Perfect rainbow polygons for colored point sets in the plane

David Flores-Peñaloza[*1], Mikio Kano[†2], Leonardo Martínez-Sandoval[‡3], David Orden[§4], Javier Tejel[¶5], Csaba D. Tóth[‖6], Jorge Urrutia[**7], and Birgit Vogtenhuber[††8]

[1] Departamento de Matemáticas, Facultad de Ciencias, Universidad Nacional Autónoma de México, Mexico.
[2] Ibaraki University, Hitachi, Ibaraki, Japan.
[3] Sorbonne Université, Institut de Mathématiques de Jussieu - Paris Rive Gauche (UMR 7586), Paris, France.
[4] Departamento de Física y Matemáticas, Universidad de Alcalá, Alcalá de Henares, Spain.
[5] Departamento de Métodos Estadísticos, IUMA, Universidad de Zaragoza, Zaragoza, Spain.
[6] Department of Mathematics, California State University Northridge, Los Angeles, CA, USA.
[7] Instituto de Matemáticas, Universidad Nacional Autónoma de México, Mexico.
[8] Institute of Software Technology, Graz University of Technology, Graz, Austria.

Abstract

Given a planar n-colored point set $S = S_1 \dot\cup \ldots \dot\cup S_n$ in general position, a simple polygon P is called a *perfect rainbow* polygon if it contains exactly one point of each color. The *rainbow index* r_n is the minimum integer m such that every n-colored point set S has a perfect rainbow polygon with at most m vertices. We determine the values of r_n for $n \leq 7$, and prove that in general $\frac{20n-28}{19} \leq r_n \leq \frac{10n}{7} + 11$.

1 Introduction

The study of colored point sets has attracted a lot of interest, and particular attention has been given to 2-, 3-, and 4-colored point sets, see [1], [2], and [4]. Let $S = S_1 \dot\cup \ldots \dot\cup S_n$ be an n-colored point set in the plane, where for every $1 \leq i \leq n$, S_i is the set of elements of S colored with color c_i. We assume that each S_i is non-empty and that S is in general position. All polygons considered here are simple polygons. An m-gon is a polygon with m vertices, and m-gons for $m = 3, 4, 5, 6, 7$ are called triangles, quadrilaterals, pentagons, hexagons, and heptagons, respectively.

Given an n-colored point set S, a polygon P is called a *perfect rainbow* polygon if it contains exactly one point of each color. We are interested in finding the smallest number r_n such that any n-colored point set has a perfect rainbow polygon with at most r_n vertices.

It is well know that for every 3-colored point set S, there exists an empty triangle such that its vertices are in S and have different colors, that is, $r_3 = 3$. In this work, we determine the exact values of r_n up to $n = 7$, which is indeed the first case where $r_n > n$. Moreover, for general n, we show lower and upper bounds on r_n. Due to space constraints, most proofs are only sketched or completely deferred to the full paper.

2 Rainbow indexes for $n \leq 7$

Theorem 1 *The rainbow indexes for $n \leq 7$ are: $r_3 = 3$, $r_4 = 4$, $r_5 = 5$, $r_6 = 6$, and $r_7 = 8$.*

Proof. We sketch the proofs for r_6 and r_7. Figure 1 illustrates the lower bounds. For the upper bound of r_6, we prove that parallel lines ℓ_3 and ℓ_4 as in Figure 2 do exist and we work out the cases there. For r_7, we proceed analogously, constructing the perfect rainbow 8-gon by adding two edges to the hexagon in order to capture a point of the seventh color. $\square$

[*] Email: dflorespenaloza@ciencias.unam.mx. Research supported by the grant UNAM PAPIIT IN117317

[†] Email: mikio.kano.math@vc.ibaraki.ac.jp. Research supported by JSPS KAKENHI Grant Number 16K05248

[‡] Email: leomtz@im.unam.mx. Research supported by the grant ANR-17-CE40-0018 of the French National Research Agency ANR (project CAPPS).

[§] Email: david.orden@uah.es. Research supported by project MTM2017-83750-P of the Spanish Ministry of Science (AEI/FEDER, UE).

[¶] Email: jtejel@unizar.es. Research supported by MINECO project MTM2015-63791-R and Gobierno de Aragón under Grant E41-17R (FEDER).

[‖] Email: csaba.toth@csun.edu. Research supported by NSF awards CCF-1422311, CCF-1423615, and DMS-1800734.

[**] Email: urrutia@matem.unam.mx. Research supported by UNAM project PAPIIT-IN102117

[††] Email: bvogt@ist.tugraz.at. Research supported by the Austrian Science Fund within the collaborative DACH project *Arrangements and Drawings* as FWF project I 3340-N35.

This project has been supported by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 734922.

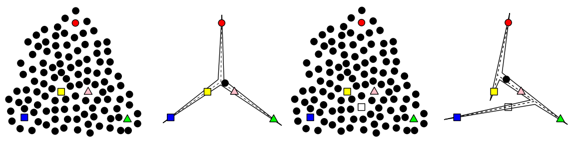

Figure 1: Lower bound constructions for r_6 and r_7.

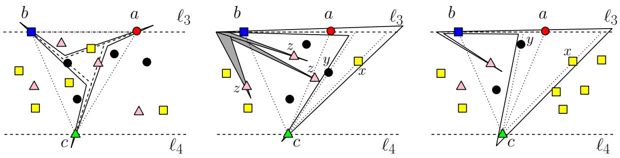

Figure 2: Cases for the upper bound of r_6.

3 Upper bound for rainbow indexes

We show in this section that for any n-colored point set, there exists a perfect rainbow polygon of size at most $\frac{10n}{7} + 11$. To that end, we first give a lemma showing that seven points (without colors) inside a vertical strip can be always covered by a tree with four vertices and a segment such that their union is inside the strip and is non-crossing (see Figure 3b).

Lemma 2 *Let $\{p_1, \ldots, p_7\}$ be the seven points of a point set S, ordered from left to right. Let B be the strip defined by the two vertical lines passing through p_1 and p_7, respectively. Then, there exist two non-crossing trees T_1 and T_2, the first one of order 4 and the second one of order 2, such that:*

(i) The union of T_1 and T_2 covers the points of S, is inside B and is non-crossing.

(ii) For every T_i, $i = 1, 2$, there exists a special leaf v_i such that the extension of the edge in T_i incident to v_i goes to the left. Moreover, if the extension at v_i hits T_j, then the extension at v_j does not hit T_i, that is, the two trees and the two extensions do not create cycles.

Theorem 3 *For any n-colored point set S, there is a perfect rainbow polygon of size at most $\frac{10n}{7} + 11$.*

Figure 3 illustrates the method to obtain such a perfect rainbow polygon. Assume that $n = 7k$. We choose n points such that each point has a different color. We divide the n points from left to right into k groups of seven points each and apply Lemma 2 to each group to cover the seven points by two trees. Then we join all trees to a long vertical segment P' placed to the left, by extending the edge adjacent to the special leaf of each tree. Finally, we build a perfect rainbow polygon by surrounding the edges of the obtained tree.

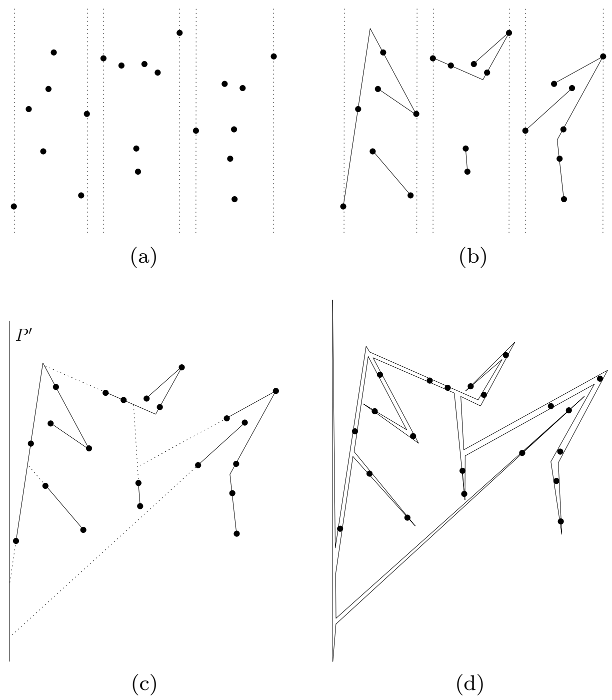

Figure 3: (a) Dividing the n points into groups of size 7. (b) Applying Lemma 2 to each group. (c) Joining all trees to the segment P'. (d) Building the perfect rainbow polygon.

4 Lower bound for rainbow indexes

For every $k \geq 3$, Dumitrescu et al. [3] constructed a set S of $n = 2k$ points in the plane such that 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*, where a segment is defined as a chain of collinear edges, and this bound is the best possible.

In this section, we use the point sets constructed in [3] to derive a lower bound for the complexity of a covering tree under a new measure that we define here. This bound, in turn, yields a lower bound on the complexity of simple polygons that contain the given points and have arbitrarily small area.

Covering Trees versus Polygons. Let T be a noncrossing geometric tree (i.e., plane straight-line tree). Similarly to [3], 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 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 the 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. 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 geometric 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 4 *Let T be a noncrossing geometric tree and $\mathcal{M}$ a partition of the edges into the minimum number of pairwise noncrossing segments. If $s \geq 2$ then for every $\varepsilon > 0$, there is 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 the 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 (i.e., 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 then 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, the area of the δ-neighborhood of T is bounded 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 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 relate point sets to covering trees.

Lemma 5 *Let S be a finite set of points in the plane, not all on a line. Then there exists an $\varepsilon > 0$ such that if S is contained in a simple polygon P with m vertices and* $\text{area}(P) \leq \varepsilon$, *then S admits a noncrossing covering tree T 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 $k \in \mathbb{N}$, there exists a simple polygon P_k with precisely m vertices such that $S \subset \text{int}(P_k)$ and $\text{area}(P_k) \leq \frac{1}{k}$. The real projective plane $P\mathbb{R}^2$ is a compactification of $\mathbb{R}^2$. By compactness, the sequence $(P_k)_{k \geq 3}$ contains a convergent subsequence of polygons in $P\mathbb{R}^2$. The limit is a weakly simple polygon P with precisely m vertices (some of which may coincide) such that $S \subset P$ and $\text{area}(P_k) = 0$. The edges of P form a set of pairwise noncrossing line segments (albeit with possible overlaps) whose union is a connected set that contains S. In particular, the union of the m edges of P form a noncrossing covering tree T for S. 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. These segments are pairwise noncrossing (since the edges of P are pairwise noncrossing), and yield a covering of T with a set $\mathcal{M}$ of pairwise nonoverlapping and noncrossing segments. Analogously to the proof of Lemma 4, 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 $\mathcal{M}$. $\square$

Construction. We use the point set constructed by Dumitrescu et al. [3]. We review some of its properties here. For every $k \in \mathbb{N}$, they construct a set of $n = 2k$ points, $S = \{a_i, b_i : i = 1, \ldots, k\}$. The pairs $\{a_i, b_i\}$ ($i = 1, \ldots, k$}) are called *twins*. The points a_i ($i = 1, \ldots, k$) lie on the parabola $\alpha = \{(x, y) : y = x^2\}$, sorted by increasing x-coordinate. The points b_i ($i = 1, \ldots, k$) lie on a convex curve β above α, such that $\text{dist}(a_i, b_i) < \varepsilon$ for a sufficiently small ε, the lines a_ib_i are almost vertical with monotonically increasing positive slopes (hence the supporting lines of any two twins intersect below α). For $i = 1, \ldots, k$, they also define pairwise disjoint disks $D_i(\varepsilon)$ of radius ε centered at a_i such that $b_i \in D_i(\varepsilon)$. 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. Finally, 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-corrdinate, the supporting lines of c_1c_4 and c_2c_3 cross to the left of these points.

Analysis. Let S be a set of $n = 2k$ points defined in [3] 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 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 chains of perfect segments; we call these *perfect chains*. We rephrase two lemmas from [3] using this terminology.

Lemma 6 *[3, 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 7 *[3, 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 (i.e., right) endpoint of pq. Then q is the endpoint of a perfect chain.*

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. A careful adaptation of a charging scheme from [3, Lemma 4] yields the following result, where t is the number of forks (with multiplicity) in $\mathcal{M}$.

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

The combination of Lemma 8 and $n = s_1 + 2s_2$ yields the following lemma.

Lemma 9 *Let S be a set of $n = 2k \geq 4$ points from [3]. Then every covering tree T of S satisfies $2s + t \geq (20n - 8)/19$.*

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

Theorem 10 *For every odd integer $m \geq 5$, there exists a finite set of m-colored points in the plane such that every perfect rainbow polygon has at least $(20m - 28)/19$ vertices.*

Proof. Let $n = m - 1$. We construct the point set $S = S_1 \dot\cup S_2$ in general position as follows. Let S_1 be the set of $n = 2k \geq 4$ points from [3], where each point has a unique color. We can prove that there is an $\varepsilon > 0$ such that if there is a simple polygon of area at most ε with $(20m - 8)/19$ vertices that contains S_1, then S_1 admits a noncrossing spanning tree and a partition of its edges into segments such that $2s + t \leq (20m - 8)/19$.

Let S_2 be the union of two disjoint $\varepsilon/(2n)$-nets for the range space of triangles, that is, every triangle of area $\varepsilon/(2n)$ or more contains at least two points in S_2. All points in S_2 have color m. Now suppose, for the sake of contradiction, that there exists a perfect rainbow polygon P with x vertices where $x < (20m - 28)/19$. Triangulate P arbitrarily into $x - 2$ triangles. The area of the largest triangle is at least $\text{area}(P)/(x - 2)$. Since this triangle contains at most one point from S_2, we have $\text{area}(P)/(x - 2) \leq \varepsilon/(2n)$, and so $\text{area}(P) \leq \varepsilon$. By the choice of ε, S_1 admits a noncrossing spanning tree and a partition of its edges into segments such that $2s + t < (20m - 8)/19$. This can be proved to be a contradiction, which completes the proof. $\square$

References

[1] S. Bereg, F. Hurtado, M.Kano, M. Kormand, D. Lara, C. Seara, R. Silveira, J. Urrutia, and K. Verbeek, Balanced partitions of 3-colored geometric sets in the plane *Discrete Applied Math.*, **181** (2015) 21–32

[2] S. Bereg, M. Kano, Balanced line for a 3-colored point set in the plane, *Electron. J. Combin.*, **19** (2012) 33.

[3] A. Dumitrescu, D. Gerbner, B. Keszegh, C. D. Tóth, Covering paths for planar point sets, *Discrete & Computational Geometry*, **51:2** (2014) 462–484.

[4] M. Kano and J. Kynčl, The hamburger theorem, *Computational Geometry: Theory and Applications*, **68** (2018) 167–173.