Discrete Comput Geom 34:47–70 (2005)
DOI: 10.1007/s00454-005-1162-6

Discrete & Computational
Geometry
© 2005 Springer Science+Business Media, Inc.

Conflict-Free Coloring of Points and Simple Regions in the Plane*

Sariel Har-Peled[1] and Shakhar Smorodinsky[2]

[1]Department of Computer Science, University of Illinois,
201 N. Goodwin Avenue, Urbana, IL 61801, USA
sariel@cs.uiuc.edu
http://www.uiuc.edu/˜sariel/

[2]Institute for Theoretische Informatik, ETH Zurich,
IFW CH-8092, Zurich, Switzerland
sshakhar@inf.ethz.ch
http://www.inf.ethz.ch/˜sshakhar

Abstract. We study conflict-free colorings, where the underlying set systems arise in geometry. Our main result is a general framework for conflict-free coloring of regions with low union complexity. A coloring of regions is conflict-free if for any covered point in the plane, there exists a region that covers it with a unique color (i.e., no other region covering this point has the same color). For example, we show that we can conflict-free color any family of n pseudo-discs with $O(\log n)$ colors.

1. Introduction

In this paper we study coloring problems related to frequency-assignment problems in cellular networks. In a geometric setting the problems are of the following two types:

CF-coloring of regions: Given a finite family $\mathcal{S}$ of n regions of some fixed type (such as discs, pseudo-discs, axis-parallel rectangles, etc.), what is the minimum integer k, such that one can assign a color to each region of $\mathcal{S}$, using a total of at most k colors, such that the resulting coloring has the following property: For each point $p \in \bigcup_{b\in\mathcal{S}} b$ there is at least one region $b \in \mathcal{S}$ that contains p in its interior, whose color is unique among all regions in $\mathcal{S}$ that contain p in their interior (in this case we say that p is being "served" by that color). We refer to such a coloring as a *conflict-free* coloring of $\mathcal{S}$ (CF-coloring in short).

* A preliminary version of this paper appeared in *Proc*. 19*th Annual ACM Symposium on Computational Geometery*, San Diego, CA, June 8–10, 2003, pages 114–123. Work on this paper by the first author was partially supported by NSF CAREER Award CCR-0132901. Work on this paper by the second author was done while he was a Ph.D. student under the supervision of Prof. Micha Sharir.

CF-coloring of a range space: A given set P of n points in $\mathbb{R}^d$ and a set $\mathcal{R}$ of ranges (for example, the set of all discs in the plane) define a so-called *range space* $(P, \mathcal{R})$. Given such a range space, what is the minimum integer k, such that one can color the points of P by k colors, so that for any $r \in \mathcal{R}$ with $P \cap r \neq \emptyset$, there is at least one point $q \in P \cap r$ that is assigned a unique color among all colors assigned to points of $P \cap r$ (in this case we say that r is "served" by that color). We refer to such a coloring as a *conflict-free* coloring of $(P, \mathcal{R})$ (CF-coloring in short).

The study of such problems, which was originated in [ELRS] and [Sm], was motivated by the problem of frequency-assignment in cellular networks. Specifically, cellular networks are heterogeneous networks with two different types of nodes: *base stations* (that act as servers) and *clients*. The base stations are interconnected by an external fixed backbone network. Clients are connected only to base stations; connections between clients and base stations are implemented by radio links. Fixed frequencies are assigned to base stations to enable links to clients. Clients, on the other hand, continuously scan frequencies in search of a base station with good reception. The fundamental problem of frequency-assignment in cellular networks is to assign frequencies to base stations so that every client, located within the receiving range of at least one station, can be served by some base station, in the sense that the client is located within the range of the station and no other station within its reception range has the same frequency. The goal is to minimize the number of assigned frequencies since the frequency spectrum is limited and costly.

Suppose we are given a set of n base stations, also referred to as *antennas*. Assume, for simplicity, that the area covered by a single antenna is given as a disc in the plane. Namely, the location of each antenna (base station) and its radius of transmission is fixed and is given (the transmission radii of the antennas are not necessarily equal). Even et al. [ELRS] have shown that one can find an assignment of frequencies to the antennas with a total of at most $O(\log n)$ frequencies such that each antenna (a base station) is assigned one of the frequencies and the resulting assignment is free of conflicts, in the preceding sense. Furthermore, it was shown that this bound is worst-case optimal. Thus, Even et al. have shown that any family of n discs in the plane has a CF-coloring with $O(\log n)$ colors and that this bound is tight in the worst case. Furthermore, such a coloring can be found in polynomial time. The approach used in [ELRS] relies strongly on the fact that the regions under consideration are discs.

In this paper we improve and extend the results of [ELRS] combining more involved probabilistic and geometric ideas. Our main result, which is delegated to Section 3.1, is a general probabilistic algorithm which CF-colors any set of n "simple" regions (not necessarily convex) whose union has "low" complexity, using a "small" number of colors. (The quoted terms are interrelated, in a manner stated more precisely in Section 3.1.) In particular, we show that if the regions under consideration have a union of near linear complexity, then they can be CF-colored using a polylogarithmic number of colors. This holds for pseudo-discs [KLPS], convex α-fat shapes [ES], and (α, β)-covered objects [Ef]. This provides the first non-trivial and near-optimal bounds for one of the problems that motivated the work of Even et al. [ELRS]. In practice, cellular antennas are directional, and the region of influence of an antenna is a circular sector with a central angle of 60°. Since such sectors are fat and convex, our results thus imply that those regions have a CF-coloring using a polylogarithmic number of colors.

In Section 3.2 we refine the results of Section 3.1, deriving better bounds for some special cases. We show that any set of n axis-parallel rectangles in the plane can be CF-colored with $O(\log^2 n)$ colors. We note that the assumption that the rectangles be axis-parallel cannot be removed, for otherwise one can construct a set $\mathcal{R}$ of n rectangles in which any CF-coloring of $\mathcal{R}$ needs n colors.

In Section 4 we study the problem of CF-coloring of range spaces, where the underlying ranges are axis-parallel rectangles in the plane, and show that any n points can be CF-colored with $O(\sqrt{n})$ colors with respect to axis-parallel rectangles (recall that in this new version we color the points of P with respect to a family of ranges, whereas in the preceding problem we colored the given regions). Using a different approach, we also obtain non-trivial upper bounds on the number of colors needed in any CF-coloring of a range space consisting of n points in $\mathbb{R}^d$ whose ranges are axis-parallel boxes. We also study the special case when all the given points form the regular $\sqrt{n} \times \sqrt{n}$ grid and show that in this case one can color the points with $O(\log n)$ colors and that this bound is worst-case optimal. This bound holds for any dimension. Namely, for any fixed d one can color the points of the d-dimensional regular $n^{1/d} \times \cdots \times n^{1/d}$ grid with $O(\log n)$ colors with respect to axis-parallel boxes. In fact, we show that the constant in the big "O" notation *does not* depend on the dimension d. We note that without the assumption that the rectangles are axis-parallel, the problem becomes uninteresting. Indeed, any planar set P of n points in general position (i.e., no three are collinear) needs n colors in any CF-coloring of P with respect to arbitrarily oriented rectangles.

Finally, in Section 5, we generalize the notion of CF-coloring of range spaces and of regions to what we call *k-CF-coloring*. That is, in the case of coloring a range space, we say that a range is "served" if there is a color that appears in the range (at least once and) at most k times, for some fix prescribed parameter k. A similar generalization of k-CF-coloring a set of regions is also studied. For example, we show that there is a range space consisting of n points for which any CF-coloring needs n colors but there exists a 2-CF-coloring with $O(\sqrt{n})$ colors (and a k-CF-coloring with $O(n^{1/k})$ colors for any fixed $k \geq 2$). We also show that any range space $(P, \mathcal{R})$ (not necessarily in geometry) with a finite VC-dimension c, can be k-CF-colored with $O(\log|P|)$ colors, for reasonably large k. This relaxation of the model is applicable in the wireless scenario since the real interference between conflicting antennas (i.e., antennas that are assigned the same frequency and overlap in their coverage area) is a function of the number of such antennas. This suggests that if for any given point, there is some frequency that is assigned to at most a "small" number of antennas that cover this point, then this point can still be served using that frequency because the interference between a small number of antennas is low. This feature is captured by the notion of k-CF-coloring.

2. Preliminaries

We briefly introduce some notations and tools used in this paper. In the following, P denotes a set of n points in $\mathbb{R}^d$, and $\mathcal{R}$ denotes a set of ranges (for example, the set of all discs in the plane). A *range space* S is a pair $(X, \mathcal{R})$, where X is a (finite or infinite) set and $\mathcal{R}$ is a (finite or infinite) family of subsets of X. If A is a subset of X then $\Pi_{\mathcal{R}}(A) = \{r \cap A : r \in \mathcal{R}\}$ is the *projection* of $\mathcal{R}$ on A. In this paper we focus on range spaces that arise naturally in combinatorial and computational geometry. One such example is the space $S = (\mathbb{R}^d, \mathcal{H})$, where $\mathcal{H}$ is the set of all half-spaces in $\mathbb{R}^d$. For a finite set of points P in $\mathbb{R}^d$ and a (finite or infinite) collection $\mathcal{R}$ of ranges, we abuse the notation slightly and refer to the pair $(P, \mathcal{R})$ as a range space, referring in fact to the range space $(P, \Pi_{\mathcal{R}}(P))$.

The "Delaunay" graph $G = G(P, \mathcal{R})$ is the graph whose vertex set is P and whose edges are all pairs (u, v) for which there exists a range $r \in \mathcal{R}$ such that $r \cap P = \{u, v\}$. We denote a range realizing an edge $(u, v) \in G$ by r_{uv}. When $\mathcal{R}$ is the set of all discs in the plane and P is a finite set of points with no four of them co-circular, the "Delaunay" graph of the range space $(P, \mathcal{R})$ coincides with the classical definition of the Delaunay triangulation of P.

A coloring $f : P \to \{1, \dots, k\}$ is a *conflict-free* coloring of $(P, \mathcal{R})$ (*CF-coloring* in short), if for any $r \in \mathcal{R}$, such that $P \cap r \neq \emptyset$, there exists a color i, for which there is a point $p \in P \cap r$, such that $f(p) = i$, and no other point of $P \cap r$ is assigned the color i. Any range r for which this property holds (regardless of whether the coloring is conflict free) is said to be *served* by the coloring. We refer to the minimum number of colors needed to CF-color $(P, \mathcal{R})$ as the *conflict-free chromatic number* of $(P, \mathcal{R})$ (or CF-chromatic number).

For a set $\mathcal{R}$ of ranges in $\mathbb{R}^d$, let $k_{\text{opt}}(n, \mathcal{R})$ denote the maximum number of colors needed for the given set $\mathcal{R}$, over all sets of n points in $\mathbb{R}^d$.

A range space $(P, \mathcal{R})$ is called *monotone* if for any $P_1 \subset P$ and for each $r \in \mathcal{R}$ with $|r \cap P_1| > 2$ there exists a range $r' \in \mathcal{R}$ such that $|r' \cap P_1| = 2$, and $r' \cap P_1 \subset r \cap P_1$. It is easy to verify that this property holds when $\mathcal{R}$ is the set of all axis-parallel rectangles in the plane.[1]

Even et al. have shown that the problem of CF-coloring a family S of n discs in the plane can be reduced to that of CF-coloring a range space $(P, \mathcal{R})$ where P is a set of n points in $\mathbb{R}^3$ and $\mathcal{R}$ is the set of all half-spaces.

A natural approach (used in [ELRS]) for CF-coloring of a monotone range space $(P, \mathcal{R})$ is to pick a large independent set L_1 in $G(P, \mathcal{R})$, color all the points of L_1 by a single color, and repeat this process on $(P \backslash L_1, \mathcal{R})$. We summarize this approach in the following algorithm:

Algorithm 1. CFcolor$(P, \mathcal{R})$: *CF-color a set P with respect to a set of ranges $\mathcal{R}$.*

$i \leftarrow 0$: *i denotes an unused color*
$P_1 \leftarrow P$
while $P_{i+1} \neq \emptyset$
  Increment: $i \leftarrow i + 1$
  Find an independent set $P_i' \subset P_i$ of $G(P_i, \mathcal{R})$:
  We elaborate subsequently on the implementation of this step.
  Color: $f(x) \leftarrow i, \forall x \in P_i'$
  Prune: $P_{i+1} \leftarrow P_i \backslash P_i'$
end while

---

[1] The interested reader might try to prove this property for the case where $\mathcal{R}$ is the set of all discs in the plane.

Let $L_i \subset P$ denote the set of points in P colored with i by Algorithm 1. We refer to L_i as the *ith layer* of $(P, \mathcal{R})$.

Lemma 2.1 [ELRS]. *The coloring of a monotone range space* $(P, \mathcal{R})$ *by* Algorithm 1 *is a valid CF-coloring of* $(P, \mathcal{R})$.

Proof. Consider a range $r \in \mathcal{R}$, such that $|P \cap r| \geq 2$. Let i be the maximal color assigned to points of P lying in r. Let $P_i \subset P$ be the set of input points at the beginning of the ith iteration, i.e., the set just before color i has been assigned. Note that $L_i \subset P_i$ and $L_i \cap r = P_i \cap r$ (since i is the maximal color in r). Clearly, if $|r \cap L_i| = 1$ then r is served and we are done.

Thus, we only have to consider the case $|r \cap L_i| > 1$. However, by the monotonicity property (applied to the subset P_i), it follows that there exists a range r' such that: (i) $|r' \cap P_i| = 2$, and (ii) $r' \cap P_i \subset r \cap P_i = r \cap L_i$.

This means that the two points of $r' \cap L_i$ form an edge in the graph $G(P_i, \mathcal{R})$. This however contradicts the fact that L_i is independent in $G(P_i, \mathcal{R})$, and thereby completes the proof of the lemma. □

To realize the usefulness of Lemma 2.1, consider the following result in [ELRS]: Let P be a set of n points and let $\mathcal{R}$ be the set of all discs in the plane. Then the chromatic number of $(P, \mathcal{R})$ is $O(\log n)$. The proof follows immediately from the fact that $(P, \mathcal{R})$ is monotone and Lemma 2.1, as $G(P, \mathcal{R})$ is just the Delaunay graph of P, which is planar (see e.g., [BKOS]), and as such it has an independent set of size at least $n/4$ (by the four colors theorem). It follows, that P has a decomposition into $O(\log n)$ layers and hence the chromatic number of $(P, \mathcal{R})$ is $O(\log n)$. (It was also shown in [ELRS] that there exists a set P of n points in the plane for which any CF-coloring of $(P, \mathcal{R})$ needs $\Omega(\log n)$ colors, and therefore this bound is worst case tight. Recently Pach and Tóth [PT] have shown that *any* set P of n points in the plane needs $\Omega(\log n)$ colors in any CF-coloring of $(P, \mathcal{R})$.)

We summarize this technique in the following lemma.

Lemma 2.2. *Let* $\mathcal{R}$ *be a set of ranges in* $\mathbb{R}^d$, *so that for any finite set* P, *the range space* $(P, \mathcal{R})$ *is monotone.*

(i) *If the Delaunay graph* $G(P, \mathcal{R})$ *contains an independent set of size at least* $\alpha|P|$, *for some fixed* $0 < \alpha < 1$, *then* $k_{\text{opt}}(n, \mathcal{R}) \leq \log n / \log(1/(1-\alpha))$.
(ii) *If* $G(P, \mathcal{R})$ *contains an independent set of size* $\Omega(|P|^{1-\epsilon})$, *for some fixed* $0 < \epsilon < 1$, *then* $k_{\text{opt}}(n, \mathcal{R}) = O(n^{\epsilon})$.

Proof. The assumption in part (i) of the lemma implies that in the ith iteration of Algorithm 1 we color at least $\alpha|P_i|$ points of P_i with the color i. This means that if we start with a set of n points, the number of iterations is at most $\log n / \log(1/(1-\alpha))$. Similarly, part (ii) of the lemma follows by observing that the number of iterations needed by Algorithm 1 is bounded by $O(n^{\epsilon})$. □

We need the following technical definition and lemma, for subsequent sections.

Definition 2.3. For a finite set V, a *k-uniform hypergraph H* on V is a pair of the form (V, E), where E is a set of subsets of V, such that each set in E is of size k (those are the *hyperedges* of H). The *degree* of a vertex $v \in V$ is the number of sets (i.e., hyperedges) of E that contain v.

A set $A \subseteq V$ is called an *independent set* if no hyperedge of E is contained in A.

Lemma 2.4.

(i) *Let G be a simple graph on n vertices with average degree δ. Then G contains an independent set of size $\Omega(n/\delta)$.*
(ii) *Let H be a k-uniform hypergraph with n vertices and average degree δ. Then H contains an independent set of size $\Omega(n/\delta^{1/(k-1)})$.*

Both facts are easy exercises in graph theory (see, e.g., [AS]).

3. CF-Coloring of Regions

In this section we consider the problem of CF-coloring of regions, and present one of the main results of this paper. We introduce a general approach that yields near-optimal bounds on the CF-chromatic number of any finite collection of regions with "low" (usually near-linear) union complexity. Our approach can also be applied to a general geometric range space (not necessarily monotone) whose Delaunay graph has "low" complexity.

3.1. *CF-Coloring of Regions with Low Union Complexity*

Let $\mathcal{R}$ be a family of regions in the plane, such that the complexity of the union of any n regions of $\mathcal{R}$ is at most $\mathcal{U}(n)$. In the following, we assume that $\mathcal{U}(n)$ is a near-linear function. This holds for pseudo-discs [KLPS] and (α, β)-covered objects [Ef]. See below for more precise statements of those bounds.

Definition 3.1. For a set $\mathcal{S}$ of n regions of $\mathcal{R}$, a subset $\widehat{S} \subseteq \mathcal{S}$ is *admissible* (with respect to $\mathcal{S}$) if any $p \in \bigcup \widehat{S}$ satisfies one of the following two conditions:

1. There is only one region of $\widehat{S}$ that covers p.
2. There exists $r \in \mathcal{S} \setminus \widehat{S}$, such that $p \in r$.

See Fig. 1.

Remark. Note, that an admissible set is also an independent set in the corresponding Delaunay graph $G = G(\mathcal{S}, E_{\mathcal{S}})$, where $E_{\mathcal{S}} = \{e \mid e \in \Phi_{\mathcal{S}}, |e| = 2\}$ and $\Phi_{\mathcal{S}} = \{\Phi_{\mathcal{S}}(p) \mid p \in \mathbb{R}^2\}$, $\Phi_{\mathcal{S}}(p) = \{r \mid r \in \mathcal{S}, p \in r\}$. Indeed, in the graph G every two regions r_1, r_2 that cover a common point, which is not covered by any other region, are connected by an edge in G. Thus, it cannot be that both r_1 and r_2 belong to an admissible set.

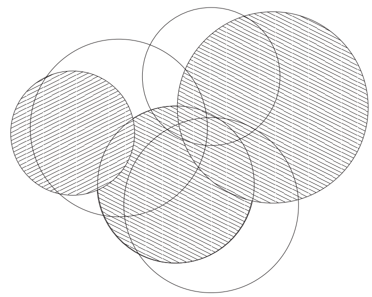

Fig. 1. A set $\mathcal{S}$ of discs and an admissible subset $\widehat{S}$ (depicted shaded).

Interestingly, there may exist an independent set in $G(\mathcal{S}, E_{\mathcal{S}})$ which is not admissible; see Fig. 2.[2] As a matter of fact, the range space $(\mathcal{S}, \Phi_{\mathcal{S}})$ is not necessarily monotone, and thus coloring the range space $(\mathcal{S}, \Phi_{\mathcal{S}})$ using Algorithm 1 is not necessarily valid, as testified by the example shown in Fig. 2.

Assume that we are given an algorithm **A** that computes, for any set of regions $\mathcal{S}$, a non-empty admissible set $\mathbf{A}(\mathcal{S})$. We can now use the algorithm **A** to CF-color the given regions: (i) Compute an admissible set $\widehat{S} = \mathbf{A}(\mathcal{S})$, and assign to all the regions in $\widehat{S}$ the color 1. (ii) Color the remaining regions in $\mathcal{S}\setminus\widehat{S}$ recursively, where in the ith stage we assign the color i to the regions in the admissible set. We denote the resulting coloring by $C_{\mathbf{A}}(\mathcal{S})$.

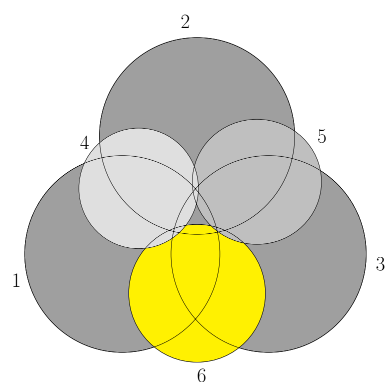

Fig. 2. The range space depicted is (V, E_V), where V is the set of circles $\{1, 2, 3, 4, 5, 6\}$ and every face of the arrangement $\mathcal{A}(V)$ induces a subset in Φ_V, which is the subset of circles of V covering this face. Clearly, the range $\{1, 2, 3\} \in \Phi_V$ but the ranges $\{1, 2\}, \{2, 3\}, \{3, 1\}$ are not in Φ_V since there is no face in the arrangement that is covered only by those pairs of circles. Thus, the range space (V, Φ_V) is not monotone. In particular, the set $\{1, 2, 3\}$ is independent in the graph $G(V, E_V)$ while it is not admissible, and as such the coloring depicted, which is clearly illegal (in the conflict-free sense), is one that Algorithm 1 might output.

[2] We are indebted to Shai Zaban for suggesting this example.

Lemma 3.2. *Given a set of regions $\mathcal{S}$, the coloring $C_{\mathbf{A}}(\mathcal{S})$ is a valid CF-coloring of $\mathcal{S}$.*

Proof. The proof is similar to that of Lemma 2.1. □

Remark. As a matter of fact, the coloring $C_{\mathbf{A}}(\mathcal{S})$ has the stronger property that every point p in $\bigcup \mathcal{S}$ is served by the maximal color among the colors of regions that contain p.

Lemma 3.3. *Let $\mathcal{R}$ be a set of n regions and let $\mathcal{U}(m)$ denote the maximum complexity of the union of any m regions of $\mathcal{R}$. Let $\mathcal{A}(\mathcal{R})$ denote the arrangement of the boundary curves of the regions in $\mathcal{R}$. Then the number of faces of the arrangement $\mathcal{A}(\mathcal{R})$ that are contained inside at most k regions of $\mathcal{R}$ (denoted by $F_{\leq k}(\mathcal{R})$) is bounded by $O\left(k^2\mathcal{U}(n/k) + n\right)$.*

Proof. We may assume that the regions of $\mathcal{R}$ are in general position, in the sense that no three distinct boundaries pass through a common point. This can be enforced by a slight perturbation of the curves, which does not decrease $F_{\leq k}(\mathcal{R})$. Let $S_{\leq k}(\mathcal{R})$ be the set of vertices of the arrangement $\mathcal{A}(\mathcal{R})$ that lie in the interior of at most k regions of $\mathcal{R}$. By the probabilistic analysis of Clarkson and Shor [CS], we have $|S_{\leq k}(\mathcal{R})| = O\left(k^2\mathcal{U}(n/k)\right)$. We charge a face f contained in at most k regions to its lowest vertex, if ∂f has vertices. Thus, the only faces unaccountable for by this charging scheme are the faces that have no vertices on their boundary. However, it is easy to check that the number of such faces is only $O(n)$, as we can charge such a face to the region of $\mathcal{R}$ that forms its outer boundary. Thus $F_{\leq k}(\mathcal{R}) = O(S_{\leq k}(\mathcal{R}) + n) = O(k^2\mathcal{U}(n/k) + n)$. □

In what follows, we assume that $\mathcal{U}(m) \geq m$ for any m and that $\mathcal{U}(m)/m$ is a monotonically non-decreasing function, so the bound in Lemma 3.3 is in fact $O\left(k^2\mathcal{U}(n/k)\right)$ in this case.

Lemma 3.4. *Let $\mathcal{R}$ be a set of n regions in the plane, so that the boundaries of any pair of them intersect in a constant number of points, and let $\mathcal{U}(m)$ denote the maximum complexity of the union of any m regions of $\mathcal{R}$. Then there exists an admissible set $\widehat{S} \subseteq \mathcal{R}$ with respect to $\mathcal{R}$, such that $|\widehat{S}| = \Omega\left(n^2/\mathcal{U}(n)\right)$.*

Proof. Let $\mathcal{A} = \mathcal{A}(\mathcal{R})$ be the arrangement of the regions of $\mathcal{R}$. Place an arbitrary point inside each face of the arrangement $\mathcal{A}$ and let P denote the resulting point set.

Let χ be a random coloring of the regions of $\mathcal{R}$ by two colors, black and white, where each region is colored independently by choosing black or white with equal probabilities. A point $p \in P$ is said to be *unsafe* if all the regions of $\mathcal{R}$ that contain p are colored black. Let P_{U} be the set of unsafe points of P. Let $\mathcal{R}_{\mathrm{B}}$ be the set of all regions of $\mathcal{R}$ which are colored black by χ. We construct a graph G over $\mathcal{R}_{\mathrm{B}}$, connecting two regions $r, r' \in \mathcal{R}_{\mathrm{B}}$ by an edge if there is an unsafe point $p \in P_{\mathrm{U}}$ that is contained inside both r and r'.

Let $e(G)$ and $v(G)$ denote, respectively, the number of edges and vertices in G. We claim that, with constant probability, $v(G) \geq n/3$ and $e(G) = O(\mathcal{U}(n))$.

Clearly, the condition $|\mathcal{R}_{\mathrm{B}}| = v(G) \geq n/3$ holds with high probability (which tends to 1 when n increases) by the Chernoff inequality (see [AS]). As for the second claim, for a point $p \in P$, let $\mathbf{d}(p)$ denote the number of regions of $\mathcal{R}$ that contain it. Clearly, the probability that p is unsafe is $1/2^{\mathbf{d}(p)}$. If p is unsafe, there are $\binom{\mathbf{d}(p)}{2}$ pairs of regions of $\mathcal{R}_{\mathrm{B}}$ whose intersections contain p, so p induces $\binom{\mathbf{d}(p)}{2}$ edges in G. Let X_p be the random variable having value 0 if p is safe, and $\binom{\mathbf{d}(p)}{2}$ if p is unsafe. Clearly, $e(G) \le \sum_{p \in P} X_p$. Thus, using linearity of expectation and Lemma 3.3, we have

$$\begin{aligned}\mathbf{E}[e(G)] &\le \sum_{p \in P} \mathbf{E}[X_p] = \sum_{\substack{p \in P \\ \mathbf{d}(p)>1}} \frac{\binom{\mathbf{d}(p)}{2}}{2^{\mathbf{d}(p)}} = O\left(\sum_{i=2}^{n} \sum_{\substack{p \in P \\ \mathbf{d}(p)=i}} \frac{i^2}{2^i}\right) = O\left(\sum_{i=2}^{n} i^2 \mathcal{U}(n/i) \cdot \frac{i^2}{2^i}\right) \\ &= O\left(\sum_{i=2}^{n} \frac{i^4}{2^i} \mathcal{U}(n)\right) = O\left(\mathcal{U}(n)\right).\end{aligned}$$

Thus, by the Markov inequality, it follows that there is a constant c, such that

$$\mathbf{Pr}[e(G) \ge c \cdot \mathcal{U}(n)] \le \tfrac{1}{4}.$$

It follows that, with constant probability, G has at least $n/3$ vertices, and its average degree is at most $6c \cdot \mathcal{U}(n)/n$. Thus, by Lemma 2.4(i), G contains an independent set of size $\Omega(n^2/\mathcal{U}(n))$. Let $\mathcal{R}'$ be this independent set. It is easy to verify that $\mathcal{R}'$ is admissible with respect to $\mathcal{R}$. Indeed, let f be a face of $\mathcal{A}(\mathcal{R})$ that is contained in at least two regions $r_1, r_2 \in \mathcal{R}'$, and let p be its representing point. Then p must be safe, so p, and thus f, is contained also in a white region, which clearly does not belong to $\mathcal{R}'$. □

Lemma 3.5. *The admissible set guaranteed by* Lemma 3.4 *can be computed in randomized expected $O(\mathcal{U}(n) \log n)$ time.*

Proof. Note that the proof of Lemma 3.4 is constructive. Assume a model of computation as in [SA] in which computing the intersection points of any pair of regions in $\mathcal{R}$, and a few similar operations, can be performed in constant time.

To construct G, first we compute a random coloring χ of the regions of $\mathcal{R}$ by black and white. Let w be the number of white regions. Next, randomly permute the regions of $\mathcal{R}$, so that all the white regions (according to χ) appear before the black regions of $\mathcal{R}$. This can be done by randomly permuting the white regions and randomly permuting the black regions, independently, and concatenating the two permutations. Let π denote this permutation. Note that π is a random permutation chosen uniformly from the set of all permutations of the elements of $\mathcal{R}$. Let r_i denote the ith region of $\mathcal{R}$ according to π.

We need to compute all the unsafe points (i.e., faces which are covered only by black regions) in $\mathcal{A}(\mathcal{R})$. This can be facilitated by computing $\mathcal{C}_i$, which is the vertical decomposition of the complement of the union of the first i regions of $\mathcal{R}$, for $i = 1, \ldots, w$. Formally, $\mathcal{C}_i$ is the vertical decomposition of $\mathbb{R}^2 \backslash (\cup_{k=1}^{i} r_k)$, for $i = 1, \ldots, w$. We construct $\mathcal{C}_w$, by using randomized incremental construction. At the ith step, we maintain $\mathcal{C}_i$, which is computed from $\mathcal{C}_{i-1}$ by inserting into it the region r_i. This involves removing vertical trapezoids of $\mathcal{C}_{i-1}$ that are covered by r_i, splitting trapezoids that intersect the boundary of r_i, and merging trapezoids that are adjacent and have common ceiling and floor curves. We stop as soon as we computed $\mathcal{C}_w$. See [SA] and [Mu] for further details on randomized incremental constructions.

For every trapezoid $\Delta \in \mathcal{C}_i$ the algorithm also maintains its "conflict-list" which is the list of all regions of $\mathcal{R}$ intersecting the interior of Δ, for $i = 1, \ldots, w$. Using those conflict-lists, we compute the arrangement $\mathcal{A}_\Delta$ of the black regions that intersect Δ, for every trapezoid $\Delta \in \mathcal{C}_w$. Next, we perform a traversal of this arrangement, and for every face of $\mathcal{A}_\Delta$, we generate the relevant edges in G.

Now that the graph G is available, computing the admissible set in G can be done by a greedy independent set algorithm, which picks the vertex v of lowest degree in G, adds it to the output set, and removes v and its neighbors from G and recurses on the remaining subgraph. One can verify that this algorithm computes an independent set in G of size $\Omega\left((v(G))^2/e(G)\right)$, where $v(G)$ and $e(G)$ are the number of vertices and edges of G, respectively. Thus, yielding the required admissible set.

We next bound the expected running time of this algorithm. It is easily seen that the number of vertical trapezoids in $\mathcal{C}_i$ is $O(\mathcal{U}(i))$, and by the Clarkson–Shor probabilistic analysis [Mu, Lemma 5.5.1], the expected average length of a conflict-list of $\mathcal{C}_i$ is $O(n/i)$. Using backward analysis (see, e.g., [SA] and [Mu]), the probability of a trapezoid of $\mathcal{C}_i$ to be created in the ith iteration is $O(1/i)$. Putting everything together, we have that the expected time to construct $\mathcal{C}_w$ is

$$O\left(\sum_{i=1}^{w} \frac{\mathcal{U}(i)}{i} \cdot \frac{n}{i}\right) = O\left(\sum_{i=1}^{n} \frac{\mathcal{U}(n)}{n} \cdot \frac{n}{i}\right) = O\left(\sum_{i=1}^{n} \frac{\mathcal{U}(n)}{i}\right) = O\left(\mathcal{U}(n) \log n\right),$$

since we assumed that $\mathcal{U}(i)/i$ is a monotone non-decreasing function.

Similarly, the expected time to compute the arrangement of the black regions inside each vertical trapezoid $\Delta \in \mathcal{C}_w$ takes $O(l_\Delta^2)$ time, where l_Δ is the size of the conflict-list of Δ. Using the Clarkson–Shor analysis [Mu, Lemma 5.5.1] again, it follows that the total expected time to compute this arrangement is $O(\mathcal{U}(n))$. Thus, it is now straightforward to construct the graph G from it. Again, computing G takes

$$E\left[\sum_{\Delta \in \mathcal{C}_w} l_\Delta^4\right] = E_w\left[O\left(\mathcal{U}(w)\left(\frac{n}{w}\right)^4\right)\right] = O\left(\mathcal{U}(n)\right)$$

time, using the Clarkson–Shor analysis [Mu, Lemma 5.5.1] for the last and final time in this proof, and observing that $w \geq n/3$ with high probability.

The greedy independent set algorithm can be implemented in linear time in the size of the graph, and as such the running time of the algorithm is dominated by the other stages.

Note that if the admissible set generated by the algorithm is too small, then the algorithm is run again until it succeeds. □

We now present several applications of Lemmas 3.4 and 3.5.

Definition 3.6 [KLPS]. A family $\mathcal{R}$ of Jordan regions in the plane is called a family of *pseudo-discs* if the boundaries of each pair of them intersect at most twice.

Theorem 3.7. *Let $\mathcal{R}$ be a family of n pseudo-discs. Then $\mathcal{R}$ admits a CF-coloring with $O(\log n)$ colors. Such a coloring can be constructed in randomized expected $O(n \log n)$ time.*

Proof. The complexity of the union of any m regions of $\mathcal{R}$ is $O(m)$ (see [KLPS]). Plugging this fact into Lemma 3.4, we have that $\mathcal{R}$ contains an admissible set $\widehat{S}$ with respect to $\mathcal{R}$ of size $\Omega(n)$. Applying Lemma 3.2, and arguing as in the proof of Lemma 2.2, we have that $\mathcal{R}$ admits a CF-coloring with $O(\log n)$ colors. □

Definition 3.8 [Ef]. A planar object c is (α, β)*-covered* if the following holds: (i) c is simply connected, and (ii) for any point $p \in \partial c$ we can place a triangle Δ fully inside c, such that p is a vertex of Δ, each angle of Δ is at least α, and the length of each edge of Δ is at least β times the diameter of c.

Theorem 3.9. *Let C be a collection of n (α, β)-covered regions in the plane, of finite description complexity, such that the boundaries of each pair of regions of C intersect in at most s points. Then C has a CF-coloring using $O(\beta_{s+2}(n)\log^3 n \log\log n)$ colors, where $\beta_{s+2}(n) = \lambda_{s+2}(n)/n$ and where $\lambda_{s+2}(n)$ is the maximum length of an s-order Davenport–Schinzel sequence from n symbols, see, e.g.,* [SA]. *This coloring can be computed in randomized expected $O(n \log^{O(1)} n)$ time, in an appropriate model of computation.*

Proof. In this case, $\mathcal{U}(n) = O(\lambda_{s+2}(n)\log^2(n)\log\log n)$ by the result of Efrat [Ef]. Thus, by Lemma 3.4, C has an admissible set of size

$$\Omega\left(\frac{n^2}{\mathcal{U}(n)}\right) = \Omega\left(\frac{n^2}{\lambda_{s+2}(n)\log^2 n \log\log n}\right) = \Omega\left(\frac{n}{\beta_{s+2}(n)\log^2 n \log\log n}\right).$$

Applying the algorithm described in Lemma 3.2, and arguing as in Lemma 2.2, it follows that we have a CF-coloring of C using

$$O(\beta_{s+2}(n)\log^3 n \log\log n)$$

colors. □

3.2. *CF-Coloring of Simple Geometric Regions in the Plane*

3.2.1. *CF-coloring of axis-parallel rectangles*

Lemma 3.10. *Let $\mathcal{R}$ be a set of n axis-parallel rectangles, all intersecting the y-axis. Then there is a CF-coloring of $\mathcal{R}$ with $O(\log n)$ colors, which can be constructed in randomized expected $O(n \log n)$ time.*

Proof. It is easy to verify that the complexity of the union of m such rectangles is $O(m)$. Hence, the result follows immediately from Lemmas 3.4 and 3.2. □

Theorem 3.11. *Let $\mathcal{R}$ be a set of n axis-parallel rectangles. Then there is a CF-coloring of $\mathcal{R}$ using $O(\log^2 n)$ colors.*

Proof. Let ℓ be a vertical line, such that at most $n/2$ rectangles of $\mathcal{R}$ lie fully to the left of ℓ, and at most $n/2$ rectangles of $\mathcal{R}$ lie fully to its right. Let $\mathcal{R}_0, \mathcal{R}_1, \mathcal{R}_2$ denote, respectively, the sets of rectangles crossed by ℓ, lying fully to its left, and lying fully to its right. By Lemma 3.10, we can CF-color the set $\mathcal{R}_0$ with $O(\log n)$ colors. We color recursively $\mathcal{R}_1$ and $\mathcal{R}_2$, using the same set of colors in both subproblems, but keeping this set disjoint from the set used to color $\mathcal{R}_0$. This gives rise to a coloring of $\mathcal{R}$ with a total of $O(\log^2 n)$ colors, which is easily seen to be a CF-coloring. □

Again, the proof is constructive, and leads to an $O(n \log^2 n)$-randomized expected time algorithm for computing the coloring.

3.2.2. *CF-coloring of half-planes*

Theorem 3.12. *There exists a collection $\mathcal{H}$ of n half-planes, for which $\Omega(\log n)$ colors are needed in any CF-coloring of $\mathcal{H}$.*

Proof. We use a standard dual transformation that maps a line l to a point l^* and a point p to a line p^*, such that p lies above (resp., below) l if and only if the line p^* lies above (resp., below) the point l^*. It is easily verified that any CF-coloring of a set $\mathcal{H} = \{l_1^+, \ldots, l_n^+\}$ of n positive half-planes is equivalent to that of a CF-coloring of a range space $(P, \mathcal{R})$, where $P = \{l_1^*, \ldots, l_n^*\}$ is the set of dual points of the boundary lines of the half-planes in $\mathcal{H}$, and $\mathcal{R}$ is the set of all negative half-planes. Thus, it suffices to show that for any integer n, there exists a set P of n points in the plane such that any CF-coloring of P with respect to negative half-planes needs at least $\Omega(\log n)$ colors. Such a construction can be obtained by placing n points on the parabola $y = x^2$; see, e.g., [ELRS]. □

Remark. It easily follows from the results of Even et al. [ELRS] that $O(\log n)$ colors always suffice for CF-coloring n half-planes.

Similar constructions show that there exists a collection $\mathcal{R}$ of n axis-parallel rectangles for which $\Omega(\log n)$ colors are needed in any CF-coloring of $\mathcal{R}$. This still leaves a logarithmic gap with the upper bound of Theorem 3.11.

In the context of range spaces, similar constructions of a set of n points in the plane show that in any CF-coloring of the given points, $\Omega(\log n)$ colors are needed when the ranges are axis-parallel rectangles.

4. CF-Coloring of Range Spaces

In this section we consider the "dual" problem of CF-coloring of points with respect to regions rather than coloring regions with respect to points.

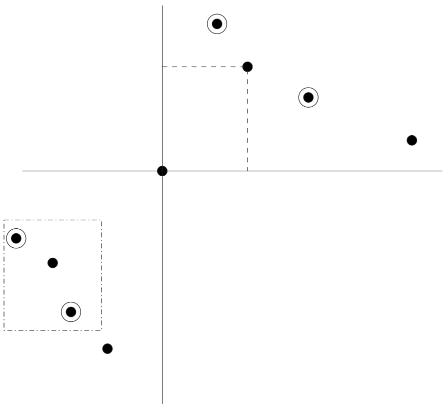

Fig. 3. A point p and the neighbors of p in two opposite quadrants in the graph $G(P, \mathcal{B}^2)$. The circled points form an independent set in this graph.

4.1. *Axis-Parallel Rectangles*

In this section we deal with the problem of CF-coloring of points in the plane, where the ranges are axis-parallel rectangles.

Theorem 4.1. *For the set $\mathcal{B}^2$ of all axis-parallel rectangles in the plane, we have $k_{\text{opt}}(n, \mathcal{B}^2) = O(\sqrt{n})$.*

Proof. Let P be a set of n points in the plane, and let $G = G(P, \mathcal{B}^2)$ denote the corresponding Delaunay graph. Note that the ranges that realize the edges of G can be taken to be those rectangles that have two points of P as opposite vertices and are otherwise disjoint from P. If there is a point $p \in P$ with degree $\geq 2\sqrt{n}$ in G, then there are two opposite quadrants around p that contain together at least $\sqrt{n}$ neighbors of p in $G(P, \mathcal{B}^2)$. See Fig. 3. Suppose, without loss of generality, that these are the upper-right and the lower-left quadrants. The neighbors of p in each of the quadrants form a monotone decreasing sequence. Choosing every other element in each sequence yields an independent set in G of size at least $\sqrt{n}/2$. Otherwise, all the points of p have degree $< 2\sqrt{n}$ in G. However, in this case, Lemma 2.4(i) implies that there is an independent set in G of size $\Omega(\sqrt{n})$. By Lemma 2.2(ii), $(P, \mathcal{B}^2)$ can be CF-colored using $O(\sqrt{n})$ colors. □

Remark. Noga Alon, Timothy Chan, János Pach, and Geza Tóth [PT] have independently noticed that the result of Theorem 4.1 can be slightly improved by a polylogarithmic factor, using more involved graph-theoretic arguments [AKS], [PT]. Their main observation is that the Delaunay graph $G(P, \mathcal{B}^2)$ has sparse neighborhoods. Namely, for any point p, the subgraph of G induced by the set N_p of the neighbors of p has size $O(|N_p|)$. The result in [AKS] implies that if a graph G has maximum degree δ and has "sparse neighborhoods" then G contains an independent set of size $\Omega(n((\log \delta)/\delta))$. Choosing $\delta = \sqrt{n \log n}$ we have: If G contains a point with degree more than δ, then by the above analysis G contains an independent set of size $\Omega(\delta)$. Otherwise, by the sparse neighborhood property of G we have that G contains an independent set of size $\Omega(n((\log \delta)/\delta)) = \Omega(\delta)$. A simple modification of the proof of Lemma 2.2(ii) implies that the number of layers into which P can be decomposed is $O(\sqrt{n}/\sqrt{\log n})$. By Lemma 2.1, $(P, \mathcal{B}^2)$ can be CF-colored using $O(\sqrt{n}/\sqrt{\log n})$ colors.

Substantially improving the result of Theorem 4.1 is the main open problem that we pose in this paper, as we currently have only a trivial lower bound of $\Omega(\log n)$.

Using a somewhat different approach, we next give an alternative proof of Theorem 4.1, which generalizes to higher dimensions.

Theorem 4.2. *Let $\mathcal{B}^d$ be the set of all axis-parallel boxes in $\mathbb{R}^d$. Then $k_{\rm opt}(n, \mathcal{B}^d) = O(n^{1-1/2^{d-1}})$.*

Note that for $d = 2$ we obtain the same bound as in Theorem 4.1.

Proof. Let P be a set of n points in $\mathbb{R}^d$, and denote the coordinates by $x_1, \dots, x_d$. Let P_1 be the ordered sequence of the points of P according to their x_1-coordinate. At the ith stage, for $i = 2, \dots, d$, let P_i be the longest monotone subsequence of P_{i-1}, according to their x_i-coordinates. By the Erdős–Szekeres theorem (see, e.g., [W]) there exists a monotone subsequence of P_{i-1} of length $\Omega(\sqrt{|P_{i-1}|})$.

Thus, P_d is a sequence of $\Omega(n^{1/2^{d-1}})$ points which is monotone in all coordinates (in each coordinate it can be either increasing or decreasing). It is easy to verify that if we pick every other point in this sequence, we obtain an independent set in $G(P, \mathcal{B}^d)$ of size $|P_d|/2 = \Omega(n^{1/2^{d-1}})$. We thus conclude, by Lemma 2.2(ii), that $k_{\rm opt}(n, \mathcal{B}^d) = O(n^{1-1/2^{d-1}})$. □

It is easy to construct the CF-coloring provided by Theorem 4.2 in time $O(n^{2-1/2^{d-1}} \log n)$: there are $O(n^{1-1/2^{d-1}})$ iterations, in each of which we compute $(d-1)$ times a longest monotone subsequence, which can be done in $O(n \log n)$ time.

In contrast to the rather weak bounds of Theorems 4.1 and 4.2, we next show that the special case where P is a grid admits a CF-coloring of (optimal) logarithmic size.

Definition 4.3. The grid $\mathcal{G}(n, d)$ is the Cartesian product $\{1, \dots, \lfloor n^{1/d} \rfloor\}^d$.

In the following, we use the fact that if two integer numbers have the same number of trailing zeros in their binary representation, then there must be a number between them that has a larger number of trailing zeros in its binary representation. Thus, we can use the number of trailing zeros in the binary representation as the color assigned to an integer, when coloring consecutive integers.

Lemma 4.4. *Let $\mathcal{I} = \mathcal{B}^1$ be the set of intervals on the real line and let* $\mathrm{cf}(i)$ *be the function defined on the positive integers and returning $j+1$ if 2^j is the largest power of 2 that divides i.*

Then the CF-chromatic number of $(\mathcal{G}(n,1),\mathcal{I})$ *is* $1+\lfloor\log n\rfloor$, *it is realized by* $\mathrm{cf}(\cdot)$, *and this bound is tight. Furthermore, for an interval* $I=[i,j]$, *the color that appears exactly once in* $I\cap\mathcal{G}(n,1)$ *is the largest number in the set* $\{\mathrm{cf}(i),\mathrm{cf}(i+1),\ldots,\mathrm{cf}(j)\}$.

Proof. We only prove the lower bound. The other claims can be easily verified. Let $f(\cdot)$ be any CF-coloring of $\mathcal{G}(n,1)=\{1,\ldots,n\}$, using the minimum number of colors. Let $h(m)$ be the minimum number of colors used by $f(\cdot)$ for coloring an interval of length m.

Consider the color appearing exactly once in the coloring $f(\cdot)$ of the interval $I=[1,n]$. Namely, there is an $i\in I$ such that $f(i)\neq f(j)$, for all $j\in I$, $j\neq i$. Let $I_l=[1,i-1]$ and $I_r=[i+1,n]$, and assume, without loss of generality, that $|I_r|\geq|I_l|$. Clearly, we have $h(n)=h(|I|)\geq 1+h(|I_r|)=h(\lceil(n-1)/2\rceil)+1$. By induction, it is now easy to prove that $h(n)\geq\lfloor\log n\rfloor+1$. □

Lemma 4.5. *The CF-chromatic number of* $(\mathcal{G}(n,d),\mathcal{B}^d)$ *is at most* $1+\lfloor\log n\rfloor$.

Proof. For $\mathbf{i}=(i_1,\ldots,i_d)\in\mathcal{G}(n,d)$, we define its color to be $f(\mathbf{i})=\sum_{j=1}^{d}\mathrm{cf}(i_j)-(d-1)$. Let R be any axis-parallel box, and let N_j be the set of integers in the projection of R onto the jth axis. Note that, for $j=1,\ldots,d$, $\mathrm{cf}(N_j)$ has a unique maximum in this range, by Lemma 4.4. Let i'_j be the index that realizes it. Clearly, $f(\mathbf{i}')$ is the maximum value achieved by $f(\cdot)$ on R, where $\mathbf{i}'=(i'_1,i'_2,\ldots,i'_d)$. Furthermore, no other point of $R\cap\mathcal{G}(n,d)$ realizes this value. Thus $f(\cdot)$ provides the required CF-coloring. To complete the proof, note that the value of $f(\cdot)$ is bounded from above by $d(1+\lfloor\log(\lfloor n^{1/d}\rfloor)\rfloor)-(d-1)\leq 1+\lfloor\log n\rfloor$. □

Lemma 4.5 is tight for $d=1$ and $d=\lfloor\log n\rfloor$. For other values of d, one can show a lower bound of $\lfloor\log n\rfloor-d$. To see this, consider any CF-coloring of $\mathcal{G}(n,d)$, and let p be the point with a unique color in the whole grid. Then there is a box that avoids p and contains almost half of the points of $\mathcal{G}(n,d)$. Analyzing carefully the number of points remaining in this box, and using induction, we obtain the asserted lower bound.

4.2. *Random Point Set Inside a Square*

In the following, let $\mathcal{U}$ denote the unit square in the plane. In this section we consider the CF-coloring of a point-set generated by picking points uniformly and independently out of $\mathcal{U}$. The ranges considered are axis-parallel rectangles.

Lemma 4.6. *Let* $P_1,P_2\subseteq\mathcal{U}$ *be two random point-sets of cardinality* m *each, assume that* P_1 *was CF-colored using* χ *colors, and let* n *be a parameter such that* $m\leq n$. *Then* $P_1\cup P_2$ *can be CF-colored using* $\chi+O(\log^3 n)$ *colors, with high probability.*

Proof. Clearly, if $m=O(\log^3 n)$, then the claim trivially holds. Otherwise, let $\epsilon=O((\log n)/m)$. By ϵ-net theory [HW], P_1 is an ϵ-net for rectangles inside the unit square

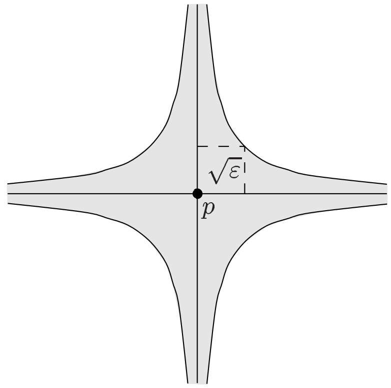

Fig. 4. The region $S(p, \epsilon)$.

under the measure of area, with high probability. Namely, any axis-parallel rectangle of area larger than ϵ, contains a point of P_1.

For a point $p \in \mathcal{U}$, let

$$S(p, \epsilon) = \{q \mid (\text{rect}(p, q)) \leq \epsilon\}$$

be the set of points that form rectangles of area at most ϵ with p, namely $q \in S(p, \epsilon)$ iff $|q_x - p_x| \cdot |q_y - p_y| \leq \epsilon$, see Fig. 4. Furthermore,

$$A = \text{area}\,(S(p, \epsilon)) \leq 4\left(\epsilon^2 + \int_{x=\sqrt{\epsilon}}^{1} \frac{\epsilon}{x}\, dx\right) = O\left(\epsilon \log \frac{1}{\epsilon}\right) = O\left(\frac{\log n}{m} \log \frac{m}{\log n}\right).$$

Clearly, $E\,[|S(p, \epsilon) \cap P_2|] = Am$, and by the Chernoff inequality,

$$\begin{aligned} \mathbf{Pr}[|S(p, \epsilon) \cap P_2| \geq Am(1 + \log n)] &\leq \left(\frac{e^{\log n}}{(1 + \log n)^{(1+\log n)}}\right)^{Am} \\ &\leq \left(\frac{e^{\log n}}{(1 + \log n)^{(1+\log n)}}\right)^{O(\log n)} \\ &\leq e^{-c \log^2 n \log\log n} \leq n^{-c \log n \log\log n}, \end{aligned}$$

where c is an appropriate constant.

Let G_2 be the graph defined over the points of P_2, connecting two points of P_2, if the diagonal rectangle they define has area smaller than ϵ. Consider a point $p \in P_2$, clearly, all its neighbors in G_2, must lie inside $S(p, \epsilon)$, and by the above discussion, with high probability, the number of neighbors of p in G_2 is bounded by

$$\nu = O(Am \log n) = O\left(\left(\epsilon \log \frac{1}{\epsilon}\right) m \log n\right) = O(\log^3 n).$$

In particular, G_2 can be colored using $\nu + 1$ colors. Let $f(\cdot)$ be the coloring of $P_1 \cup P_2$ resulting from coloring the points of P_1 by their given colors, and coloring P_2 by additional $O(\log^3 n)$ colors, as specified by the coloring of G_2.

We claim that $f(\cdot)$ is a CF-coloring of $P_1 \cup P_2$. Indeed, let $R \subseteq \mathcal{U}$ be an arbitrary axis-parallel rectangle. If $R \cap P_1 \neq \emptyset$, then we are done, because the given coloring of P_1 is conflict free. Furthermore, if area$(R) \geq \epsilon$, then it contains a point of P_1, as P_1 is an ϵ-net.

Thus, it must be that area$(R) \leq \epsilon$, $R \cap P_1 = \emptyset$, and $R \cap P_2 \neq \emptyset$. However, by construction of G_2, all the pairs of points of $R \cap P_2$ are connected in G_2, thus $f(\cdot)$ assigns all of them unique colors.

It follows that $f(\cdot)$ is a CF-coloring of $P_1 \cup P_2$ with high probability. □

Theorem 4.7. *Let P be a set of n points picked randomly and uniformly out of the unit square $\mathcal{U}$. Then, with high probability, for the range space formed by axis-parallel rectangles, the set P has a CF-coloring using $O(\log^4 n)$ colors.*

Proof. Order P in an arbitrary order, and let P_i be the first 2^i points of P. Now, repeatedly apply Lemma 4.6 to $P_i \backslash P_{i-1}$ and P_{i-1}, for $i = 1, \ldots, \lceil \log n \rceil$. □

Observe that the property of having an empty axis-parallel rectangle is uniquely defined by the ordering of the given points in each coordinate. It follows, that instead of picking points randomly in the unit square, we can just generate the points by picking a random permutation π of $1, \ldots, n$, and placing the ith point at $(i, \pi(i))$. One can modify the proof of Theorem 4.7 so that it also holds in this setting. This results in an identical result with a combinatorial proof instead of a geometric one.

5. Relaxing the Notion of CF-Coloring

In this section we generalize the notion of CF-coloring of a range space and show a relation between the problem of CF-coloring a range space and its VC-dimension. We also generalize the notion of CF-coloring of regions. To simplify the presentation we ignore, in this section, the issue of algorithmic construction of the coloring. Nevertheless, all upper bounds in this section are constructive, and can be easily computed in polynomial time.

5.1. *k-CF-Coloring of a Range-Space*

Definition 5.1 (k-CF-Coloring of a Range Space). Let $(P, \mathcal{R})$ be a range space in $\mathbb{R}^d$. A function $\chi : P \to \{1, \ldots, i\}$ is a *k-CF-coloring* of $(P, \mathcal{R})$ if for every $r \in \mathcal{R}$ with $r \cap P \neq \emptyset$ there exists a color j such that $1 \leq |\{p \in P \cap r | \chi(p) = j\}| \leq k$; that is, for every possible non-empty range r there exists at least one color j such that j appears (at least once and) at most k times among the colors assigned to points of $P \cap r$.

Let $k_{\text{opt}}(n, k, \mathcal{R})$ denote the minimum number of colors needed for a k-CF-coloring of $(P, \mathcal{R})$, maximized over all sets P of size n.

Note that a 1-CF-coloring of a range space is just a CF-coloring.

5.1.1. *CF-coloring of balls in three dimensions*

Lemma 5.2. *Let $\mathcal{R}$ be the set of balls in three dimensions. Then $k_{\mathrm{opt}}(n, \mathcal{R}) = n$. The same holds for the set $\mathcal{R}$ of half-spaces in $\mathbb{R}^d$, for $d > 3$.*

Proof. Take P to be a set of n points on the positive portion of the moment curve $\gamma = \{(t, t^2, t^3) | t \geq 0\}$ in $\mathbb{R}^3$. It is easy to verify that any pair of points $p, q \in P$ are connected in the Delaunay triangulation of P [Er], implying that there exists a ball whose intersection with P is $\{p, q\}$. Thus, all points must be colored using different colors.

The second claim follows by lifting P into the standard paraboloid in $\mathbb{R}^4$ by the map $(x, y, z) \mapsto (x, y, z, x^2 + y^2 + z^2)$. A ball in $\mathbb{R}^3$ is mapped to a half-space in $\mathbb{R}^4$ so that a point p lies in the ball if and only if its image lies in the half-space. It follows that n colors are necessary in any CF-coloring of the image of P. This clearly extends to any dimension $d \geq 4$. □

Theorem 5.3. *Let $\mathcal{R}$ be the set of all balls in $\mathbb{R}^3$. Then $k_{\mathrm{opt}}(n, k, \mathcal{R}) = O(n^{1/k})$, for any fixed constant $k \geq 1$.*

Proof. The proof technique is a generalization of the ideas introduced in Section 2. Indeed, let P be any set of n points in $\mathbb{R}^3$, and construct a $(k + 1)$-uniform hypergraph $\mathcal{H} = (P, \mathcal{E})$, where $\mathcal{E}$ is the collection of all subsets of P of size $k + 1$ that are realizable by a range in $\mathcal{R}$. By the Clarkson–Shor technique, it is easy to see that $|\mathcal{E}| = \mathcal{O}(n^2)$, where the constant of proportionality depends on k. Thus, the average degree of $\mathcal{H}$ is $O(n)$ and therefore, by Lemma 2.4(ii), there exists an independent set $P' \subset P$ of size $\Omega(n^{1-1/k})$. (Note that independence means that any ball that contains at least $k+1$ points of P', must also contain a point from $P \backslash P'$; this equivalence follows by an appropriate extension of the monotonicity property of balls.) We can color all points of P' by a single color, say 1, and iterate on $P \backslash P'$, similar to Algorithm 1. Thus, the total number of colors we use is $O(n^{1/k})$. It is easy to see (similar to Lemma 2.1) that this coloring is a valid k-CF-coloring of $(P, \mathcal{R})$. □

5.2. *k-CF-Coloring of Range Spaces with Finite VC-Dimension*

Definition 5.4. Let $S = (X, \mathcal{R})$ be a range space. The *Vapnik–Chervonenkis* dimension (or *VC-dimension*) of S, denoted by $\mathrm{VC}(S)$, is the maximal cardinality of a subset $P' \subset P$ such that $\{P' \cap r | r \in \mathcal{R}\} = 2^{P'}$ (such a subset is said to be *shattered*). If there are arbitrarily large shattered subsets in X, then $\mathrm{VC}(S)$ is defined to be ∞. See [AS] and [PA] for discussion of VC-dimension and its applications.

There are many range spaces with finite VC-dimension that arise naturally in combinatorial and computational geometry. One such example is the range space $S = (\mathbb{R}^d, \mathcal{H}_d)$, where $\mathcal{H}_d$ is the family of all (open) half-spaces in $\mathbb{R}^d$. Any set of $d + 1$ affinely independent points is shattered in this space, and, by Radon's theorem, no set of $d + 2$ points is shattered. Therefore $\mathrm{VC}(S) = d + 1$. As a matter of fact, all range spaces used in this paper have finite VC-dimension.

Since all the range spaces studied in this paper have finite VC-dimension, and since some of them can be CF-colored only with n colors, there is no direct relationship between a finite VC-dimension of a range space and the existence of a CF-coloring of that range space with a small number of colors. In this subsection we show that such a relationship does exist, if we consider k-CF-coloring with a reasonably large k.

We first introduce a general framework for k-CF-coloring of a range space $S = (X, \mathcal{R})$.

Definition 5.5. A subset $X' \subset X$ is *k-admissible* with respect to S if for any range $r \in \mathcal{R}$ with $|r \cap X'| > k$ we have $r \cap (X \backslash X') \neq \emptyset$.

Note that, assuming a monotonicity property of the ranges in $\mathcal{R}$ (i.e., if a subset S is realizable by a range, then its has a subset of size k which is realizable by some range in $\mathcal{R}$), a k-admissible set is simply an independent set in the hypergraph $(X, \mathcal{E})$, where $\mathcal{E}$ is the set of all hyperedges consisting of $k + 1$ elements of X that can be realized by a range in $\mathcal{R}$.

Assume that we are given an algorithm $\mathbf{A}$ that computes, for any range space $S = (X, \mathcal{R})$, a non-empty k-admissible set $X' = \mathbf{A}(S)$. We can now use the algorithm $\mathbf{A}$ to k-CF-color the given range space: (i) Compute an admissible set $X' = \mathbf{A}(S)$, and assign to all the elements in X' the color 1. (ii) Color the remaining elements in $X \backslash X'$ recursively, where in the ith stage we assign the color i to the points in the resulting k-admissible set. We denote the resulting coloring by $C_{\mathbf{A}}(S)$.

The proof of the following lemma is similar to that of Lemma 2.1, and is omitted.

Lemma 5.6. *Given a range space $S = (X, \mathcal{R})$, the coloring $C_{\mathbf{A}}(S)$ is a valid k-CF-coloring of S.*

Lemma 5.7. *Let $S = (X, \mathcal{R})$, with $|X| = n$, be a finite range space with VC-dimension d. For any $k \geq d$ there exists a k-admissible set $X' \subset X$ with respect to S of size $\Omega(n^{1-(d-1)/k})$.*

Proof. Any coloring of X is valid as far as the small ranges of $\mathcal{R}$ are concerned; namely, those are the ranges that contain at most k points. Thus, let $\mathcal{R}'$ be the set of ranges of $\mathcal{R}$ of size larger than k. By Sauer's lemma [Sa] we have that $|\mathcal{R}'| \leq |\mathcal{R}| \leq n^d$.

Next, we randomly color X by black and white, where an element is being colored in black with probability p, where p would be specified shortly. Let I be the set of points of X colored in black. If a range $r \in \mathcal{R}'$ is colored only in black, we remove one of the points of r from I. Let I' be the resulting set. Clearly, I' is a k-admissible set for $(X, \mathcal{R})$.

Furthermore, by linearity of expectation, the expected size of I' is at least

$$pn - \sum_{r \in \mathcal{R}'} p^{|r|} \geq pn - \sum_{r \in \mathcal{R}'} p^{k+1} \geq pn - p^{k+1} n^d.$$

Setting $p = ((k + 1)n^{d-1})^{-1/k}$, we have that the expected size of I' is at least $pn - p^{k+1}n^d = pn(1 - 1/(k + 1)) = \Omega(n^{1-(d-1)/k})$, as required. □

For the case of geometric range spaces, one might be able to get better bounds than the one guaranteed by Lemma 5.7. See Theorem 5.3 for such an example.

Theorem 5.8. *Let $S = (X, \mathcal{R})$, with $|X| = n$, be a finite range space with VC-dimension d. Then for $k \geq d \log n$ there exists a k-CF-coloring of S with $O(\log n)$ colors.*

Proof. By Lemma 5.7 the range space S contains a k-admissible set of size at least $n/2$. Plugging this fact to the algorithm suggested by Lemma 5.6 completes the proof of the theorem. □

As remarked above, Theorem 5.8 applies to all the range spaces studied in this paper. Note also that Lemma 5.7 gives us a tradeoff between the number of colors and the threshold size of the coloring. As such, the bound of Theorem 5.8 is just one of a family of such bounds implied by Lemma 5.7.

5.3. *k-CF-Coloring of Regions*

Definition 5.9 (k-CF-Coloring of Regions). Let $\mathcal{R}$ be a collection of regions in $\mathbb{R}^d$. A function $\chi : \mathcal{R} \to \{1, \ldots, i\}$ is a *k-CF-coloring* of $\mathcal{R}$ if for every point $p \in \bigcup \mathcal{R}$ there exists a color j such that $1 \leq |\{r \in \mathcal{R} | p \in r, \chi(r) = j\}| \leq k$; that is, for every possible point p in the union of $\mathcal{R}$ there exists at least one color j such that j appears (at least once and) at most k times among the colors assigned to the regions of $\mathcal{R}$ that contain p.

As above, we note that a 1-CF-coloring of a set of regions $\mathcal{R}$ is just a CF-coloring of $\mathcal{R}$.

Consider a CF-coloring of a set of balls in $\mathbb{R}^3$. Note that the union of a set of n balls can have $\Omega(n^2)$ complexity and one cannot apply the technique developed in Section 3.1 to obtain non-trivial bounds on the number of colors needed for a 1-CF-coloring of such a set of balls, or other regions with high union complexity. However, as we will show in this section, one can obtain non-trivial bounds on the number of colors needed for k-CF-coloring a set of regions in $\mathbb{R}^3$ with near-quadratic union complexity, for any $k \geq 2$. The approach that we use generalizes to any fixed dimension.

Let $\mathcal{R}$ be a family of regions in $\mathbb{R}^3$, such that the complexity of union of any n regions of $\mathcal{R}$ is at most $\mathcal{U}(n)$. In the following, we assume that $\mathcal{U}(n)$ is a monotone increasing function of n and that $\mathcal{U}(n) = \Omega(n^2)$. This holds for balls with $\mathcal{U}(n) = \Theta(n^2)$ (see, e.g., [SA]).

Definition 5.10. For a set $\mathcal{S}$ of n regions, a subset $\widehat{S} \subseteq \mathcal{S}$ is *k-admissible* with respect to $\mathcal{S}$ if any $p \in \bigcup \widehat{S}$ satisfies one of the following two conditions:

1. There are at most k regions of $\widehat{S}$ that cover p.
2. There exists $r \in \mathcal{S} \setminus \widehat{S}$, such that $p \in r$.

Assume that we are given an algorithm $\mathbf{A}$ that computes, for any set $\mathcal{S}$ of regions in a given family, a non-empty k-admissible set $\mathbf{A}(\mathcal{S})$ with respect to $\mathcal{S}$. We can then use the algorithm $\mathbf{A}$ for k-CF-coloring the given regions as follows: (i) Compute a k-admissible set $\widehat{S} = \mathbf{A}(\mathcal{S})$ with respect to $\mathcal{S}$, and assign to all the regions in $\widehat{S}$ the color 1. (ii) Color the remaining regions in $\mathcal{S}\setminus\widehat{S}$ recursively, using colors ≥ 2. We denote the resulting coloring by $C_{\mathbf{A}}(\mathcal{S})$.

Lemma 5.11. *Given a set of regions $\mathcal{S}$, the coloring $C_{\mathbf{A}}(\mathcal{S})$ is a valid k-CF-coloring of $\mathcal{S}$.*

The proof is similar to that of Lemmas 2.1 and 3.2. The following result extends Lemma 3.3 to three dimensions.

Lemma 5.12. *Let $\mathcal{R}$ be a set of n regions in $\mathbb{R}^3$ of constant description complexity and let $\mathcal{U}(m)$ denote the maximum complexity of the boundary of the union of any m regions of $\mathcal{R}$, with $\mathcal{U}(m) = \Omega(m^2)$. Then the number $F_{\leq i}(\mathcal{R})$ of three-dimensional cells of the arrangement $\mathcal{A}(\mathcal{R})$ that are contained in at most i regions of $\mathcal{R}$ is $O\left(i^3\mathcal{U}(n/i)\right)$.*

Proof. Let $S_{\leq i}(\mathcal{R})$ be the set of vertices of the arrangement $\mathcal{A}(\mathcal{R})$ (of the boundary surfaces of the regions in $\mathcal{R}$) that lie in the interior of at most i regions of $\mathcal{R}$. By the Clarkson–Shor technique [CS], we have $|S_{\leq i}(\mathcal{R})| = O\left(i^3\mathcal{U}(n/i)\right)$. We charge a cell contained in at most i regions to its lowest vertex, assuming it has a vertex. Thus, the only cells unaccountable for by this charging scheme are the cells that have no vertices on their boundary. However, it is easy to check that the number of such cells is bounded by $O(n^2)$. Thus

$$F_{\leq i}(\mathcal{R}) = O(S_{\leq i}(\mathcal{R}) + n^2) = O(i^3\mathcal{U}(n/i) + n^2) = O(i^3\mathcal{U}(n/i)),$$

by our assumptions on $\mathcal{U}(n)$. □

Lemma 5.13. *Let $\mathcal{R}$ be a set of n regions in $\mathbb{R}^3$, and let $\mathcal{U}(m)$ denote the maximum complexity of the union of any m regions of $\mathcal{R}$, such that $\mathcal{U}(m) = \Omega(m^2)$ and $\mathcal{U}(\cdot)$ is monotone increasing. Then there exists a k-admissible set $\widehat{S} \subseteq \mathcal{R}$ with respect to $\mathcal{R}$, such that $|\widehat{S}| = \Omega(n^{1+1/k}/\mathcal{U}(n)^{1/k})$.*

Proof. The proof follows closely the ideas of the proof of Lemma 3.4 with a slight twist. Let $\mathcal{A} = \mathcal{A}(\mathcal{R})$ be the arrangement of the (boundary surfaces of the) regions of $\mathcal{R}$. Place an arbitrary point inside each (three-dimensional) cell of the arrangement $\mathcal{A}$ and let P denote the resulting point set.

Let χ be a random coloring of the regions of $\mathcal{R}$, by two colors, black and white, where each region is colored independently by choosing black or white with equal probabilities. A point $p \in P$ is said to be *unsafe* if all the regions of $\mathcal{R}$ that contain p are colored black. Let P_{U} be the set of unsafe points of P. Let $\mathcal{R}_{\mathrm{B}}$ be the set of all regions of $\mathcal{R}$ which are colored black by χ. We construct a $(k+1)$-uniform hypergraph H over $\mathcal{R}_{\mathrm{B}}$, whose set of hyperedges consist of all $(k+1)$-tuples of regions $r_1, \ldots, r_{k+1} \in \mathcal{R}_{\mathrm{B}}$ for which there is an unsafe point $p \in P_{\mathrm{U}}$ in $\bigcap_{j=1}^{k+1} r_j$.

Let $e(H)$ and $v(H)$ denote, respectively, the number of hyperedges and vertices of H. We claim that, with constant probability, $v(H) \geq n/3$ and $e(H) = O(\mathcal{U}(n))$.

Clearly, the condition $|\mathcal{R}_{\mathrm{B}}| = v(H) \geq n/3$ holds with high probability by the Chernoff inequality (see, e.g., [AS]). Similar to the proof of Lemma 3.4, the probability that p is unsafe is $1/2^{\mathbf{d}(p)}$, where $\mathbf{d}(p)$ is the number of regions containing p. If p is unsafe, there are $\binom{\mathbf{d}(p)}{k+1}$ $(k+1)$-tuples of regions of $\mathcal{R}_{\mathrm{B}}$ whose intersection contains p, so p induces $\binom{\mathbf{d}(p)}{k+1}$ hyperedges in H. Let X_p be the random variable having value 0 if p is safe, and $\binom{\mathbf{d}(p)}{k+1}$ if p is unsafe. Clearly, $e(H) \leq \sum_{p \in P} X_p$. Thus, using linearity of expectation and Lemma 5.12, we have

$$\begin{aligned}\mathbf{E}[e(H)] \leq \sum_{p \in P} \mathbf{E}[X_p] &= \sum_{\substack{p \in P \\ \mathbf{d}(p) \geq k+1}} \frac{\binom{\mathbf{d}(p)}{k+1}}{2^{\mathbf{d}(p)}} = O\left(\sum_{i=k+1}^{n} \sum_{\substack{p \in P \\ depth(p)=i}} \frac{i^{k+1}}{2^i} \right) \\ &= O\left(\sum_{i=k+1}^{n} i^3 \mathcal{U}(n/i) \cdot \frac{i^{k+1}}{2^i} \right) = O\left(\sum_{i=k+1}^{n} \frac{i^{k+4}}{2^i} \mathcal{U}(n) \right) = O\left(\mathcal{U}(n)\right).\end{aligned}$$

Thus, by the Markov inequality, it follows that there is a constant c, such that

$$\mathbf{Pr}[e(H) \geq c \cdot \mathcal{U}(n)] \leq \tfrac{1}{4}.$$

It follows that, with constant probability, H has at least $n/3$ vertices, and its average degree is at most $(k+1)3c \cdot \mathcal{U}(n)/n$. Thus, by Lemma 2.4(ii), H contains an independent set of size

$$\Omega\left(\frac{n}{(\mathcal{U}(n)/n)^{1/k}} \right) = \Omega\left(\frac{n^{1+1/k}}{(\mathcal{U}(n))^{1/k}} \right).$$

It is easy to verify that any such independent set is k-admissible with respect to $\mathcal{R}$. This completes the proof of the lemma. □

Note that when $\mathcal{U}(n) = O(n^2)$ we have a k-admissible set of size $\Omega(n^{1-1/k})$.

Theorem 5.14. *Let $\mathcal{R}$ be a set of n balls in $\mathbb{R}^3$. For any $k \geq 2$, there exists a k-CF-coloring of $\mathcal{R}$ with a total of at most $O(n^{1/k})$ colors.*

Proof. By Lemma 5.13 there exists a k-admissible set $\mathcal{R}'$ with respect to $\mathcal{R}$ of size $\Omega(n^{1-1/k})$. Plugging this fact into the algorithm suggested by Lemma 5.11 completes the proof. □

Remark. A closer inspection of the analysis of the proof of Lemma 5.13 shows that the lemma generalizes to any dimension $d \geq 3$, provided that we assume that $\mathcal{U}(m) = \Omega(m^{d-1})$.

Theorem 5.15. *Let $\mathcal{R}$ be a set of n regions in $\mathbb{R}^d$ with the property that the complexity of the union of any m regions of $\mathcal{R}$ is at most $\mathcal{U}(m)$, where $\mathcal{U}(m) = \Omega(m^{d-1})$ and is monotone increasing. Then there exists a k-admissible set $\widehat{S} \subseteq \mathcal{R}$ with respect to $\mathcal{R}$, such that $|\widehat{S}| = \Omega(n^{1+1/k}/\mathcal{U}(n)^{1/k})$.*

Remark. The condition that $\mathcal{U}(m) = \Omega(n^{d-1})$ can be dropped, using a more careful analysis based on the Clarkson–Shor technique. We omit details of this improvement.

6. Conclusions

We proved several results on CF-coloring of points and regions. There are numerous problems for further research suggested by our results. In particular, the main open problems we pose in this paper for further research are:

1. *Substantially* improve the bounds on the CF-chromatic number of points in the plane with respect to axis-parallel rectangles.
2. Improve the bounds on the k-CF-chromatic number of points, with respect to balls, in $\mathbb{R}^d$, for $d \geq 3$ and $k \geq 2$.
3. Improve the bounds on the number of colors needed for k-CF-coloring of n balls in $\mathbb{R}^d$, for $d \geq 3$ and $k \geq 2$.
4. Develop deterministic algorithms for CF-coloring. One natural approach is to try to use discrepancy [C], [Ma].
5. Develop a kinetic coloring framework for moving points (or regions in the dual case).
6. Develop a dynamic coloring framework for supporting the more general case where points (or regions in the dual case) can be inserted and deleted.

Acknowledgments

The authors thank Pankaj Agarwal, Noga Alon, Boris Aronov, Timothy Chan, Jeff Erickson, Guy Even, János Pach, Rom Pinchasi, and Micha Sharir for helpful discussions of the problems studied in this paper. In particular, the authors thank Micha Sharir for his invaluable and helpful comments on preliminary versions of this paper.

References

[AKS] N. Alon, M. Krivelevich, and B. Sudakov. Coloring graphs with sparse neighborhoods. *J. Combin. Theory Ser. B*, 77:73–82, 1998.

[AS] N. Alon and J. H. Spencer. *The Probabilistic Method*, 2nd edn. Wiley Interscience, New York, 2000.

[BKOS] M. De Berg, M. Van Kreveld, M. H. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*, 2nd edn. Springer-Verlag, New York, 2000.

[C] B. Chazelle. *The Discrepancy Method*. Cambridge University Press, Cambridge, 2000.

[CS] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.

[Ef] A. Efrat. The complexity of the union of (α, β)-covered objects. In *Proc. 15th Annu. ACM Sympos. Comput. Geom.*, pages 134–142, 1999.

[ELRS] G. Even, Z. Lotker, D. Ron, and S. Smorodinsky. Conflict-free colorings of simple geometric regions with applications to frequency assignment in cellular networks. *SIAM J. Comput.*, 33(1):94–136, 2003.

[Er] J. Erickson. Nice point sets can have nasty delaunay triangulations. *Discrete Comput. Geom.*, 30(1):109–132, 2003.

[ES] A. Efrat and M. Sharir. On the complexity of the union of fat convex objects in the plane. *Discrete Comput. Geom.*, 23:171–189, 2000.

[HW] D. Haussler and E. Welzl. ε-Nets and simplex range queries. *Discrete Comput. Geom.*, 2:127–151, 1987.

[KLPS] K. Kedem, R. Livne, J. Pach, and M. Sharir. On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles. *Discrete Comput. Geom.*, 1:59–71, 1986.

[Ma] J. Matoušek. *Geometric Discrepancy*. Springer-Verlag, New York, 1999.

[Mu] K. Mulmuley. *Computational Geometry: An Introduction through Randomized Algorithms*. Prentice-Hall, Englewood Cliffs, NJ, 1994.

[PA] J. Pach and P. K. Agarwal. *Combinatorial Geometry*. Wiley, New York, 1995.

[PT] J. Pach and G. Tóth. Conflict free colorings. In *Discrete and Computational Geometry, The Goodman–Pollack Festschrift*, pages 665–671. Springer-Verlag, Heidelberg, 2003.

[Sa] N. Sauer. On the density of families of sets. *J. Combin. Theory Ser. A*, 13:145–147, 1972.

[SA] M. Sharir and P. K. Agarwal. *Davenport–Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York, 1995.

[Sm] S. Smorodinsky. Combinatorial Problems in Computational Geometry. Ph.D. thesis, School of Computer Science, Tel Aviv University, Tel Aviv, 2003.

[W] D. B. West. *Introduction to Graph Theory*, 2nd edn., Prentice-Hall, Englewood Cliffs, NJ, 2001.

Received August 13, 2003, *and in revised form June* 3, 2004. *Online publication April* 7, 2005.