

Coloring Geometric Range Spaces

Greg Aloupis · Jean Cardinal · Sébastien Collette ·
Stefan Langerman · Shakhar Smorodinsky

Received: 11 August 2007 / Revised: 1 September 2008 / Accepted: 8 September 2008 /
Published online: 26 September 2008
© Springer Science+Business Media, LLC 2008

Abstract We study several coloring problems for geometric range-spaces. In addition to their theoretical interest, some of these problems arise in sensor networks. Given a set of points in \mathbb{R}^2 or \mathbb{R}^3 , we want to color them so that every region of a certain family (e.g., every disk containing at least a certain number of points) contains points of many (say, k) different colors. In this paper, we think of the number of colors and the number of points as functions of k . Obviously, for a fixed k using k colors, it is not always possible to ensure that every region containing k points has all colors present. Thus, we introduce two types of relaxations: either we allow the number of colors used to increase to $c(k)$, or we require that the number of points in each region increases to $p(k)$.

Symmetrically, given a finite set of regions in \mathbb{R}^2 or \mathbb{R}^3 , we want to color them so that every point covered by a sufficiently large number of regions is contained in regions of k different colors. This requires the number of covering regions or the number of allowed colors to be greater than k .

The research of G. Aloupis, J. Cardinal, S. Collette and S. Langerman was supported by the Communauté française de Belgique—Actions de Recherche Concertées (ARC).

G. Aloupis (✉) · J. Cardinal · S. Collette · S. Langerman
Université Libre de Bruxelles (U.L.B.), CP212, Bld. du Triomphe, 1050 Brussels, Belgium
e-mail: greg@cg.scs.carleton.ca

J. Cardinal
e-mail: jcardin@ulb.ac.be

S. Collette
e-mail: secollet@ulb.ac.be

S. Langerman
e-mail: slanger@ulb.ac.be

S. Smorodinsky
Ben-Gurion University, Be'er Sheva 84105, Israel
e-mail: shakhar@math.bgu.ac.il

The goal of this paper is to bound these two functions for several types of region families, such as halfplanes, halfspaces, disks, and pseudo-disks. This is related to previous results of Pach, Tardos, and Tóth on decompositions of coverings.

Keywords Coloring · Covering · Decompositions · Geometric hypergraphs

1 Introduction

A large amount of research in discrete geometry concerns algorithms and bounds for sampling and decomposing sets of simple geometric objects. The now standard notion of ε -net [7] and the theory of geometric discrepancy [5] are famous illustrations of the importance of this field.

In this contribution, we are interested in coloring (or, equivalently, partitioning) finite sets of points in \mathbb{R}^2 or \mathbb{R}^3 so that any region (within a specified family) that contains at least some fixed number of points also contains a significant number of distinctly colored points. For example, we study the following problem: *Does there exist a constant α such that given any set of points in the plane, it is always possible to color the points with k colors so that any halfplane containing at least αk points contains a point of each color?* In Sect. 2 we answer this question on the affirmative.

In all cases we consider, the number of colors and the number of points required within regions are significantly smaller than the size of the point set. We also allow the number of available colors and the number of required distinct colors to be different. We ask, for instance, *does there exist a constant α such that given a set of points in the plane, it is always possible to color the points with αk colors so that any halfplane containing at least k points also contains points of k distinct colors?* We show that this is true as well and even give a small upper bound on the value of α . We ask similar questions for other types of regions such as disks and pseudo-disks (simple closed Jordan regions, the boundaries of any pair of which intersect at most twice).

These types of problems can be seen as coloring range spaces induced by intersections of sets of points with geometric objects. The corresponding dual range spaces are those obtained by considering a finite set of regions in \mathbb{R}^2 or \mathbb{R}^3 and defining the ranges as the subsets of all regions containing a given point, for every possible point. We also consider coloring problems on these kinds of range spaces. The types of problems we ask when dealing with dual range spaces are analogous to the preceding questions. For instance: *Does there exist a constant α such that given any set of disks in the plane, it is always possible to color the disks with αk colors while ensuring that any point contained in at least k disks is contained in disks of k distinct colors?*

Before stating the results formally, we introduce several definitions.

Definitions A range space (or hypergraph) is a pair (S, R) where S is a set (called the ground set) and R is a set of subsets of S . In this paper, we consider finite restrictions of infinite geometric range spaces of the form $S = (\mathbb{R}^d, \mathcal{R})$ for $d = 2$ or 3 , where \mathcal{R} is an infinite family of regions of \mathbb{R}^d . Such a finite restriction is a range space (S, R) where the ground set S is a finite set of points in \mathbb{R}^d and the set of ranges R is the collection of subsets of S defined by the intersection of S with elements of \mathcal{R} : $R = \{S \cap r : r \in \mathcal{R}\}$.

We also consider the corresponding *dual range spaces*, denoted by $\tilde{\mathcal{S}}$, of the form $\tilde{\mathcal{S}} = (\mathcal{R}, \{r(p) : p \in \mathbb{R}^d\})$, where $r(p) = \{r \in \mathcal{R} : p \in r\}$ is the set of regions containing the point p . The finite restrictions of these dual range spaces are of the form $(\mathcal{S}, \{r(p) \cap \mathcal{S} : p \in \mathbb{R}^d\})$, where $\mathcal{S} \subset \mathcal{R}$ is finite.

A *coloring* of a range space is an assignment of colors to the elements of the ground set. A c -coloring is a coloring that uses exactly c colors. A range is said to be k -colorful with respect to some coloring if it contains at least k elements of distinct color. We are interested in the following two types of functions for a range space \mathcal{S} :

1. $c_{\mathcal{S}}(k)$ is the minimum number for which there always exists a $c_{\mathcal{S}}(k)$ -coloring of any finite restriction of \mathcal{S} , such that every range r is $\min\{|r|, k\}$ -colorful.
2. $p_{\mathcal{S}}(k)$ is the minimum number for which there always exists a k -coloring of any finite restriction of \mathcal{S} such that every range of size at least $p_{\mathcal{S}}(k)$ is k -colorful.

Note that $c_{\mathcal{S}}(k)$ and $p_{\mathcal{S}}(k)$ are nondecreasing functions. The functions apply to any range space \mathcal{S} or its dual $\tilde{\mathcal{S}}$. The goal of this paper is to provide upper bounds on $c_{\mathcal{S}}(k)$, $p_{\mathcal{S}}(k)$, $c_{\tilde{\mathcal{S}}}(k)$, and $p_{\tilde{\mathcal{S}}}(k)$ for various types of families \mathcal{R} of regions in \mathbb{R}^2 and \mathbb{R}^3 . We always consider the ground sets to be large enough, that is, unbounded in terms of k .

Previous Results The functions defined above are related to two previously studied problems. The first one is a decomposition of f -fold coverings in the plane: given a covering of the plane (or of some simple compact subset such as a square) by a set of regions such that every point is covered by at least f regions, is it possible to decompose it into two disjoint coverings? This question was first asked by Pach in 1980 [11]. It is similar to deciding whether $p_{\tilde{\mathcal{S}}}(2) \leq f$ for the dual range space $\tilde{\mathcal{S}}$ defined by the considered family of regions, the difference being that we do not assume that all points are f -covered. This difference is important in some cases, for instance, it is known that all $(d+1) \cdot f$ -covers of d -space by halfspaces decompose into f covers [12], but the proof does not directly yield a bound for $p_{\tilde{\mathcal{S}}}(2)$.

For \mathcal{T} the range space defined by translates of a centrally symmetric convex polygon, Pach [11] proved that $p_{\mathcal{T}}(2)$ exists. His proof yields an exponential upper bound on $p_{\mathcal{T}}(k)$. Then Pach and Tóth [14] improved this by showing that $p_{\mathcal{T}}(k) = O(k^2)$ and $p_{\tilde{\mathcal{T}}}(k) = O(k^2)$. So for these types of regions, a covering can be decomposed into k coverings if each point is covered at least ck^2 times for some constant c . On the negative side, for the range space induced by arbitrary disks (denoted by \mathcal{D}), Mani and Pach [10] proved that even $p_{\mathcal{D}}(2)$ is unbounded: for any constant k , there exists a set of points that cannot be 2-colored so that all open disks containing at least k points contain one point of each color (see also [15]). In [15], a similar result is obtained for $p_{\tilde{\mathcal{A}}}(2)$, where \mathcal{A} is the range space induced by the family of either strips (sets defined by all points enclosed between two parallel lines), axis-aligned rectangles, or translates of concave quadrilaterals. More precisely, for any integer f , there exist f -fold coverings with strips, or axis-aligned rectangles, or concave quadrilaterals, that cannot be decomposed into two coverings. The fact that $p_{\tilde{\mathcal{S}}}(2)$ is unbounded implies that for all $k > 2$, $p_{\tilde{\mathcal{S}}}(k)$ is unbounded as well, since any bound for the latter would imply a bound for the former by merging color classes. The previous impossibilities constitute our main motivation for introducing some slack and defining the problem of $c(k)$ -coloring a finite range space such that ranges are k -colorful, with $k \leq c(k)$.

The second previously studied problem is that of computing the chromatic number of geometric hypergraphs defined as the minimum number of colors needed to make all ranges nonmonochromatic, i.e., 2-colorful [17]. One of the main results of that contribution is that any dual range space induced by a finite set of pseudo-disks admits an $O(1)$ -coloring that makes all ranges 2-colorful. Hence, for the family of pseudo-disks \mathcal{P} , $c_{\tilde{\mathcal{P}}}(2) = O(1)$. It was also shown that for the special case of disks, $c_{\tilde{\mathcal{D}}}(2) = 4$.

A recent result of Chen, Pach, Szegedy, and Tardos [6, Theorem 3] implies that for any constants c, p , the following holds: there exists a point set such that for any c -coloring of its elements, we can find an axis-aligned rectangle containing at least p points, all of which have the same color. In fact, this holds with high probability for random point sets. This implies that $c_{\mathcal{A}}(k)$ and $p_{\mathcal{A}}(k)$, where \mathcal{A} is the range space induced on \mathbb{R}^2 by the set of all axis-aligned rectangles, are unbounded. A previous result of Kříž and J. Nešetřil [8] on the chromatic number of Hasse diagrams of 2-dimensional posets also implies that $c_{\mathcal{A}}(2)$ is unbounded.

Furthermore, Pach, and Tardos [13] proved that for any n , there exists a set of n axis-parallel rectangles in the plane such that one needs $\Omega(\log n)$ colors for coloring the rectangles such that no point is covered by a monochromatic set, hence that $c_{\tilde{\mathcal{A}}}(2)$ is unbounded, implying that $c_{\tilde{\mathcal{A}}}(k)$ is unbounded.

Our Results In Sect. 2, we consider the range space $\mathcal{HP} = (\mathbb{R}^2, \mathcal{R})$, where \mathcal{R} is the set of all halfplanes. We prove that $c_{\mathcal{HP}}(k) \leq 3k - 2$ and $p_{\mathcal{HP}}(k) \leq 4k - 1$. In other words, we can ensure that a halfplane contains k points of different colors in two ways: either we k -color the point set but require that the halfplane contains at least $4k - 1$ points, or we allow the point set to be $(3k - 2)$ -colored. We also give lower bounds on the two functions.

In Sect. 3, we consider the range space $\mathcal{HS} = (\mathbb{R}^3, \mathcal{R})$, where \mathcal{R} is the set of all halfspaces, that is, sets of points on one side of a given plane. We prove that $c_{\mathcal{HS}}(k) = O(k)$, i.e., that points in \mathbb{R}^3 can always be colored with $O(k)$ colors so that any halfspace containing at least k points contains k distinct colors; and that $c_{\tilde{\mathcal{HS}}}(k) = O(k)$, i.e., halfspaces in \mathbb{R}^3 can be colored so that every point in the intersection of at least k of them is contained in halfspaces of k distinct colors.

We provide a number of results on range spaces defined by disks and pseudo-disks in Sect. 4. For the range space \mathcal{D} defined by disks, we prove that $c_{\mathcal{D}}(k) = O(k)$ by mapping disks in the plane to lower halfspaces in \mathbb{R}^3 and using the result of Sect. 3. For a dual range space $\tilde{\mathcal{P}}$ defined by a family of pseudo-disks, we prove that $c_{\tilde{\mathcal{P}}}(k) = O(k)$. So we can always color a set of pseudo-disks with $O(k)$ colors in such a way that any point covered by r pseudo-disks is covered by $\min\{r, k\}$ pseudo-disks with distinct colors. As a consequence, since halfplanes could be viewed as a special case of pseudo-disks, we directly have $c_{\tilde{\mathcal{HP}}}(k) = O(k)$. We also show that $c_{\mathcal{P}}(k) = O(k)$, the proof of which involves similar arguments.

By lifting a 2D point set to the unit paraboloid $z = x^2 + y^2$ in 3D, every lower halfspace in 3D isolates a set of points which is contained in a disk in the original set of points, and thus $p_{\mathcal{HS}}(k) \geq p_{\mathcal{D}}(k)$. We also prove that $p_{\tilde{\mathcal{HS}}}(k) = p_{\mathcal{HS}}(k)$: coloring halfspaces is equivalent, via projective duality, to coloring points with respect to halfspaces.

Table 1 Results for range spaces induced by various families \mathcal{R} of regions. The symbol \star indicates new results; the symbol ∞ indicates a function unbounded in terms of k

\mathcal{S}	$c_{\mathcal{S}}(k)$	$p_{\mathcal{S}}(k)$	$c_{\tilde{\mathcal{S}}}(k)$	$p_{\tilde{\mathcal{S}}}(k)$
Halfplanes	$\leq 3k - 2$ (Theorem 1) \star $\geq 3k/2$ (Theorem 2) \star	$\leq 4k - 1$ (Theorem 3) \star $\geq 4k/3$ (Theorem 4) \star	$O(k)$ (Theorem 6) \star	$\leq 8k - 3$ (Corollary 1) \star
Halfspaces in \mathbb{R}^3	$O(k)$ (Theorem 5) \star	∞ (Implied by disks)	$O(k)$ (Corollary 2) \star	∞ (Implied by disks)
Translates of a cent. sym. convex polygon	$O(k)$ (Theorem 7) \star	$O(k^2)$ [14]	$O(k)$ (Theorem 6) \star	$O(k^2)$ [14]
Axis-aligned rectangles	∞ [6]	∞ [6]	∞ [13]	∞ [15]
Disks	$O(k)$ (Corollary 3, Theorem 5) \star	∞ (open disks [15])	4 when $k = 2$ [17] $\leq 24k + 1$ (Remark 1) \star	
Pseudo-disks	$O(k)$ (Theorem 7) \star	∞ (open disks [15])	$O(1)$ when $k = 2$ [17] $O(k)$ (Theorem 6) \star	

All the proofs are constructive, and polynomial-time algorithms can easily be derived from them. The results are summarized in Table 1.

Application to Sensor Networks Let \mathcal{R} be a collection of sensors, each of which monitors the area within a surrounding disk. Assume further that each sensor has a battery life of one time unit. The goal is to monitor a given planar region A for as long as possible. If we activate all sensors in \mathcal{R} simultaneously, A will be monitored for only one time unit. This can be improved if \mathcal{R} can be partitioned into c pairwise disjoint subsets, each of which covers A . Each subset can be used in turn, allowing us to monitor A for c units of time. Obviously if there is a point in A covered by only c sensors, then we cannot partition \mathcal{R} into more than c families. Therefore it makes sense to ask the following question: what is the minimum number $p(k)$ for which we know that if every point in A is covered by $p(k)$ sensors, then we can partition \mathcal{R} into k pairwise disjoint covering subsets? This is exactly the type of problem that we described. Suppose that we now make the assumption that switching a sensor on or off does not consume energy. Then our Theorem 6 implies that we can assign time slots to each sensor, using a total of $O(k)$ slots, so that any point p that is covered by at least k sensors will be monitored at least a constant fraction of the time.

For more on the relation between these partitioning problems and sensor networks, see the paper of Venkatasubramanian et al. [3].

2 Halfplanes

In this section we study the case where the family \mathcal{R} is the set of all halfplanes in \mathbb{R}^2 . We denote by $\mathcal{HP} = (\mathbb{R}^2, \mathcal{R})$ the corresponding infinite range space.

It is not always possible to color a set of points S with k colors such that every halfplane of size k (that is, containing k points of S) is k -colorful, even for $k = 2$. The simplest example consists of an odd number of points in convex position. This is our main motivation for allowing either the number of colors or the range size to be greater than k .

For the proof of Theorems 1 and 3, the notion of Tukey depth is used:

Definition 1 [20] Given a set S of points in \mathbb{R}^d , the *Tukey depth* of a point p (not necessarily in the set) is the minimum number of points of S in a halfplane that contains p .

It is well known that for any set of n points in the plane, there exists a point in \mathbb{R}^2 at depth $t \geq n/3$. The *depth- k region* is the set of all points at Tukey depth k or more. It is easily seen that this region is the intersection of all halfplanes containing more than $n - k$ points of S , and therefore its boundary is a convex polygon. For more information on properties of Tukey depth, see [2]. We now turn to some useful observations regarding depth- k regions.

Lemma 1 Let S be a finite set of more than $3k$ points in \mathbb{R}^2 . Then every open halfplane not intersecting the depth- k region of S contains at most $2k - 2$ points of S . The corresponding closed halfplane contains at least k points.

Proof Let Π be an open halfplane not intersecting the depth- k region. Without loss of generality, we assume that its bounding line ℓ is tangent to the depth- k polygon. Let Π' be the corresponding closed halfplane. Π' contains at least k points since the point of tangency belongs to Π' and has depth k . On the other hand, ℓ contains either a side of the polygon or precisely one of its vertices, v .

In the former case, Π contains less than k points because its complement contains more than $n - k$ points. In the latter case, Π is contained in the union of two open halfplanes, Π_1 and Π_2 , whose bounding lines pass through v and its two neighbors in the polygon (respectively). Since each of Π_1 and Π_2 contains at most $k - 1$ points, Π contains at most $2k - 2$ points. \square

We define the *orientation of a halfplane* as the absolute angle of the inward normal of the line bounding it. Thus, for example, the orientation of the halfplane defined by all points lying above the x -axis is $\frac{\pi}{2}$.

Let p be a point of S lying outside the depth- k region. It is easily seen that the set of orientations of all closed halfplanes that are tangent to the depth- k region and that contain p form a closed (circular) interval of length at most π . Thus, each point may be represented as an interval on the unit circle. Let \mathcal{A} be the set of arcs corresponding to points in S outside or on the boundary of the depth- k region, and let \mathcal{A}' be the set of interiors of the elements of \mathcal{A} .

Lemma 2 Every point on the unit circle is covered by at most $2k - 1$ arcs of \mathcal{A}' , and every point that is not the endpoint of an arc is covered by at least k arcs. Furthermore, the minimum number of segments covering any point is at most $k - 1$.

Proof Every point p on the unit circle represents the orientation of a closed halfplane Π' tangent to the depth- k region. Thus if p is not the endpoint of an arc, then the number of arcs that cover p is at least the number of points in Π' , which is at least k by Lemma 1. As in the proof of Lemma 1, if the boundary ℓ of the halfplane contains a vertex v but no edge of the depth k region, then Π' is contained in the union of v and two open halfplanes Π_1 and Π_2 whose bounding lines pass through v and its two neighboring edges in the polygon. Since each of Π_1 and Π_2 contains at most $k - 1$ points, and there might be a point at v , Π' contains at most $2k - 1$ points. If ℓ contains an edge of the depth k region, then all points on ℓ correspond to either empty arcs or to the endpoint of some arc. Thus the arcs that cover p correspond to points in the open halfplane Π bounded by ℓ and their number is at most $k - 1$. \square

Theorem 1 $c_{HP}(k) \leq 3k - 2$. That is, we can color any set of points in the plane with $3k - 2$ colors such that any halfplane containing h points is $\min\{h, k\}$ -colorful.

Proof A proper coloring of a set of arcs on the unit circle is an assignment of colors to the arcs such that no pair of arcs of the same color overlap. In [19] it was proved that every set of arcs on the unit circle has a proper coloring with $m + M$ colors, where m (respectively, M) is the minimum (respectively, maximum) number of arcs covering each point of the circle. Combining this with Lemma 2, we conclude that the corresponding set \mathcal{A}' can be $(3k - 2)$ -colored. Accordingly we can color the points (outside the depth- k region) of S that correspond to \mathcal{A}' . The remaining points are colored arbitrarily. Thus there exists a $(3k - 2)$ -coloring of S such that every open halfplane not intersecting—but tangent to—the depth- k region is colorful (the colors of points inside that halfplane are pairwise distinct).


Now it remains to prove that every halfplane of size h is $\min\{h, k\}$ -colorful. Given such a halfplane Π , there are two cases: (i) Π does not intersect the depth- k region. Then it is strictly contained in an open halfplane Π' whose boundary line is tangent to the depth- k region, and thus no two points in it are colored with the same color. (ii) Π intersects the depth- k -region. Then it contains a closed halfplane Π' tangent to it. If the point p on the circle corresponding to Π' is not the endpoint of an arc, then Π' contains at least k points of different colors. If p is the endpoint of an arc, then Π' contains at least all points corresponding to arcs that cover a point infinitesimally to the left of p , which also have at least k different colors. \square

We also get the following lower bound.

Theorem 2 $c_{HP}(k) \geq 3k/2$. That is, coloring a set of points so that any halfplane containing h points is $\min\{h, k\}$ -colorful may require $3k/2$ colors.

Proof This lower bound is achieved by a point set consisting of three subsets of $k/2$ points, aligned on three halflines with a common origin and respective angles 0 , $2\pi/3$, and $4\pi/3$ (see Fig. 1). These $3k/2$ points must all have different colors, otherwise we can find a halfplane that contains two such groups, and less than k colors. The number of points can be made arbitrarily larger by adding points close to the origin of the halflines. \square

Fig. 1 Point set used in the proofs of Theorems 2 and 4



We now consider the depth- $2k$ region. As described above, points outside the depth- $2k$ region are associated with a set of closed arcs, \mathcal{A} , on the unit circle. Recall that each arc in \mathcal{A} has length at most π and that by Lemma 1 every point on the unit circle is covered by at least $2k$ arcs.

Lemma 3 *Let \mathcal{A} be a set of arcs of length at most π on the unit circle. If every point on the circle is covered at least $2k$ times, then \mathcal{A} has a k -colorful k -coloring.*

Proof As noticed by Pach [12], a $2k$ -covering of the unit circle with arcs of length at most π is decomposable into k disjoint coverings (this is seen by repeatedly removing a minimal covering of the unit circle). Therefore we can assign one color to all arcs within each covering, so that every point on the circle is covered by k colors. \square

Theorem 3 $p_{HP}(k) \leq 4k - 1$. *That is, we can color any set of points in the plane with k colors such that any halfplane containing at least $4k - 1$ points is k -colorful.*

Proof Let \mathcal{A} be the set of arcs corresponding to the points that lie outside or on the boundary of the depth- $2k$ region. By Lemma 3, \mathcal{A} can be made k -colorful, as it covers every point of the unit circle at least $2k$ times. This means that there exists a k -coloring of S such that every closed halfplane tangent to the depth- $2k$ region is k -colorful. As we consider large point sets in comparison to k , there always exists a depth- $2k$ region (specifically, as long as $n \geq 6k$).

Let Π be a halfplane containing at least $4k - 1$ points. Π must intersect (or touch) the depth- $2k$ region, because every open halfplane tangent to the region contains at most $4k - 2$ points, by Lemma 1. Thus Π contains a closed halfplane Π' whose boundary is tangent to the depth- $2k$ region. By construction, Π' must be k -colorful, and therefore so must be Π . \square

We get a lower bound using the same example as for $c_{HP}(k)$.

Theorem 4 $p_{HP}(k) \geq 4k/3$. *That is, coloring a set of points with k colors so that any halfplane containing at least p points is k -colorful may require that $p \geq 4k/3$.*

Proof Consider a set of $3(k + x)/2$ points forming a star with three branches of size $(k + x)/2$, denoted by G_1 , G_2 , and G_3 (see Fig. 1). We show that a k -coloring of this set where any halfplane of size $k + x$ contains all k colors exists only if $x \geq k/3$.

Suppose that there exists such a k -coloring. We denote $S_1 = G_1 \cup G_2$, $S_2 = G_2 \cup G_3$, and $S_3 = G_3 \cup G_1$. The sets S_1 , S_2 , and S_3 are contained in halfplanes and have $k + x$ points each. Thus each of them must contain all k colors.

The set G_1 contains $(k + x)/2$ points and thus points of at most $(k + x)/2$ different colors. As S_1 contains all k colors, G_2 has points of at least $k - (k + x)/2 = (k - x)/2$ colors not represented in G_1 . The same reasoning holds for S_3 and G_3 . Hence G_3 must also contain at least $(k - x)/2$ colors not represented in G_1 , and at least $(k - x)/2$ of these colors are the same as those in G_1 .

Now consider S_2 . As it contains both G_2 and G_3 , which have $(k - x)/2$ colors in common, it can contain up to $k + x - (k - x)/2 = (k + 3x)/2$ distinct colors. As we know that S_2 contains all k colors, we deduce $k \leq (k + 3x)/2$ and thus $x \geq k/3$. \square

Corollary 1 $p_{\widetilde{\mathcal{HP}}}(k) \leq 8k - 3$. That is, we can color any set of halfplanes with k colors so that any point in the plane covered by $8k - 3$ halfplanes is contained in halfplanes of k different colors.

Proof If we restrict ourselves to lower halfplanes, $p_{\widetilde{\mathcal{HP}}}(k)$ is equal to $p_{\mathcal{HP}}(k)$ by projective duality. So if we are given a set of halfplanes (lower and upper), every point which is covered $8k - 3$ times is covered at least $4k - 1$ times by either lower halfplanes or upper halfplanes. Thus we can color the lower and the upper halfplanes independently, using Theorem 3, and obtain $p_{\widetilde{\mathcal{HP}}}(k) \leq 8k - 3$. \square

3 Halfspaces in \mathbb{R}^3

In this section, we deal with the case where \mathcal{R} consists of all halfspaces in \mathbb{R}^3 . We call $\mathcal{HS} = (\mathbb{R}^3, \mathcal{R})$ the corresponding infinite range space and consider the value of $c_{\mathcal{HS}}(k)$.

The depth- k region in \mathbb{R}^3 is a convex polyhedron.

Lemma 4 Given a set of more than $4k$ points in \mathbb{R}^3 , every open halfspace not intersecting the depth- k polyhedron and whose bounding plane is tangent to the depth- k polyhedron contains at most $3k - 3$ points. The corresponding closed halfspace contains at least k points.

Proof The proof is similar to that of Lemma 1 in 2D. We consider open and closed halfspaces tangent to the depth- k polyhedron. By the definition of the depth- k polyhedron, any such closed halfspace contains at least k points. A halfspace is either tangent at a vertex, an edge, or a face of the polyhedron; if an open halfspace is tangent at a face, it contains at most $k - 1$ points; if an open halfspace is tangent at an edge (respectively, a vertex), it is contained in the union of two (respectively, three) open halfspaces tangent at a face of the polyhedron. \square

We consider halfspaces defined by planes tangent to the depth- k polyhedron. Each normal vector to one of these planes corresponds to precisely one halfspace and defines one point on the unit sphere \mathbb{S}^2 . We define R_p for every $p \in S$ as the set of

points on the unit sphere corresponding to normal vectors of tangent halfspaces containing p . We show that these sets are pseudo-disks, and provide an upper bound on $c_{\mathcal{H}_S}(k)$ by coloring their intersection graph.

Lemma 5 *Let p and q be two points of S outside the depth- k polyhedron. Then,*

1. R_p and R_q are connected subsets of \mathbb{S}^2 .
2. The boundaries of R_p and R_q intersect at most twice.

Proof The first property follows directly from the convexity of the depth- k polyhedron. Given a point p outside the depth- k region, any convex combination of the normal vectors of all planes tangent to the polyhedron and incident to p define a halfspace containing p .

To prove that the boundaries of R_p and R_q intersect at most twice, we look at all planes tangent to the polyhedron and incident to both p and q . These map to points that are on the boundary of both R_p and R_q . As p and q are distinct, they define a line. Through this line, there exist at most two planes tangent to the depth- k polyhedron. \square

The proof of the next theorem uses the following definition and lemma from graph theory [9].

Definition 2 A simple graph $G = (V, E)$ is called k -degenerate for some positive integer k if every (vertex-induced) subgraph of G has a vertex of degree at most k .

We use the standard notion of chromatic number $\chi(G)$ of a graph G defined as the number of colors needed to color the vertices of G so that no edge is monochromatic.

Lemma 6 *Let $G = (V, E)$ be a k -degenerate graph. Then $\chi(G) \leq k + 1$.*

Proof Proceed by induction on $n = |V|$. Let $v \in V$ be a vertex of degree at most k . By the induction hypothesis, the graph $G \setminus v$ (obtained by removing v and all of its incident edges from G) is $(k + 1)$ -colorable. Since v has at most k neighbors, there is always a free color that can be assigned to v which is distinct from the colors of its neighbors. \square

Theorem 5 $c_{\mathcal{H}_S}(k) = O(k)$. *That is, we can color any set of points in \mathbb{R}^3 with $O(k)$ colors so that any halfspace containing h points is $\min\{h, k\}$ -colorful.*

Proof Let $\mathcal{A} = \{R_x | x \in S, \text{ outside or on the surface of the depth-}k \text{ polyhedron}\}$. By Lemma 5, we know that \mathcal{A} is a set of pseudo-disks on \mathbb{S}^2 . Let \mathcal{A}' be the corresponding open pseudo-disks. By Lemma 4, we also know that every point belongs to at most $3k - 2$ regions of \mathcal{A}' .

By a lemma of Sharir [16], the complexity of a planar arrangement of bounding curves of n pseudo-disks such that any point belongs to the interior of at most i pseudo-disks is $O(ni)$. We observe that this lemma also applies to pseudo-disks on a sphere. This can be seen by projecting the arrangement of pseudo-disks R_p on the unit sphere onto the plane. We can use, for instance, a stereographic projection from

the point x on the top of the sphere to a horizontal plane below it, thereby sending x to infinity. The point x is contained in at most $3k - 2$ pseudo-disks, which are reverted in the projection. We therefore have an arrangement of pseudo-disks boundaries in the plane such that each point is contained in at most $6k - 4$ pseudo-disks.

Thus by this lemma, the complexity of the bounding curves in \mathcal{A}' is $O(nk)$. Now consider the intersection graph of \mathcal{A}' . This graph is $O(k)$ -degenerate. To see this, consider a pair of intersecting regions $r_1, r_2 \in \mathcal{A}'$. Either the boundaries of r_1 and r_2 intersect at some vertex, in which case we know that there are $O(nk)$ such vertices, or one of the regions, say r_1 , is contained in r_2 . However, since every point belongs to at most $O(k)$ regions, every region is contained in at most $O(k)$ other regions, hence the total number of such pairs of regions is at most $O(nk)$. Thus the number of edges in the intersection graph is $O(nk)$. This is true for every induced subgraph, and hence by Lemma 6, this graph is $O(k)$ -colorable. A similar observation has already been made by Chan [4].

Now it remains to prove that every halfspace of size $h \geq k$ is k -colorful. Given such a halfspace, there are two possibilities. Either the halfspace does not intersect the depth- k polyhedron. Then it is strictly contained in an open halfspace tangent to the polyhedron, and thus every point it contains has a unique color. Or the halfspace intersects the polyhedron and thus contains a closed halfspace tangent to it. Then it contains at least k different colors. \square

Corollary 2 $c_{\widetilde{\mathcal{HS}}}(k) = O(k)$. That is, we can color any set of halfspaces in \mathbb{R}^3 with $O(k)$ colors so that any point in the intersection of at least k of them is covered by k different colors.

Proof Given a set of halfspaces in \mathbb{R}^3 , we first consider the subset of lower halfspaces. Each such lower halfspace is the set of point below a given plane. By projective duality, a set of planes can be mapped to a set of points so that a point is above k planes if and only if in the projective dual a plane is above k points. In other words, by applying Theorem 5 in the dual, we derive a coloring for the lower halfspaces in the primal, which is correct as the inclusion relation (above-below) is preserved by projective duality: every lower halfspace containing at least k points in the primal is a point covered by k halfspaces in the projective dual.

We apply the same technique to color the upper halfspaces, using different colors. Now if a point is in ℓ lower halfspaces and u upper halfspaces, and $\ell + u \geq k$, then it is covered by halfspaces with k different colors: either $\ell < k$ and $u < k$, and we have $\ell + u$ different colors, or one of them is at least k , and we have k different colors. \square

4 Disks and Pseudo-disks

In this section we consider the case where the ranges in \mathcal{R} are disks or pseudo-disks. We denote by $\mathcal{D} = (\mathbb{R}^2, \mathcal{R})$ the range space for disks and by $\widetilde{\mathcal{D}}$ its dual, where the ground set is the set of disks and the ranges are the subsets of all disks having a common point. Similarly, we use the notation \mathcal{P} and $\widetilde{\mathcal{P}}$ for the range spaces defined by pseudo-disks.

The proof given above for halfspaces in \mathbb{R}^3 can be used to prove that $c_{\mathcal{D}}(k) = O(k)$. This is seen by a standard lifting transformation of disks and points in the plane to points and halfspaces in \mathbb{R}^3 that preserves the incidence relations.

Corollary 3 $c_{\mathcal{D}}(k) = O(k)$.

Proof Given a set S of points in \mathbb{R}^2 , we proceed by lifting the points onto the parabola of equation $z = x^2 + y^2$ in \mathbb{R}^3 . It is known that any disk in \mathbb{R}^2 is the projection onto the plane xy of the intersection between the parabola and a lower halfspace in \mathbb{R}^3 . The result follows by applying Theorem 5 to this set. \square

In the following, we give a bound for the value of $c_{\tilde{\mathcal{P}}}(k)$, where $\tilde{\mathcal{P}}$ is the dual range space defined by pseudo-disks. Similar to the proof of Theorem 5, we analyze the degeneracy of a graph induced by a finite set of regions.

Definition 3 Let S be a finite family of simple closed Jordan regions in the plane. We denote by $G_k(S)$ the graph on S whose edges are all pairs $r, s \in S$ such that there exists a point p that belongs to $r \cap s$ and at most k other regions of S .

Lemma 7 Let S be a family of pseudo-disks. Then $G_k(S)$ is $O(k)$ -degenerate, and hence the chromatic number of $G_k(S)$ is at most $O(k)$.

We aim to show that the number of edges in any (vertex-induced) subgraph of G with m vertices is at most $O(km)$, and therefore, the average degree in any induced subgraph is at most $O(k)$. Thus, there must exist a vertex of degree at most $O(k)$ in any induced subgraph. Hence, $G_k(S)$ is $O(k)$ -degenerate, and by Lemma 6 it is $O(k)$ -colorable as asserted. We need the following lemmas:

Lemma 8 There exists a constant c such that for any set S of n pseudo-disks, $G_0(S)$ has at most cn edges.

Proof See, for instance, the proof of Lemma 5.1 in [17]. \square

Lemma 9 Let S be a family of n pseudo-disks, and let $G = (S, E)$ be a subgraph of the intersection graph of S . (Thus E is a subset of the set of all pairs of regions from S that have a nonempty intersection.) For each edge $e = (a, b) \in E$, choose a point $p_e \in a \cap b$ that belongs to the intersection of a and b . Let X be the set of all pairs (e, r) such that $e \in E$ and $r \in S \setminus \{a, b\}$ contains the point p_e chosen for the edge e . Suppose that $|E| > 4cn$, where c is the constant from Lemma 8. Then $|X| \geq \frac{|E|^2}{4cn}$.

Proof The proof proceeds in two steps. In the first step, we prove the following bootstrapping inequality:

$$|X| \geq |E| - cn.$$

In the second step we use a random sampling argument similar to the one used for the Crossing Lemma (see, e.g., [1]).

The proof of the first step proceeds by induction on $|E| - cn$. For the case

$$|E| - cn \leq 0,$$

the claim is trivial. Assume that the claim holds for some nonnegative integer k (namely, for $|E|$ and n satisfying $|E| - cn = k$). Suppose that $|E| - cn = k + 1$. Since $|E| > cn$, Lemma 8 implies that there must exist a region $r \in S$ and an edge $e \in E$ which generates at least one configuration $(e, r) \in X$ (namely, that point p_e belongs to r , for otherwise X is empty, meaning that there is no edge of $G_k(S)$ for any $k > 0$; thus the graph is a subgraph of $G_0(S)$ and the number of edges in E by Lemma 8 is at most cn). After removing e from E we are left with $|E| - 1$ edges, n regions, and a set X' of configurations, where $|X| \geq |X'| + 1$. We have $|E| - 1 - cn = k$, so we can apply the induction hypothesis to obtain $|X'| \geq |E| - 1 - cn$. Thus $|X| \geq |X'| + 1 \geq |E| - cn$. This completes the proof of the first step.

Let X denote the set of configurations, as above. We take a random sample S' of the regions in S by choosing each region independently with some fixed probability p (to be determined later on). Let E' denote the subset of edges in E , all of whose defining regions are in S' . Let $n' = |S'|$; $m' = |E'|$, and let $X' \subset X$ denote the subset of configurations in X all of whose defining regions a, b , and r are in S' . By the above bootstrapping inequality, we have $|X'| \geq m' - cn'$. Note that $|X'|$, m' , and n' are random variables, so the above inequality holds for their expectations as well. Hence, using linearity of expectation, $\mathbf{E}[|X'|] \geq \mathbf{E}[m'] - c \mathbf{E}[n']$. It is easily seen that $\mathbf{E}[n'] = pn$. We have $\mathbf{E}[m'] = p^2|E|$ and $\mathbf{E}[|X'|] = p^3|X|$. Indeed, the probability that a given edge $e \in E$ belongs to E' is the probability that the two regions defining e are chosen in S' , which is p^2 for any fixed $e \in E$. Similarly, the probability that a configuration of a region $r \in S$ that contains a point p_e is counted in X' is p^3 . Substituting these values in the above inequality, we get $p^3|X| \geq p^2|E| - cpn$, or $|X| \geq \frac{|E|}{p} - \frac{cn}{p^2}$. This inequality holds for any $0 < p \leq 1$, and we choose $p = 2cn/|E|$ (by assumption, $p \leq 1$) to obtain $|X| \geq \frac{|E|^2}{4cn}$. This completes the proof of the lemma. \square

Proof of Lemma 7 Let X denote the set of configurations as above when E is the set of edges of $G_k(S)$ and for each edge $e \in E$, p_e is the point witnessing that $e \in E$ (that is, p_e is a point that belongs to the regions defining e and at most k other regions of S). By Lemma 9 we have:

$$|X| \geq \frac{|E|^2}{4cn}.$$

On the other hand, note that by the definition of $G_k(S)$, any point p_e can belong to at most k regions of S , so obviously

$$|X| \leq k|E|.$$

Combining the two bounds we have: $|E| \leq 4ckn$. Thus the sum of degrees of vertices in the graph $G_k(S)$ is at most $8ckn$, so the average degree is at most $8ck$. Thus there always exists a vertex with degree at most $8ck$, hence $G_k(S)$ is $8ck$ -degenerate. This completes the proof of the lemma. \square

Theorem 6 $c_{\tilde{\mathcal{P}}}(k) = O(k)$.

Proof We know by Lemma 7 that there exists a constant c such that $G_k(S)$ is ck -degenerate. We show that we can color the pseudo-disks in S with $ck + 1$ color so that for any point p with depth $d(p)$, the set of disks E_p containing p is $\min\{d(p), k\}$ -colorful. We use $ck + 1$ colors to color pseudo-disks inductively. The proof is by induction on $|S| = n$. Let $r \in S$ be a region whose degree in $G_k(S)$ is at most ck . By Lemma 7, there exists such a region. The induction hypothesis is that $S \setminus \{r\}$ admits a valid coloring. To complete the inductive step, we must assign a color to r so that the new coloring is still valid. Note that by the inductive hypothesis, points that belong to r and at least k other regions are already contained in some k regions (in $S \setminus \{r\}$), all colors of which are distinct. Hence, the color of r will not affect the validity for those points. We may only run into trouble for those points $p \in r$ that are contained in at most i (for $i \leq k - 1$) other regions. However, note that any region containing p is a neighbor of r in $G_k(S)$ by definition. Note also that by the induction hypothesis, all regions containing such a point p get distinct colors. Moreover, since the number of neighbors of r in $G_k(S)$ is at most ck , we can color r with a color distinct from all its neighbors in $G_k(S)$. Thus for any point in r that belongs to exactly i (for $i \leq k - 1$) other regions, all regions covering this point including r will have distinct color. This completes the inductive step and hence the proof of the theorem. \square

Remark 1 For the special case of real disks, it can be shown that the constant in Lemma 8 is $c = 3$ (we omit the details here). Thus by Lemma 7, the graph $G_k(S)$ is $24k$ -degenerate. Hence in the special case of real disks, we have that $c_{\tilde{\mathcal{D}}}(k) \leq 24k + 1$.

For the version in the primal range space in which we color points rather than regions, we can also prove the following:

Theorem 7 $c_{\mathcal{P}}(k) = O(k)$.

Proof The proof is very similar to the proof of Theorem 6 and uses the same ingredients. The analogue of Lemma 8 is provided in [18]. \square

Acknowledgements The authors wish to thank János Pach for helpful discussions about the problems studied in this paper and the anonymous referee for his insightful comments. We also thank the participants to the algorithmic lunches at the ULB Computer Science Department.

References

1. Aigner, M., Ziegler, G.M.: Proofs from The Book. Springer, Berlin (1998)
2. Aloupis, G.: Geometric measures of data depth. In: DIMACS Series in Discrete Mathematics and Theoretical Computer Science, pp. 147–158. Am. Math. Soc., Providence (2006)
3. Buchsbaum, A., Efrat, A., Jain, S., Venkatasubramanian, S., Yi, K.: Restricted strip covering and the sensor cover problem. In: ACM–SIAM Symposium on Discrete Algorithms (SODA'07) (2007)
4. Chan, T.M.: Low-dimensional linear programming with violations. SIAM J. Comput. **34**(4), 879–893 (2005)
5. Chazelle, B.: The Discrepancy Method: Randomness and Complexity. Cambridge University Press, Cambridge (2000)

6. Chen, X., Pach, J., Szegedy, M., Tardos, G.: Delaunay graphs of point sets in the plane with respect to axis-parallel rectangles. Manuscript (2006)
7. Haussler, D., Welzl, E.: Epsilon-nets and simplex range queries. *Discrete Comput. Geom.* **2**, 127–151 (1987)
8. Kříž, I., Nešetřil, J.: Chromatic number of Hasse diagrams, eyebrows and dimension. *Order*, **8**(1), 41–48 (1991)
9. Lick, D.R., White, A.T.: k -degenerate graphs. *Can. J. Math.* **12**, 1082–1096 (1970)
10. Mani, P., Pach, J.: Decomposition problems for multiple coverings with unit balls. Manuscript (1986)
11. Pach, J.: Decomposition of multiple packing and covering. In: 2. Kolloq. über Diskrete Geom., pp. 169–178. Inst. Math. Univ. Salzburg (1980)
12. Pach, J.: Personal communication (2007)
13. Pach, J., Tardos, G.: Personal communication (2006)
14. Pach, J., Tóth, G.: Decomposition of multiple coverings into many parts. In: Proc. of the 23rd ACM Symposium on Computational Geometry, pp. 133–137 (2007)
15. Pach, J., Tardos, G., Tóth, G.: Indecomposable coverings. In: The China–Japan Joint Conference on Discrete Geometry, Combinatorics and Graph Theory (CJCDGCGT 2005). Lecture Notes in Computer Science, pp. 135–148. Springer, Berlin (2007)
16. Sharir, M.: On k -sets in arrangement of curves and surfaces. *Discrete Comput. Geom.* **6**, 593–613 (1991)
17. Smorodinsky, S.: On the chromatic number of some geometric hypergraphs. *SIAM J. Discrete Math.* (to appear)
18. Smorodinsky, S., Sharir, M.: Selecting points that are heavily covered by pseudo-circles, spheres or rectangles. *Comb. Probab. Comput.* **13**(3), 389–411 (2004)
19. Tucker, A.: Coloring a family of circular arcs. *SIAM J. Appl. Math.* **229**(3), 493–502 (1975)
20. Tukey, J.: Mathematics and the picturing of data. In: Proceedings of the International Congress of Mathematicians, vol. 2, pp. 523–531 (1975)