &= 2(s_0 + s_1 + s_2) + t \\
&\geq (s_0 + s_1 + t) + (s_1 + 2s_2) \\
&= \frac{2[8s_0 + 9s_1 + 4(t + 1)] + 3s_0 + s_1 + 11t - 8}{19} + n \\
&\geq \frac{2s_2 + s_1 - 8}{19} + n \\
&= \frac{n - 8}{19} + n \\
&= \frac{20n - 8}{19},
\end{aligned}$$

as claimed. \square

We are now ready to prove the main result of this section.

Theorem 21. *For every integer $k \geq 5$, there exists a finite set of k -colored points in the plane such that every perfect rainbow polygon has at least $\frac{40\lfloor(k-1)/2\rfloor-8}{19}$ vertices.*

Proof. Assume first that k is odd, and let S be the set of $k - 1 = 2j \geq 4$ points from [16]. If T is a noncrossing covering tree for S minimizing $2s + t = m$, then by Theorem 4, there exists a k -colored point set \widehat{S} built from S such that every perfect rainbow polygon for \widehat{S} has at least m vertices. By Lemma 20, every noncrossing covering tree of S satisfies $2s + t \geq \frac{20(k-1)-8}{19}$, hence every perfect rainbow polygon for \widehat{S} has at least $\frac{20(k-1)-8}{19} = \frac{40\lfloor(k-1)/2\rfloor-8}{19}$ vertices.

Assume now that k is even. From the $(k - 1)$ -colored point set \widehat{S} built previously, we can obtain a k -colored point set \widehat{S}' by adding a new point with a different color. Since every perfect rainbow polygon for \widehat{S} has at least $\frac{20(k-2)-8}{19} = \frac{40\lfloor(k-1)/2\rfloor-8}{19}$ vertices, then every perfect rainbow polygon for \widehat{S}' also has at least $\frac{40\lfloor(k-1)/2\rfloor-8}{19}$ vertices. \square

6 Conclusions

In this paper, we studied the perfect rainbow polygon problem and we proved that the rainbow index of k satisfies $\frac{40\lfloor(k-1)/2\rfloor-8}{19} \leq \text{rb-index}(k) \leq 10\lfloor\frac{k}{7}\rfloor + 11$, for $k \geq 5$. We also showed that $k = 7$ is the first value such that $\text{rb-index}(k) \neq k$. Our bounds are based on the equivalence between perfect rainbow polygons and noncrossing covering trees.

Several open questions arise in relation to this problem. For instance, we conjecture that given a colored point set S , finding a minimum perfect rainbow polygon for S is NP-hard. Another interesting question is to close the gap between the lower and upper bounds on the rainbow index.

Acknowledgments

David Flores, David Orden, Javier Tejel, Jorge Urrutia, and Birgit Vogtenhuber are supported by the H2020-MSCA-RISE project 734922-CONNECT. Research by David Flores-Peñaloza was supported by the grant UNAM PAPIIT IN117317. Research by Mikio Kano was supported by JSPS KAKENHI Grant Number 16K05248. Research by Leonardo Martínez-Sandoval was supported by the grant ANR-17-CE40-0018 of the French National Research Agency ANR (project CAPPS). Research by David Orden was supported by project MTM2017-83750-P of the Spanish Ministry of Science (AEI/FEDER, UE) and by Project PID2019-104129GB-I00 /

AEI / 10.13039/501100011033 of the Spanish Ministry of Science and Innovation. Research by Javier Tejel was supported by MINECO project MTM2015-63791-R, Gobierno de Aragón under Grant E41-17R (FEDER), and Project PID2019-104129GB-I00 / AEI / 10.13039/501100011033 of the Spanish Ministry of Science and Innovation. Research by Csaba D. Tóth was supported by NSF awards CCF-1422311, CCF-1423615, and DMS-1800734. Research by Jorge Urrutia was supported by UNAM project PAPIIT IN102117. Research by Birgit Vogtenhuber was supported by the Austrian Science Fund within the collaborative DACH project *Arrangements and Drawings* as FWF project I 3340-N35.

References

- [1] Bernardo M. Ábrego, Silvia Fernández-Merchant, Mikio Kano, David Orden, Pablo Pérez-Lantero, Carlos Seara, and Javier Tejel. $K_{1,3}$ -covering red and blue points in the plane. *Discrete Mathematics & Theoretical Computer Science*, 21(3):6, 2019. doi:10.23638/DMTCS-21-3-6.
- [2] A. Karim Abu-Affash, Paz Carmi, Matthew J. Katz, and Yohai Trabelsi. Bottleneck non-crossing matching in the plane. *Computational Geometry*, 47(3):447–457, 2014. doi:10.1016/j.comgeo.2013.10.005.
- [3] Greg Aloupis, Jean Cardinal, Sébastien Collette, Erik D. Demaine, Martin L. Demaine, Muriel Dulieu, Ruy Fabila Monroy, Vi Hart, Ferran Hurtado, Stefan Langerman, Maria Saumell, Carlos Seara, and Perouz Taslakian. Non-crossing matchings of points with geometric objects. *Computational Geometry*, 46(1):78–92, 2013. doi:10.1016/j.comgeo.2012.04.005.
- [4] Greg Aloupis, Jean Cardinal, Sébastien Collette, Shinji Imahori, Matias Korman, Stefan Langerman, Oded Schwartz, Shakhar Smorodinsky, and Perouz Taslakian. Colorful strips. *Graphs and Combinatorics*, 27(3):327–339, 2011. doi:10.1007/s00373-011-1014-5.
- [5] Esther M. Arkin, Joseph S. B. Mitchell, and Christine D. Piatko. Minimum-link watchman tours. *Information Processing Letters*, 86(4):203–207, 2003. doi:10.1016/S0020-0190(02)00502-1.
- [6] Luis Barba and Patrick Schneider. Sharing a pizza: Bisecting masses with two cuts. In Joachim Gudmundsson and Michiel H. M. Smid, editors, *Proc. 29th Canadian Conference on Computational Geometry (CCCG)*, pages 174–178, 2017. URL: <http://2017.cccg.ca/proceedings/CCCG2017.pdf>.
- [7] Sergey Bereg, Ferran Hurtado, Mikio Kano, Matias Korman, Dolores Lara, Carlos Seara, Rodrigo I. Silveira, Jorge Urrutia, and Kevin Verbeek. Balanced partitions of 3-colored geometric sets in the plane. *Discrete Applied Mathematics*, 181:21–32, 2015. doi:10.1016/j.dam.2014.10.015.
- [8] Sergey Bereg and Mikio Kano. Balanced line for a 3-colored point set in the plane. *The Electronic Journal of Combinatorics*, 19(1):P33, 2012. URL: <http://www.combinatorics.org/ojs/index.php/eljc/article/view/v19i1p33>.
- [9] Sergey Bereg, Feifei Ma, Wencheng Wang, Jian Zhang, and Binhai Zhu. On some matching problems under the color-spanning model. *Theoretical Computer Science*, 786:26–31, 2019. doi:10.1016/j.tcs.2018.08.008.

- [10] Sergei Bespamyatnikh, David G. Kirkpatrick, and Jack Snoeyink. Generalizing ham sandwich cuts to equitable subdivisions. *Discrete & Computational Geometry*, 24(4):605–622, 2000. doi:10.1007/s004540010065.
- [11] Ahmad Biniiaz, Prosenjit Bose, Anil Maheshwari, and Michiel H. M. Smid. Plane bichromatic trees of low degree. *Discrete & Computational Geometry*, 59(4):864–885, 2018. doi:10.1007/s00454-017-9881-z.
- [12] Pavle V. M. Blagojevic, Günter Rote, Johanna K. Steinmeyer, and Günter M. Ziegler. Convex equipartitions of colored point sets. *Discrete & Computational Geometry*, 61(2):355–363, 2019. doi:10.1007/s00454-017-9959-7.
- [13] Björn Brodén, Mikael Hammar, and Bengt J. Nilsson. Guarding lines and 2-link polygons is APX-hard. In *Proc. 13th Canadian Conference on Computational Geometry (CCCG)*, pages 45–48, 2001. URL: <http://www.cccg.ca/proceedings/2001/mikael-2351.ps.gz>.
- [14] Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars. *Computational Geometry: Algorithms and Applications*. Springer, Berlin, 3rd edition, 2008. doi:10.1007/978-3-540-77974-2.
- [15] Olivier Devillers, Ferran Hurtado, Gyula Károlyi, and Carlos Seara. Chromatic variants of the Erdős-Szekeres theorem on points in convex position. *Computational Geometry*, 26(3):193–208, 2003. doi:10.1016/S0925-7721(03)00013-0.
- [16] Adrian Dumitrescu, Dániel Gerbner, Balázs Keszegh, and Csaba D. Tóth. Covering paths for planar point sets. *Discrete & Computational Geometry*, 51(2):462–484, 2014. doi:10.1007/s00454-013-9563-4.
- [17] Adrian Dumitrescu and Minghui Jiang. On the approximability of covering points by lines and related problems. *Computational Geometry*, 48(9):703–717, 2015. doi:10.1016/j.comgeo.2015.06.006.
- [18] Adrian Dumitrescu and Rick Kaye. Matching colored points in the plane: Some new results. *Computational Geometry*, 19(1):69–85, 2001. doi:10.1016/S0925-7721(01)00007-4.
- [19] Chenglin Fan, Jun Luo, Wencheng Wang, Farong Zhong, and Binhai Zhu. On some proximity problems of colored sets. *Journal of Computer Science and Technology*, 29(5):879–886, 2014. doi:10.1007/s11390-014-1475-0.
- [20] Rudolf Fleischer and Xiaoming Xu. Computing minimum diameter color-spanning sets is hard. *Information Processing Letters*, 111(21-22):1054–1056, 2011. doi:10.1016/j.ipl.2011.07.015.
- [21] Andreas F. Holmsen, Jan Kynčl, and Claudiu Valculescu. Near equipartitions of colored point sets. *Computational Geometry*, 65:35–42, 2017. doi:10.1016/j.comgeo.2017.05.001.
- [22] Riko Jacob and Gerth Stølting Brodal. Dynamic planar convex hull. Manuscript, 2019. arXiv:1902.11169.
- [23] Wenqi Ju, Chenglin Fan, Jun Luo, Binhai Zhu, and Ovidiu Daescu. On some geometric problems of color-spanning sets. *Journal of Combinatorial Optimization*, 26(2):266–283, 2013. doi:10.1007/s10878-012-9458-y.

- [24] Atsushi Kaneko and Mikio Kano. Discrete geometry on red and blue points in the plane: A survey. In Boris Aronov, Saugata Basu, János Pach, and Micha Sharir, editors, *Discrete and Computational Geometry: The Goodman-Pollack Festschrift*, pages 551–570. Springer, Berlin, 2003. doi:10.1007/978-3-642-55566-4_25.
- [25] Mikio Kano and Jan Kynčl. The hamburger theorem. *Computational Geometry*, 68:167–173, 2018. doi:10.1016/j.comgeo.2017.06.012.
- [26] Mikio Kano, Kazuhiro Suzuki, and Miyuki Uno. Properly colored geometric matchings and 3-trees without crossings on multicolored points in the plane. In Jin Akiyama, Hiro Ito, and Toshinori Sakai, editors, *Discrete and Computational Geometry and Graphs (JCDCGG)*, volume 8845 of *LNCS*, pages 96–111. Springer, 2014. doi:10.1007/978-3-319-13287-7.
- [27] Stefan Kratsch, Geevarghese Philip, and Saurabh Ray. Point line cover: The easy kernel is essentially tight. *ACM Trans. Algorithms*, 12(3):40:1–40:16, 2016. doi:10.1145/2832912.
- [28] V. S. Anil Kumar, Sunil Arya, and H. Ramesh. Hardness of set cover with intersection 1. In Ugo Montanari, José D. P. Rolim, and Emo Welzl, editors, *Proc. 27th International Colloquium on Automata, Languages and Programming (ICALP)*, volume 1853 of *LNCS*, pages 624–635. Springer, 2000. doi:10.1007/3-540-45022-X_53.
- [29] Jan Kynčl, János Pach, and Géza Tóth. Long alternating paths in bicolored point sets. *Discrete Mathematics*, 308(19):4315–4321, 2008. doi:10.1016/j.disc.2007.08.013.
- [30] Jiří Matoušek. *Geometric Discrepancy* Springer-Verlag, volume 18 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1999. doi:10.1007/978-3-642-03942-3.
- [31] Jiří Matoušek. *Using the Borsuk-Ulam Theorem*. Universitext. Springer-Verlag, Berlin, 2003. doi:10.1007/978-3-540-76649-0.
- [32] Nimrod Megiddo and Arie Tamir. On the complexity of locating linear facilities in the plane. *Operations Research Letters*, 1(5):194–197, 1982. doi:10.1016/0167-6377(82)90039-6.
- [33] Jonas Prunte. Minimum diameter color-spanning sets revisited. *Discrete Optimization*, 34, 2019. doi:10.1016/j.disopt.2019.07.001.
- [34] Toshinori Sakai. Balanced convex partitions of measures in \mathbb{R}^2 . *Graphs and Combinatorics*, 18(1):169–192, 2002. doi:10.1007/s003730200011.