Coloring Geometric Range Spaces

Greg Aloupis[1], Jean Cardinal[1], Sébastien Collette[⋆1], Stefan Langerman[⋆⋆1], and Shakhar Smorodinsky[⋆⋆⋆2]

[1] Université Libre de Bruxelles, CP212, Bld. du Triomphe, 1050 Brussels, Belgium.
Partially supported by the Communauté française de Belgique - ARC.
E-mail: {greg.aloupis,jcardin,secollet,slanger}@ulb.ac.be
[2] Institute of Mathematics, Hebrew University, Givat-Ram, Jerusalem 91904, Israel
E-mail: shakhar@cims.nyu.edu

Abstract. Given a set of points in $\mathbb{R}^2$ or $\mathbb{R}^3$, we aim to color them such that every region of a certain family (for instance disks) containing at least a certain number of points contains points of many different colors. Using k colors, it is not always possible to ensure that every region containing k points contains all k colors. Thus, we introduce two relaxations: either we allow the number of colors to increase to $c(k)$, or we require that the number of points in each region increases to $p(k)$. We give upper bounds on $c(k)$ and $p(k)$ for halfspaces, disks, and pseudo-disks. We also consider the dual question, where we want to color regions instead of points. This is related to previous results of Pach, Tardos and Tóth on decompositions of coverings.

1 Introduction

In this contribution, we are interested in coloring 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 Section 2 we answer this question on the affirmative.

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 this is true as well. We ask similar questions for other types of regions such as disks and pseudo-disks

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 $\mathbb{R}^3$,

⋆ Chargé de Recherches du FRS-FNRS.
⋆⋆ Chercheur Qualifié du FRS-FNRS.
⋆⋆⋆ http:// www.cims.nyu.edu/∼shakhar/

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?*

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. Here, we consider finite restrictions of infinite geometric range spaces of the form $\mathcal{S} = (\mathbb{R}^d, \mathcal{R})$ 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 $\widetilde{\mathcal{S}}$, of the form $\widetilde{\mathcal{S}} = (\mathcal{R}, \{r(p) : p \in \mathbb{R}^d\})$, 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 $(S, \{r(p) \cap S : p \in \mathbb{R}^d\})$, where $S \subset \mathcal{R}$ 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 k-*colorful* if it contains at least k elements of distinct color. We are interested in the following two 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 monotone non-decreasing functions. The goal of this paper is to provide upper bounds on $c_{\mathcal{S}}(k)$, $p_{\mathcal{S}}(k)$, $c_{\widetilde{\mathcal{S}}}(k)$, and $p_{\widetilde{\mathcal{S}}}(k)$ for various families of regions.

Previous results The functions defined above are related to two previously studied problems. The first one is the decomposition of *f-fold coverings* in the plane: given a covering of the plane 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 [6]. It is similar to deciding whether $p_{\widetilde{\mathcal{S}}}(2) \le f$ for the dual range space $\widetilde{\mathcal{S}}$ 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 but the proof does not directly yield a bound for $p_{\widetilde{\mathcal{S}}}(2)$. For $\mathcal{T}$ the range space defined by translates of a centrally symmetric convex polygon, Pach and Tóth [10] recently proved that $p_{\mathcal{T}}(k) = O(k^2)$ and $p_{\widetilde{\mathcal{T}}}(k) = O(k^2)$. 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 $\mathcal{D}$), Pach, Tardos, and Tóth [9] 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. In the same paper, a similar result is obtained for $p_{\widetilde{\mathcal{A}}}(2)$ where $\mathcal{A}$ is the range space induced by the family of either strips or axis-aligned rectangles. The fact that $p_{\widetilde{\mathcal{S}}}(2)$ is unbounded implies that for every $k > 2$, $p_{\widetilde{\mathcal{S}}}(k)$ 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 $k \le c(k)$.

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 polychromatic, that is, 2-colorful [12]. One of the main results of that contribution is that any dual range space induced by a finite set of pseudo-disks admits a $O(1)$-coloring that makes all ranges 2-colorful. Hence, for the family of pseudo-disks $\mathcal{P}$, $c_{\widetilde{\mathcal{P}}}(2) = O(1)$. A recent result of Chen, Pach, Szegedy and Tardos ([4], Thm. 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. This implies that $c_{\mathcal{A}}(k)$ and $p_{\mathcal{A}}(k)$ are unbounded, where $\mathcal{A}$ is the range space induced on $\mathbb{R}^2$ by the set of all axis-aligned rectangles.

Furthermore, Pach and Tardos [8] 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. Thus, $c_{\widetilde{\mathcal{A}}}(2) = \infty$, implying $c_{\widetilde{\mathcal{A}}}(k) = \infty$.

Our results In Section 2, we consider the range space $\mathcal{H} = (\mathbb{R}^2, \mathcal{R})$, where $\mathcal{R}$ is the set of all halfplanes. We prove that $c_{\mathcal{H}}(k) \le 3k - 2$, and $p_{\mathcal{H}}(k) \le 4k - 1$. 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.

In Section 3, we consider the range space $\mathcal{L} = (\mathbb{R}^3, \mathcal{R})$, where $\mathcal{R}$ is the set of all *lower halfspaces*. We prove that $c_{\mathcal{L}}(k) = O(k)$; and that $c_{\widetilde{\mathcal{L}}}(k) = O(k)$.

We provide a number of results on range spaces defined by disks and pseudo-disks in Section 4. For the range space $\mathcal{D}$ defined by disks, we prove that $c_{\mathcal{D}}(k) = O(k)$ by mapping disks in $\mathbb{R}^2$ to lower halfspaces in $\mathbb{R}^3$ and using the result of Section 3. For a dual range space $\widetilde{\mathcal{P}}$ defined by pseudo-disks we prove that $c_{\widetilde{\mathcal{P}}}(k) = O(k)$. Since halfplanes are a special case of pseudo-disks, we directly have $c_{\widetilde{\mathcal{H}}}(k) = O(k)$. We also show that $c_{\mathcal{P}}(k) = O(k)$, with 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 $p_{\mathcal{L}}(k) \ge p_{\mathcal{D}}(k)$. We also prove that $p_{\widetilde{\mathcal{L}}}(k) = p_{\mathcal{L}}(k)$: coloring lower halfspaces is equivalent in the projective dual to coloring points with respect to lower halfspaces.

All the proofs are constructive, and polynomial-time algorithms can easily be derived from them. The results are summarized in the following table, where the symbol $\star$ indicates new results; and the symbol ∞ indicates a function unbounded in terms of k.

$\mathcal{S}$	$c_{\mathcal{S}}(k)$	$p_{\mathcal{S}}(k)$	$c_{\widetilde{\mathcal{S}}}(k)$	$p_{\widetilde{\mathcal{S}}}(k)$
halfplanes	$\leq 3k-2$ (Thm. 1)$\star$	$\leq 4k-1$ (Thm. 2)$\star$	$O(k)$ (Thm. 4)$\star$	$\leq 8k-3$ (Cor. 1)$\star$
lower halfspaces in $\mathbb{R}^3$	$O(k)$ (Thm. 3)$\star$	∞ (Implied by disks)	$O(k)$ (Cor. 2)$\star$	∞ (Implied by disks)
translates of a cent. sym. convex polygon	$O(k)$ (Thm. 5)$\star$	$O(k^2)$ [10]	$O(k)$ (Thm. 4)$\star$	$O(k^2)$ [10]
axis-aligned rectangles	∞ [4]	∞ [4]	∞ [8]	∞ [9]
disks	$O(k)$ (Cor. 3, Thm. 3)$\star$	∞ (open disks [9])	$\leq 24k+1$ (Rem. 1)$\star$	
pseudo-disks	$O(k)$ (Thm. 5)$\star$	∞ (open disks [9])	$O(k)$ (Thm. 4)$\star$	

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. For more on the relation between these partitioning problems and sensor networks, see the paper of Buchsbaum *et al.* [2].

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 $\mathcal{H} = (\mathbb{R}^2, \mathcal{R})$ 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 (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 2 the notion of Tukey depth is used.

Definition 1 ([15]). *Given a set S of points in $\mathbb{R}^d$, the* Tukey depth *of a point p (not necessarily in the set) is the maximum integer t with the property that every halfspace containing p contains at least t points of S.*

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. 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 and the bounding line of which is tangent to 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 such that its bounding line ℓ is tangent to the depth-k polygon, and 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; their 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. □

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 arc 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 same set of arcs but open (in particular, degenerate arcs that consisted of only one point are removed).

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 which have their bounding lines passing 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 $k-1$. □

Theorem 1. $c_{\mathcal{H}}(k) \leq 3k-2$. *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 [14] it was proved that every set of arcs on the unit circle has a proper coloring with $m+M$ colors, where m (resp. M) is the minimum (resp. 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, meaning that it is strictly contained in an open halfplane Π' which has its boundary line 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 and thus 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. □

We now consider the depth-$2k$ region. As described in the preceding, 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 Pach noticed [7], a $2k$-covering of the unit circle with arcs of length at most π is decomposable into k disjoint coverings (by repeatedly removing a minimal covering of the unit circle). Thus we can assign one color to all arcs within each covering, so that each point on the circle is covered by k colors. □

Theorem 2. $p_{\mathcal{H}}(k) \leq 4k-1$. *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' with its boundary tangent to the depth-$2k$ region. By construction, Π' must be k-colorful and therefore so must Π. □

Corollary 1. $p_{\widetilde{\mathcal{H}}}(k) \leq 8k-3$. *That is, we can color any set of halfplanes with* k *colors such 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, then $p_{\widetilde{\mathcal{H}}}(k) = p_{\mathcal{H}}(k)$ 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 2 and obtain: $p_{\widetilde{\mathcal{H}}}(k) \leq 8k-3$. □

3 Lower Halfspaces in $\mathbb{R}^3$

Here, we deal with the case where $\mathcal{R}$ consists of all *lower halfspaces* in $\mathbb{R}^3$. We call $\mathcal{L} = (\mathbb{R}^3, \mathcal{R})$ the corresponding infinite range space and consider the value of $c_{\mathcal{L}}(k)$. The depth-k region in $\mathbb{R}^3$ is bounded by 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 which has a bounding plane 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 $\mathbb{R}^2$. We consider open and closed halfspaces tangent to the depth-k polyhedron and note that any tangent closed halfspace contains at least k points otherwise a point of the depth-k polyhedron has depth less than k. 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 (a vertex resp.) it is contained in the union of two (three resp.) open halfspaces tangent at a face of the polyhedron. □

In what follows, we consider lower halfspaces defined by planes tangent to the depth-k polyhedron. Each normal vector to one of these planes corresponds to precisely one lower halfspace and defines one point on the unit sphere. We map the points from the unit sphere onto the xy plane so that every lower halfspace corresponds to a single point in $\mathbb{R}^2$. This representation is used in the remainder of the section.

Lemma 5. *Let R_x denote the set of points in $\mathbb{R}^2$ corresponding to lower half-spaces tangent to the depth-k polyhedron and containing $x \in S$. Let p and q be two points of S outside the depth-k polyhedron. Then,*

1. *R_x is a connected subset of $\mathbb{R}^2$.*
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 x outside the depth-k region, any convex combination of the normal vectors of all planes tangent to the polyhedron and incident to x define a halfspace containing x.

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 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. □

The proof of the next theorem uses the following definition and lemma [5]. We use the standard notion of chromatic number $\chi(G)$ of a graph G, defined as the minimum number of colors needed to color the vertices so that no edge is monochromatic.

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.*

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 color that can be assigned to v, and that is distinct from the colors of its neighbors. □

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

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

By a lemma of Sharir [11], the complexity of an arrangement of the set of bounding curves of n pseudo-disks such that any point belongs to the interior of at most i of the pseudo-disks is $O(ni)$. Thus 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 $3k-2$ regions, every region is contained in at most $3k-3$ 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 was made by Chan [3].

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, meaning that 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, meaning that it contains at least k different colors. □

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

Proof. Given a set of halfspaces in $\mathbb{R}^3$, we consider their bounding planes. By projective duality, a set of planes can be mapped to a set of points, such 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 3 in the dual, we derive a coloring for the 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 dual. □

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 notations $\mathcal{P}$ and $\widetilde{\mathcal{P}}$ for the range spaces defined by pseudo-disks.

The proof given above for lower 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 3 to this set. □

In the following, we give a bound for the value of $c_{\widetilde{\mathcal{P}}}(k)$, where $\widetilde{\mathcal{P}}$ is the dual range space defined by pseudo-disks. Similar to the proof of Theorem 3, 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 $\mathbb{R}^2$. We denote by $G_k(S)$ the graph on S where the 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 [12]. □

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 non-empty 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 [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 positive 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, for which all defining regions are in S'. Let $n' = |S'|; m' = |E'|$, and let $X' \subset X$ denote the subset of configurations in X for which all the 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 $|X| \geq |E|^2/4cn$. This completes the proof of the lemma. □

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 $e \in E$ (i.e., 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: $|X| \geq |E|^2/4cn$.

On the other hand, note that by 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. □

Theorem 4. $c_{\widetilde{\mathcal{P}}}(k) = O(k)$

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 such 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 for which the 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. □

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 $c_{\widetilde{\mathcal{D}}}(k) \leq 24k + 1$.

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

Theorem 5. $c_{\mathcal{P}}(k) = O(k)$

Proof. The proof is very similar to the proof of Theorem 4 and uses the same ingredients. The analog of Lemma 8 is provided in [13]. □

Acknowledgements. The authors thank Janos Pach, as well as the participants to the algorithmic lunches at the ULB CS Department, for helpful discussions.

References

1. M. Aigner and G. M. Ziegler. *Proofs from the book*. Springer, 1998.
2. A. Buchsbaum, A. Efrat, S. Jain, S. Venkatasubramanian, and K. Yi. Restricted strip covering and the sensor cover problem. In *ACM-SIAM Symposium on Discrete Algorithms (SODA'07)*, 2007.
3. T. M. Chan. Low-dimensional linear programming with violations. *SIAM Journal on Computing*, 34(4):879–893, 2005.
4. X. Chen, J. Pach, M. Szegedy, and G. Tardos. Delaunay graphs of point sets in the plane with respect to axis-parallel rectangles. manuscript, 2006.
5. D. R. Lick and A. T. White. k-degenerate graphs. *Canadian Journal on Mathematics*, 12:1082–1096, 1970.
6. J. Pach. Decomposition of multiple packing and covering. In *2. Kolloq. über Diskrete Geom.*, pages 169–178. Inst. Math. Univ. Salzburg, 1980.
7. J. Pach. Personal communication. 2007.
8. J. Pach and G. Tardos. Personal communication. 2006.
9. J. Pach, G. Tardos, and G. Tóth. Indecomposable coverings. In *The China–Japan Joint Conference on Discrete Geometry, Combinatorics and Graph Theory (CJCDGCGT 2005)*, Lecture Notes in Computer Science, pages 135–148, 2007.
10. J. Pach and G. Tóth. Decomposition of multiple coverings into many parts. In *Proc. of the 23rd ACM Symposium on Computational Geometry*, pages 133–137, 2007.
11. M. Sharir. On k-sets in arrangement of curves and surfaces. *Discrete & Computational Geometry*, 6:593–613, 1991.
12. S. Smorodinsky. On the chromatic number of some geometric hypergraphs. *SIAM Journal on Discrete Mathematics*, to appear.
13. S. Smorodinsky and M. Sharir. Selecting points that are heavily covered by pseudo-circles, spheres or rectangles. *Combinatorics, Probability and Computing*, 13(3):389–411, 2004.
14. A. Tucker. Coloring a family of circular arcs. *SIAM Journal of Applied Mathematics*, 229(3):493–502, 1975.
15. J. Tukey. Mathematics and the picturing of data. In *Proceedings of the International Congress of Mathematicians*, volume 2, pages 523–531, 1975.