Conflict-Free Colorings of Simple Geometric Regions with Applications to Frequency Assignment in Cellular Networks *

Guy Even† Zvi Lotker† Dana Ron† Shakhar Smorodinsky‡

November 6, 2002

Abstract

Motivated by a frequency assignment problem in cellular networks, we introduce and study a new coloring problem that we call *Minimum Conflict-Free Coloring* (Min-CF-Coloring). In its general form, the input of the Min-CF-coloring problem is a set system $(X, \mathcal{S})$, where each $S \in \mathcal{S}$ is a subset of X. The output is a coloring χ of the sets in $\mathcal{S}$ that satisfies the following constraint: for every $x \in X$ there exists a color i and a unique set $S \in \mathcal{S}$, such that $x \in S$ and $\chi(S) = i$. The goal is to minimize the number of colors used by the coloring χ.

Min-CF-coloring of general set systems is not easier than the classic graph coloring problem. However, in view of our motivation, we consider set systems induced by simple geometric regions in the plane.

In particular, we study disks (both congruent and non-congruent), axis-parallel rectangles (with a constant ratio between the smallest and largest rectangle), regular hexagons (with a constant ratio between the smallest and largest hexagon), and general congruent centrally-symmetric convex regions in the plane. In all cases we have coloring algorithms that use $O(\log n)$ colors (where n is the number of regions). Tightness is demonstrated by showing that even in the case of unit disks, $\Theta(\log n)$ colors may be necessary. For rectangles and hexagons we also obtain a constant-ratio approximation algorithm when the ratio between the largest and smallest rectangle (hexagon) is a constant.

We also consider a dual problem of CF-coloring points with respect to sets. Given a set system $(X, \mathcal{S})$, the goal in the dual problem is to color the elements in X with a minimum number of colors so that every set $S \in \mathcal{S}$ contains a point whose color appears only once in S. We show that $O(\log |X|)$ colors suffice for set systems in which X is a set of points in the plane, and the sets are intersections of X with scaled translations of a convex region. This result is used in proving that $O(\log n)$ colors suffice in the primal version.

*A preliminary version of this paper appeared in FOCS-2002.

†Dept. of Electrical Engineering Systems, Tel-Aviv University, Tel-Aviv 69978, Israel. E-mail:{`guy, zvilo, danar`}`@eng.tau.ac.il`.

‡School of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel. E-mail:`smoro@tau.ac.il`.

1 Introduction

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; links 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. This scanning takes place automatically and enables smooth transitions between links when a client is mobile. Consider a client that is within the reception range of two base-stations. If these two base-stations are assigned the same frequency, then mutual interference occurs, and the links between the client and each of these conflicting base-stations are rendered too noisy to be used. A base-station may serve a client provided that the reception is strong enough and interference from other base stations is weak enough. The fundamental problem of frequency assignment in cellular network is to assign frequencies to base-stations so that every client is served by some base-station. The goal is to minimize the number of assigned frequencies since spectrum is limited and costly.

We consider the following abstraction of the above problem which we refer to as the *minimum conflict-free (CF) coloring problem.*

Definition 1 *Let X be a fixed domain (e.g., the plane) and let $\mathcal{S}$ be a collection of subsets of X (e.g., disks whose centers correspond to base-stations). A function $\chi : \mathcal{S} \rightarrow \mathbb{N}$ is a* CF-coloring *of $\mathcal{S}$ if, for every $x \in \bigcup_{S \in \mathcal{S}} S$, there exists a color $i \in \mathbb{N}$, such that $\{S \in \mathcal{S} : x \in S \text{ and } \chi(S) = i\}$ contains a single subset $S \in \mathcal{S}$.*

The goal in the minimum CF-coloring problem is to find a CF-coloring that uses as few colors as possible. It is not hard to verify that in its most general form defined above, this problem is not easier than vertex coloring in graphs, and is even as hard to approximate. An adaptation of the NP-completeness proof of minimum coloring of intersection graphs of unit disks by [CCJ90] proves that even CF-coloring of unit disks (or unit squares) in the plane is NP-complete. Since this proof is based on a reduction from coloring planar graphs, it follows that approximating the minimum number of colors required in a CF-coloring of unit disks is NP-hard for an approximation ratio of $\frac{4}{3} - \varepsilon$, for every $\varepsilon > 0$.

1.1 Our Results

We restrict our attention to set systems $(X, \mathcal{R})$ where X is a set of points in the plane and $\mathcal{R}$ is a family of subsets of X that are defined by the intersections of X with geometric regions in the plane (e.g., disks). We refer to the members of $\mathcal{R}$ as *ranges*, and to $(X, \mathcal{R})$ as a *range-space.*

1.1.1 CF-coloring of Disks

Given a set of disks $\mathcal{S}$, the *size-ratio* of $\mathcal{S}$ is the ratio between the largest and the smallest radius of disks in $\mathcal{S}$. For simplicity we assume that the smallest radius is 1. The *local density* of a set of disks $\mathcal{S}$ is the maximum number of centers of disks in $\mathcal{S}$ that are contained in a square of diameter 1. We denote the local density of $\mathcal{S}$ by $\phi(\mathcal{S})$. For a set of centers $X \subset \mathbb{R}^2$, and for any given radius r, let $\mathcal{S}_r(X)$ denote the set of (congruent) disks having radius r whose centers are the points in X.

Our main results for coloring disks are stated in the following theorem.

Theorem 1 *1. Given a finite set $\mathcal{S}$ of disks with size-ratio ρ, there exists a polynomial-time algorithm that computes a CF-coloring of $\mathcal{S}$ using $O\left(\min\{(\log \rho) \cdot (\log \phi(\mathcal{S})), \log |\mathcal{S}|\}\right)$ colors.*

2. Given a finite set of centers $X \subset \mathbb{R}^2$, there exists a polynomial-time algorithm that computes a coloring χ of X using $O(\log |X|)$ colors, such that χ is a CF-coloring of $\mathcal{S}_r(X)$ for every radius r.

Tightness of Theorem 1 is shown by presenting, for any given integer n, a set $\mathcal{S}$ of n unit disks with $\phi(\mathcal{S}) = n$ for which $\Omega(\log n)$ colors are necessary in every CF-coloring.

In the first part of Theorem 1 the disks are not necessarily congruent. That is, the size-ratio ρ may be bigger than 1. In the second part of Theorem 1, the disks are congruent (i.e., the size-ratio equals 1). However, the common radius is not determined in advance. Namely, the order of quantifiers in the second part of the theorem is as follows: Given the locations of the disk centers, the algorithm computes a coloring of the centers (of the disks) such that this coloring is conflict-free *for every* radius r. We refer to such a coloring as a *uniform* CF-coloring.

Uniform CF-coloring has an interesting interpretation in the context of cellular networks. Assume that base-stations are located in the disk centers X. Assume that a client located at point P has a reception range r. The client is served provided that the disk centered at P with radius r contains a base-station that transmits in a distinct frequency among the base stations within that disk.

Thus, uniform CF-coloring models frequency assignment under the setting of isotropic base-stations that transmit with the same power and clients with different reception ranges. Moreover, the coloring of the base-stations in a uniform CF-coloring is independent of the reception ranges of the clients.

Building on Theorem 1, we also obtain two bi-criteria CF-coloring algorithms for disks having the same (unit) radius. In both cases we obtain colorings that use very few colors. In the first case this comes at a cost of not serving a small area that is covered by the disks (i.e., an area close to the boundary of the union of the disks). In the second case we serve all the area, but we allow the disks to have a slightly larger radius. A formal statement of these bi-criteria results follows.

Theorem 2 *For every $0 < \varepsilon < 1$ and every finite set of centers $X \subset \mathbb{R}^2$, there exist poly-time algorithms that compute colorings as follows:*

1. *A coloring χ of $\mathcal{S}_1(X)$ using $O\left(\log \frac{1}{\varepsilon}\right)$ colors for which the following holds: The area of the set of points in $\bigcup \mathcal{S}_1(X)$ that are not served with respect to χ is at most an ε-fraction of the total area of $\mathcal{S}_1(X)$.*
2. *A coloring of $\mathcal{S}_{1+\varepsilon}(X)$ that uses $O\left(\log \frac{1}{\varepsilon}\right)$ colors such that every point in $\bigcup \mathcal{S}_1(X)$ is served.*

In other words, in the first case, the portion of the total area that is not served is an exponentially small fraction as a function of the number of colors. In the second case, the increase in the radius of the disks is exponentially small as a function of the number of colors.

1.1.2 An $O(1)$-Approximation for CF-Coloring of Rectangles and Regular Hexagons

Let $\mathcal{R}$ denote a set of axis-parallel rectangles. Given a rectangle $R \in \mathcal{R}$, let $w(R)$ ($h(R)$, resp.) denote the width (height, resp.) of R. The *size-ratio* of $\mathcal{R}$ is defined by $\max\left\{\frac{w(R_1)}{w(R_2)}, \frac{h(R_1)}{h(R_2)}\right\}_{R_1,R_2 \in \mathcal{R}}$.

The size ratio of a collection of regular hexagons is simply the ratio of the longest side length and the shortest side length.

Theorem 3 *Let $\mathcal{R}$ denote either a set of axis-parallel rectangles or a set of axis-parallel regular hexagons. Let ρ denote the size-ratio of $\mathcal{R}$ and let $\chi_{\mathrm{opt}}(\mathcal{R})$ denote an an optimal CF-coloring of $\mathcal{R}$.*

1. *If $\mathcal{R}$ is a set of rectangles, then there exists a poly-time algorithm that computes a CF-coloring χ of $\mathcal{R}$ such that $|\chi(\mathcal{R})| = O((\log \rho)^2 \cdot |\chi_{\mathrm{opt}}(\mathcal{R})|)$.*
2. *If $\mathcal{R}$ is a set of hexagons, then there exists a poly-time algorithm that computes a CF-coloring χ of $\mathcal{R}$ such that $|\chi(\mathcal{R})| = O((\log \rho) \cdot |\chi_{\mathrm{opt}}(\mathcal{R})|)$.*

For a constant size-ratio ρ, Theorem 3 implies a constant approximation algorithm.

1.1.3 Uniform CF-coloring of Congruent Centrally-Symmetric Convex Regions

Consider a convex region C and a point O. *Scaling* by a factor $r > 0$ with a respect to a center O is the transformation that maps every point $P \neq O$ to the point P' along the ray emanating from O towards P such that $|P'O| = r \cdot |PO|$. The center point O is a fixed point of the transformation of the scaling. We denote the image of C with respect to such a scaling by $C_{r,O}$. Given a point x and a scaling factor $r > 0$, we denote by $C_{r,O}(x)$ the image of $C_{r,O}$ obtained by the translation that maps O to x. We refer to C' as a *scaled translation* of C if there exist points x, O and a scaling factor $r > 0$ such that $C' = C_{r,O}(x)$. Given a set of centers X and a scaling factor $r > 0$, the set $\mathcal{C}_{r,O}(X)$ denotes the set of scaled translations $\{C_{r,O}(x)\}_{x \in X}$.

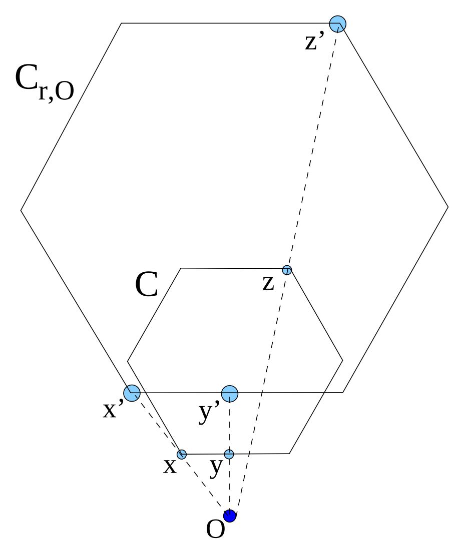

Figure 1: An example of a scaled translation of a regular hexagon C, with respect to the point O, where the scaling factor r is 2. Here the points x, y and z on the small hexagon C are mapped to the points x', y', and z', respectively, on the larger hexagon $C_{r,O}$. The dashed lines correspond to the rays emanating from O towards the points x, y, and z.

A region $C \in \mathbb{R}^2$ is *centrally-symmetric* if there exists a point O (called the *center*) such that the transformation of reflection about O is a bijection of C onto C. Note that disks, rectangles, and regular hexagons are all convex centrally-symmetric regions.

The following theorem generalizes the uniform coloring result presented in Part 2 of Theorem 1 to sets of centrally-symmetric convex regions that are congruent via translations.

Theorem 4 *Let C denote a centrally-symmetric convex region with a center point O. Given a finite set of centers $X \subset \mathbb{R}^2$, there exists a coloring χ of X that uses $O(\log|X|)$ colors, such that χ is a CF-coloring of $\mathcal{C}_{r,O}(X)$, for every scaling factor r.*

A poly-time constructive version of Theorem 4 holds when the region C is "well behaved", e.g., a disk, an ellipsoid, or a polygon. (More formally, a poly-time algorithm for computing Delaunay graphs of arrangements of regions $\mathcal{C}_{r,O}(X)$ is needed.)

1.2 Techniques

1.2.1 A Dual Coloring Problem: CF-Coloring of Points with respect to Ranges

In order to prove Theorem 1 we consider the following coloring problem, which is dual to our original coloring problem described in Definition 1:

Definition 2 *Let $(X, \mathcal{R})$ denote a range space. A function $\chi : X \to \mathbb{N}$ is a* CF-coloring *of X with respect to $\mathcal{R}$ if, for every $R \in \mathcal{R}$, there exists a color $i \in \mathbb{N}$, such that the set $\{x \in R : \chi(x) = i\}$ contains a single point.*

Note that in the original definition of CF-coloring (Definition 1), we were interested in coloring ranges (regions) so as to serve points contained in the ranges, while in Definition 2 we are interested in coloring points so as to "serve" ranges containing the points.

We give a general framework for CF-coloring points with respect to sets of ranges $\mathcal{R}$, and provide a sufficient condition under which a coloring using $O(\log|X|)$ colors can be achieved. This condition is stated in terms of a special graph constructed from $(X, \mathcal{R})$. This graph is the standard Delaunay graph when X is a set of points in the plane and $\mathcal{R}$ is a set of ranges obtained by intersections with disks. We then study several cases in which the condition is satisfied. Theorem 1 and Theorem 4 follow by reduction to these cases. We believe that Theorem 5 stated below (from which Theorem 4 is easily derived), is of independent interest.

Theorem 5 *Let C be a compact convex region in the plane, and let X be a finite set of points in the plane. Let $\mathcal{R} \subseteq 2^X$ denote the set of ranges obtained by intersecting X with all scaled translations of C. Then there exists a CF-coloring of X with respect to $\mathcal{R}$ using $O(\log|X|)$ colors.*

Recently, Pach and Toth [PT02] proved that $\Omega(\log|X|)$ colors are required for CF-coloring *every* set X of points in the plane with respect to disks.

1.2.2 CF-Coloring of Chains

A *chain* $\mathcal{S}$ is a collection of subsets, each assigned a unique index in $\{1, \ldots, |\mathcal{S}|\}$ for which the following holds. For every (discrete) interval $[i, j]$, $1 \le i \le j \le |\mathcal{S}|$, there exists a point $x \in \bigcup_{S \in \mathcal{S}} S$, such that the sub-collection of subsets that contains the point x equals the sub-collection of subsets indexed from i to j. Moreover, for every point $x \in \bigcup_{S \in \mathcal{S}} S$, the set of indexes of subsets that contain the point x is an interval. For an illustration, see Figure 4. We show that chains of unit disks (resp., unit squares and hexagons) are tight examples of Theorem 1 (resp., Theorem 3); namely, every CF-coloring of a chain must use $\Omega(\log|\mathcal{S}|)$ colors, and it is possible to CF-color every chain using $O(\log|\mathcal{S}|)$ colors.

Chains also play an important role in our approximation algorithm for CF-coloring rectangles (and hexagons). Loosely speaking, our coloring algorithm works by decomposing the set of rectangles into chains. An important component in our analysis is understanding and exploiting the intersections between pairs of different chains. Specifically, we show how different types of pairs of chains (see Figures 7 and 10) can "help" each other so as to go below the upper bound on the number of colors required to color chains, which is logarithmic in their size.

1.3 Related Problems

As noted above, minimum CF-coloring of general set systems is not easier (even to approximate) than vertex-coloring in graphs. The latter problem is of course known to be NP-hard, and is even hard to approximate [FK98]. The problem remains hard for the special case of unit disks (and squares), and it is even NP-hard to achieve an approximation ratio of $\frac{4}{3} - \varepsilon$, for every $\varepsilon > 0$ (by an adaptation of [CCJ90]).

Marathe *et al.* [MBH+95] studied the problem of vertex coloring of intersection graphs of unit disks. They presented an approximation algorithm with an approximation ratio of 3. Motivated by channel assignment problems in radio networks, Krumke *et al.* [KMR01] presented a 2-approximation algorithm for the problem of distance-2 coloring problem in families of graphs that generalize intersection graphs of disks.

A natural variant of Min-CF-coloring is Min-CF-multi-coloring. Given a collection $\mathcal{S}$ of sets, a CF-multi-coloring of $\mathcal{S}$ is a mapping χ from $\mathcal{S}$ to *subsets* of colors. The requirement is that for every point $x \in \bigcup_{S \in \mathcal{S}} S$, there exist a color i such that $\{S : \ x \in S, \ i \in \chi(S)\}$ contains a single subset. The Min-CF-multi-coloring problem is related to the problem of minimizing the number of time slots required to broadcast information in a single-hop radio network. In view of this relation, it has been observed by Bar-Yehuda ([B01], based on [BGI92]), that every set-system $(X, \mathcal{S})$ can be CF-multi-colored using $O(\log |X| \cdot \log |\mathcal{S}|)$ colors.

Mathematical optimization techniques have been used to solve a family of frequency assignment problems that arise in wireless communication (for a comprehensive survey see [AHK+01]). We elaborate why these frequency assignment problems do not capture Min-CF-coloring. Basically, such frequency assignment problem are modeled using *interference* or *constraint* graphs. The vertices correspond to base-stations, and edges correspond to interference between pairs of base-stations. Each edge (v, w) is associated with a *penalty function* $p_{v,w} : \mathbb{N} \times \mathbb{N} \to \mathbb{R}$, so that if v is assigned frequency $i \in \mathbb{N}$ and w is assigned frequency $j \in \mathbb{N}$, then a penalty of $p_{v,w}(i, j)$ is incurred. A typical constraint is to bound the maximum penalty on every edge. A typical cost function is the number of frequencies used. CF-coloring cannot be modeled in this fashion because CF-coloring allows for conflicts between base-stations provided that another base-station serves the "area of conflict". Even models that use non-binary constraints (see [DBJC98]) do not capture CF-coloring. We note that the above models take into account interferences between close frequencies, while we have ignored this issue for sake of simplicity. We can however incorporate some variants of such constraints. For example, in the case of unit disks we can easily impose the constraint that for every point x, the frequency assigned to the disk that serves x, differs by at least $\delta_{\min}$ from the frequency assigned to every other disk covering x. By applying Theorem 1 and multiplying each color by $\delta_{\min}$, we can satisfy the above constraint while using $O\left(\min\left\{(\log \rho) \cdot (\log \phi(\mathcal{S})), \log |\mathcal{S}|\right\} \cdot \delta_{\min}\right)$ colors (and there is an example that exhibits tightness).

Frequency assignment problems in cellular networks as well as the positioning problem of base-stations have been vastly studied. See [AKM+01, GGRV00, H01] for other models and many references. Finally, we refer to [HS02, SM02] for further work on CF-coloring problems.

Further Research. Among the open problems related to our results are: (1) Is there a constant approximation algorithm for Min-CF-coloring of unit disks and disks in general? (2) Is it possible to extend our results to Min-CF-coloring with capacity constraints defined as follows: every base-station is given a capacity that bounds the number of clients that it can serve.

Organization. In Section 2, preliminary notions and notation are presented. In Section 3 we describe our results for CF-coloring points with respect to range spaces: We describe a general framework and several applications. In Section 4 we prove our results for CF-coloring of disks (Theorems 1 and 2), which build on results from Section 3. Tightness of Theorem 1 is established in Section 6, and Theorem 4 is proved in Section 5. Our $O(1)$–approximation algorithm for rectangles is provided in Section 7. In Section 8 we discuss how a very similar algorithm can be applied to color regular hexagons. Finally, in Section 9 we derive a couple of additional related results.

2 Preliminaries

2.1 Combinatorial Arrangements

A finite set $\mathcal{R}$ of regions (in the plane) induces the following *equivalence relation*. Every two points x, y in the plane belong to the same class if and only if they reside in exactly the same subset of regions in $\mathcal{R}$. That is, x and y are in the same equivalence class if $\{R \in \mathcal{R} : \ x \in R\} = \{R \in \mathcal{R} : \ y \in R\}$. We refer to each such equivalence class as a *cell*. The set of all cells induced by $\mathcal{R}$ is denoted by *cells*$(\mathcal{R})$. With a slight abuse of notation, we view the pair $(\textit{cells}(\mathcal{R}), \mathcal{R})$ as a range space. To be precise, $(\textit{cells}(\mathcal{R}), \mathcal{R})$ is the the following range space: (a) the ground set is equal to a representative from every cell, and (b) the ranges are the intersections of sets in $\mathcal{R}$ with the ground set. We henceforth refer to the range space $(\textit{cells}(\mathcal{R}), \mathcal{R})$ as the *combinatorial arrangement* induced by $\mathcal{R}$; we denote this combinatorial arrangement by $\mathcal{A}(\mathcal{R})$.

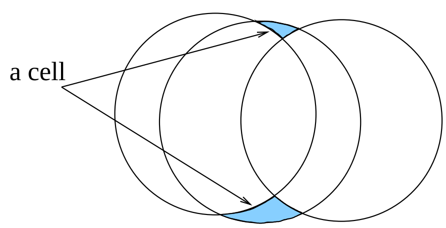

Figure 2: An arrangement of disks. The marked cell corresponds to the regions that are contained in the middle disk, and only in that disk.

The definition of a combinatorial arrangement differs from that of a topological arrangement (where one considers the subdivision into connected components induced by the ranges). For example, Figure 2 depicts a collection of disks. The two shadowed regions constitute a single cell in the combinatorial arrangement induced by the disk. In the definition of a topological arrangement these regions are considered as two separate cells. We often consider combinatorial arrangements of the form $(V, \mathcal{R})$, where $V \subset \textit{cells}(\mathcal{R})$. We refer, in short, to combinatorial arrangements as arrangements.

2.2 Primal and Dual Range Spaces

Consider a range space $(X, \mathcal{R})$. The dual set system is $(\mathcal{R}, X^*)$, where $X^* = \{N(x)\}_{x \in X} \subseteq 2^{\mathcal{R}}$ and $N(x) = \{R \in \mathcal{R} \ : \ x \in R\}$. One may represent a set system by a bipartite graph $(X \cup \mathcal{R}, E)$, with an edge (x, R) if $x \in R$. Under this representation, the dual set-system corresponds to the bipartite graph in which the roles of the two sides of the vertex set are interchanged. Isomorphism of set systems is equivalent to the isomorphism of the bipartite graph representations of the corresponding set systems.

Let $\mathcal{T}$ denote a set of regions in the plane. We use $\mathcal{T}$ to denote a set of regions with some common property; for example, the set of all unit disks, or a set of axis-parallel unit squares. Given a set of points X and a region R (such as a disk), when referring to R as a range (namely, a subset of X) we actually mean $R \cap X$.

A range space $(X, \mathcal{R})$ is a $\mathcal{T}$-type range space if $\mathcal{R} \subseteq \mathcal{T}$. We are interested in situations in which the dual of a $\mathcal{T}$-type range space is isomorphic to a $\mathcal{T}$-type range space.

Definition 3 *A set of regions $\mathcal{T}$ is* self dual *if the dual range space of every $\mathcal{T}$-type range space is isomorphic to a $\mathcal{T}$-type range space.*

For example, it is not hard to verify that the set of all unit disks is self dual. On the other hand, the set of all disks (or even disks of two different radiuses) is not self dual.

The following claim states a condition on $\mathcal{T}$ that is sufficient for $\mathcal{T}$ to be self dual when X is a set of points in the plane.

Claim 1 *Let C be a fixed centrally-symmetric region in the plane, and let $\mathcal{T}$ be the set of all regions congruent (via translation, not rotation) to C. Then $\mathcal{T}$ is self dual.*

Proof: Given a $\mathcal{T}$-type range space $(X, \mathcal{R})$, let Y denote the set of centers of the ranges in $\mathcal{R}$. Let $\mathcal{C}(X)$ denote the set of regions congruent to C centered at points of X. The range space $(Y, \mathcal{C}(X))$ is obviously a $\mathcal{T}$-type range space. To see that this system is isomorphic to the dual range space $(\mathcal{R}, X^*)$, we identify every range $R \in \mathcal{R}$ with its center. Since C is centrally-symmetric, it follows that $y \in C(x)$ if and only if $x \in C(y)$, for every two points x, y. This means that a center $y \in Y$ is in $C(x)$ if and only if the range $C(y)$ contains the point $x \in X$. Hence, for every point $x \in X$, the set $C(x) \cap Y$ equals the set of centers of ranges in $N(x)$, and the claim follows. □

As a corollary of Claim 1 we obtain.

Corollary 6 *Let $\mathcal{T}$ be a set of regions that satisfy the premises of Claim 1. Then CF-coloring arrangements of $\mathcal{T}$-type regions is equivalent to CF-coloring points with respect a $\mathcal{T}$-type set of ranges.*

We rely on Corollary 6 in the proof of Part 2 of Theorem 1 and in the proof of Theorem 4.

3 CF-Coloring Points With Respect to Range Spaces

In this section we present CF-coloring algorithms for points with respect to ranges. The colorings require $O(\log n)$ colors, where n denotes the number of points.

3.1 Intuition

We begin by presenting a simple special case of the general framework. In this case we consider a set X of n points that lie on a straight line and the ranges are intersections of X with disks.

Suppose that we wish to decide which points are colored by the color 1, and then proceed by deciding which points are colored by the color 2, and so on. Let X_i denote the set of points that are colored by the color i. Let $X_{<i}$ (resp. $X_{\leq i}$) denote the set $\bigcup_{j<i} X_j$ (resp. $\bigcup_{j \leq i} X_j$). When

determining X_i, we must make sure that the following condition holds: For every disk D, either (i) D is served by a point colored $j < i$ (i.e. $\exists j < i : |D \cap X_j| = 1$), or (ii) $D \cap X_i$ contains at most one point, or (iii) D contains a point that is not colored yet (i.e. $D \not\subseteq X_{\leq i}$). Correctness follows because if parts (i) and (ii) do not hold, then the point that will serve d will be colored by a color greater than i. In fact, a coloring that follows the above rule has the following property: For every disk, the highest color of a point contained in a disk has multiplicity 1.

The question is how can we guarantee that such an algorithm uses only $O(\log n)$ colors. For example, if every X_i consists of a single point, then obviously correctness holds, but each point is colored by a different color, so n colors are used. To obtain $O(\log n)$ colors, we show that in each stage it is possible to select at least half of the remaining points (i.e., $|X_i| \geq \frac{1}{2} \cdot |X \setminus X_{<i}|$).

The choice of X_i when the points lie on a straight line is simply to pick every other point. By convexity, if a disk D contains two (or more) points from X_i, then it must contain all the points in between these two points. Between every two points in X_i there must exist at least one point not in $X_{\leq i}$. It follows that the condition required from X_i holds, and hence points on a straight line can be colored by $\log n$ colors.

3.2 A General Framework

We start by presenting a general framework for CF-coloring a set X of points with respect to a set $\mathcal{R} \subseteq 2^X$ of ranges, and describe sufficient conditions under which the resulting coloring uses $O(\log n)$ colors. Since every range $R \in 2^X$ that contains a single point from X is trivially served by that point, we assume that every range in $\mathcal{R}$ contains at least two points from X.

Definition 4 *A partition (X_1, X_2) of X is $\mathcal{R}$-useful if $X_1 \neq \emptyset$ and*

$$\text{or} \quad S \cap X_2 \neq \forall S \in \mathcal{R} \ : \ |S \cap X_1| = 1 \quad \text{or} \quad S \cap X_2 \neq \emptyset.$$

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

$i \leftarrow 0$. (i denotes an unused color)
while $X \neq \emptyset$ **do**
  Find an $\mathcal{R}$-useful decomposition (X_1, X_2) of X. (We elaborate subsequently on the implementation of this step.)
  Color: $\forall x \in X_1 \ : \ \chi(x) \leftarrow i$.
  Project: $X \leftarrow X_2$ and $\mathcal{R} \leftarrow \{S \cap X_2 : S \in \mathcal{R},\ |S \cap X_1| \neq 1 \text{ and } |S \cap X_2| \geq 2\}$.
  Increment: $i \leftarrow i + 1$.
end while

Claim 2 *The coloring of X computed by* CF-color$(X, \mathcal{R})$ *is a CF-coloring of X with respect to $\mathcal{R}$.*

Proof: Consider a range $S \in \mathcal{R}$. Let i denote the last iteration in which $X \cap S \in \mathcal{R}$. In other words, in the ith iteration, the $\mathcal{R}$-useful decomposition (X_1, X_2) of X satisfies either $|X_1 \cap S| = 1$ or $|X_2 \cap S| = 1$. In the first case, S is served by the single element $x \in X_1 \cap S$ (which is colored i). In the second case, S is served by the single element $x \in X_2 \cap S$ (which will be assigned some color $i' > i$). Observe that if at iteration i the range space $\mathcal{R}$ becomes empty while X is not empty, then the partition $(X, \emptyset)$ is $\mathcal{R}$-useful, and all the remaining points can be colored with the color $i + 1$. □

Note that Algorithm CF-color computes a CF-coloring in which every range $S \in \mathcal{R}$ is served by the point $x \in S$ for which the color $\chi(x)$ is maximal.

3.2.1 Sufficient Conditions for Using $O(\log |X|)$ Colors

Algorithm CF-color uses $O(\log |X|)$ colors if in every iteration $|X_1| = \Omega(|X|)$. We formalize a condition guaranteeing that $|X_1|$ is a constant fraction of $|X|$. The condition is phrased in terms of a special graph that is a attached to the range space $(X, \mathcal{R})$. We refer to ranges $S \in \mathcal{R}$ as *minimal* if they are minimal with respect to inclusion. Recall that we assume that for every $S \in \mathcal{R}$, $|S| \geq 2$ (since ranges of size one are served trivially).

Definition 5 *A* Delaunay graph *of the set system* $(X, \mathcal{R})$ *is a graph* $DG_{\mathcal{R}}(X, E)$*, defined as follows. For every minimal* $S \in \mathcal{R}$*, pick a pair* $u, v \in S$ *and define* $e(S) = (u, v)$*. The edge set* E *is defined by* $E = \{e(S)\}_{S \in \mathcal{R}}$*.*

A Delaunay graph of a set system is not uniquely defined if there exist minimal ranges that contain more than two points. To simplify the presentation, we abuse notation and refer to the Delaunay graph of a set system as if it were unique. It can be shown that in the case of points in $\mathbb{R}^2$ and ranges obtained by intersections with disks, the definition of a Delaunay graph is the standard definition [BKOS97].

Claim 3 *If* $X_1 \subseteq X$ *is an independent set in* $DG_{\mathcal{R}}$*, then the partition* $(X_1, X \setminus X_1)$ *is* $\mathcal{R}$*-useful.*

Proof: Assume for the sake of contradiction that there exists an independent set X_1 such that $(X_1, X \setminus X_1)$ is not an $\mathcal{R}$-useful decomposition of X. That is, there exists a range $S \in \mathcal{R}$ such that $|S \cap X_1| \neq 1$ and $S \cap (X \setminus X_1) = \emptyset$. Note that assuming that $S \cap (X \setminus X_1) = \emptyset$ necessarily implies that $S \subseteq X_1$, and so we may replace the first condition ($|S \cap X_1| \neq 1$) by $|S \cap X_1| \geq 2$.

Let S' denote a minimal range that is a subset of S (hence $S' \subseteq X_1$). By the definition of the set of edges E in the Delaunay graph $DG_{\mathcal{R}}$ of $(X, \mathcal{R})$, it follows that there is an edge $e(S')$ between two points in S'. But this contradicts the assumption that X_1 is an independent set, and the claim follows. □

The method we use to show that Delaunay graphs have large independent sets is to show that Delaunay graphs are planar. Another easy way to show that there exists a large independent set is, for example, to show that the number of edges is linear.

Claim 4 *If in each iteration of the algorithm, the Delaunay graph of* $(X, \mathcal{R})$ *is planar, then Algorithm 1 uses* $O(\log |X|)$ *colors.*

Proof: By Claim 3, it suffices to show that, in every iteration of Algorithm CF-color, the Delaunay graph has an independent set X_1 that satisfies $|X_1| = \Omega(|X|)$. The existence of a large independent set X_1 in the Delaunay graph $DG_{\mathcal{R}}(X, E)$ follows from the planarity of $DG_{\mathcal{R}}$. Planarity implies that the graph is 4-colorable, and therefore, the largest color-class is an independent set of size at least $|X|/4$. (Note that $DG_{\mathcal{R}}$ can be 4-colored in polynomial time.) □

One could easily color planar graphs using 6 colors since the minimum degree is at most 5. This means that a greedy algorithm could be used to find an independent set of size at least $|X|/6$.

In the rest of this section we apply Algorithm CF-color to three types of range spaces: disks in the plane, half-spaces in $\mathbb{R}^3$, and homothetic centrally-symmetric convex regions in the plane. For each of these cases we prove that the premise of Claim 4 is satisfied. That is, that the Delaunay graph of the corresponding range space is planar. Moreover, for disks, half space in $\mathbb{R}^3$, and polygons, the corresponding Delaunay graphs are computable in polynomial time, which implies that Algorithm CF-color is polynomial.

3.3 Disks in the Plane

Lemma 5 *Let X denote a set of n points in the plane. Let $\mathcal{R}$ denote the collection of all subsets of X of size at least two obtained by intersecting X with a disk. Then it is possible to color X with respect to $\mathcal{R}$ using $O(\log n)$ colors.*

The Delaunay graph that we attach to the set system $(X, \mathcal{R})$ is a exactly the standard Delaunay graph of a planar point set [BKOS97, Thm 9.6 (ii)]. The Delaunay graph of a planar point set X is planar [BKOS97, Thm 9.5]. Hence Lemma 5 directly follows by applying Claim 4.

3.4 Half-Spaces in $\mathbb{R}^3$

Given a hyper-plane H (not parallel to the z-axis), the positive half-space H^+ is the set of all points that either lie on or are above H. We denote by $\mathcal{H}^+$ the set of all positive half-spaces in $\mathbb{R}^3$.

Lemma 6 *Let X be a set of n points in $\mathbb{R}^3$. Let $\mathcal{R}$ denote the collection of all subsets of X of size at least two obtained by intersecting X with a half-space in $\mathcal{H}^+$. Then there exists a CF-coloring of X with respect to $\mathcal{R}$ that uses $O(\log n)$ colors.*

We make the following simplifying assumption: The points in X are in *convex position*. If not, then all the points of X that are in the interior of the convex hull may be colored by a unique "passive" color. The coloring of non-extreme points by a passive color means, in effect, that these interior points are removed. This reduction is justified by the fact that every half-space H^+ that intersects the convex hull of X must contain an extreme point of X. The coloring will be a CF-coloring of the extreme points of X with respect to positive half-spaces, and hence $X \cap H^+$ will be served as well.

Claim 7 *Every minimal range in the range space $(X, \mathcal{R})$ is a pair of points.*

Proof: Consider a range $R \in \mathcal{R}$ defined by half space H^+. Translate H upwards as much as possible so that every further translation upward reduces the range defined by the positive half-space to less than 2 points. Let H_1 denote the plane parallel to H obtained by this translation. Let R_1 denote the range corresponding to the positive half-space H_1^+. If R_1 contains more than two points, then either R_1 is contained in the plane H_1 or all but one of the points in R_1 are in the plane H_1. Assume that $R_1 \subset H_1$. Consider a line ℓ in H_1 that passes through two adjacent vertices u, v (i.e., an edge) in the polygon corresponding to the (two-dimensional) convex hull of R_1 relative to the plane H_1. Tilt the plane H_1 slightly where the line ℓ serves as the axis of rotation. It is possible to rotate H_1 so that the resulting plane H_2 satisfies $X \cap H_2^+ = \{u, v\}$. A similar argument applies if there is a single point in $R_1 \setminus H_1$, and the claim follows. $\square$

Proof of Lemma 6: Claim 7 implies that the Delaunay graph $DG_{\mathcal{R}} = (X, E)$ of the range space $(X, \mathcal{R})$ is defined by $(u, v) \in E$ if and only if there exists a positive half-space H^+ such that $X \cap H^+ = \{u, v\}$. Let $CH(X)$ denote the convex hull of X, and let $G' = (X, E')$ denote the skeleton graph of the convex hull of X. Namely, a pair (u, v) is an edge in E' if and only if there exists a supporting plane H of $CH(X)$ such that $H \cap X = \{u, v\}$. It is well known that the skeleton graph G' is planar. By definition, the edge set of the Delaunay graph is contained in the edge set of the skeleton graph. Hence the Delaunay graph is planar and by Claim 4, X can be CF-colored with respect to $\mathcal{R}$ using $O(\log |X|)$ colors. $\square$

3.5 Scaled Translations of a Convex Region in the Plane

In this Subsection we prove Theorem 5. We first introduce some definitions and notation.

For a closed region C let ∂C denote the boundary of C and let $\mathring{C}$ denote the interior of C. We next recall the definition of homothecy (c.f. [C69, p. 68]).

Definition 6 *A transformation $\tau : \mathbb{R}^2 \to \mathbb{R}^2$ is a* homothecy *if there exists a point O (called the* homothetic center*) and a nonzero real number λ (called the* similitude ratio*) such that:*

1. *O is a fixed point of τ (namely, $O = \tau(O)$),*
2. *every point $P \neq O$ is mapped to a point $\tau(P)$ where: (i) $\tau(P)$ is on the line OP, and (ii) the length of the segment $O\tau(P)$ satisfies: $|O\tau(P)| = \lambda \cdot |OP|$.*

We denote that C' is a scaled translation of C by $C' \sim C$. For a homothetic transformation $\tau : \mathbb{R}^2 \to \mathbb{R}^2$, we denote the image of a set $S \subseteq \mathbb{R}^2$ under τ by $\tau(S)$. Note that if the similitude ratio of a homothecy τ is positive, then $\tau(C) \sim C$.

Definition 7 *A range $S \in \mathcal{R}$ is* induced *by a region C if $S = C \cap X$. A range $S \in \mathcal{R}$ is* boundary-induced *by a closed region C if $S = \partial C \cap X$ and $\mathring{C} \cap X = \emptyset$.*

Recall that for the purpose of CF-coloring, ranges that contain one point as well as the empty range are trivial. Hence, we do not consider the empty set and subsets that contain a single point to be ranges. Therefore, we define the range space $\mathcal{R}$ induced by a collection of regions $\mathcal{C}$ by:

$$\mathcal{R} = \{C \cap X : C \in \mathcal{C} \text{ and } |C \cap X| \geq 2\}.$$

It follows that minimal ranges contain at least two points.

Let C denote a compact convex region C in the plane. Let $X \subset \mathbb{R}^2$ denote a finite set of points in the plane. Let $(X, \mathcal{R})$ denote the range space induced by the set of all scaled translations of C. By Claim 4, in order to prove Theorem 5, it suffices to prove that the Delaunay graph of $(X, \mathcal{R})$ is planar. To this end we first show:

Claim 8 *Every minimal range $S \in \mathcal{R}$ is boundary-induced by a region $C' \sim C$.*

Proof: Since S is a range, there exists a scaled translation $C_S \sim C$ such that $X \cap C_S = S$. By contracting C_S, if necessary, we may guarantee that the boundary of C_S contains a point from S. The interior of C_S contains at most one point of S. Otherwise, by an infinitesimal contraction we are left with a range $S' \subsetneq S$ that contains at least two points, thus contradicting the minimality of S.

We now show how to find a region $C' \sim C$ such that all of S lies on the boundary of C'. Let $x \in S$ denote a point on the boundary of C_S. If S is not boundary-induced by C_S, then there is a unique point $y \in S \cap \mathring{C}_S$. We apply the following homothecy τ. Let y' denote the intersection point of the boundary of C_S with the half-open ray emanating from x towards y. Set x to be the homothetic center, and set the similitude ratio to be the ratio $|xy|/|xy'|$. By definition of τ, both x and y are on the boundary of C'. By minimality of S, it follows that $C' \cap X = S$. By definition of τ and convexity of C, it follows that $C' \subseteq C_S$. If a point $z \in S$ is in the interior of C', then it is in the interior of C_S, hence $z = y$, which contradicts $y \in \partial C'$. It follows that every point in S is in the boundary of C', and the claim follows. □

We now show that a planar drawing of the Delaunay graph $DG_{\mathcal{R}} = (X, E)$ is obtained if its edges are drawn as straight line segments. Consider two edges $(x_0, y_0), (x_1, y_1) \in E$. For $i = 0, 1$, assume that $x_i, y_i \in S_i$, for a minimal range $S_i \in \mathcal{R}$, where $S_0 \neq S_1$. Let $C_i \sim C$ denote scaled translations of C such that S_i is boundary-induced by C_i. If $C_0 \cap C_1 = \emptyset$, then the segments x_0y_0 and x_1y_1 do not cross each other. If $C_0 \cap C_1 \neq \emptyset$, then the boundaries ∂C_0 and ∂C_1 intersect.

We first consider the case that ∂C does not contain a straight side. Namely, no three points on ∂C are co-linear. Under this assumption, since C_i is a scaled translation of C, for $i = 0, 1$, it follows that $\partial C_0 \cap \partial C_1$ contains at most two points.

If $\partial C_0 \cap \partial C_1$ contains a single point p, then, one can separate the convex regions C_0 and C_1 using a straight line passing through p. This separating line implies that the segments x_0y_0 and x_1y_1 cannot cross each other.

If $\partial C_0 \cap \partial C_1$ contains two points, denote these points by p and q. The boundary ∂C_i is partitioned into two simple curves each delimited by the points p and q; one curve is contained in is $\partial C_i \setminus \mathring{C}_{1-i}$ and the second curve is $\partial C_i \cap C_{1-i}$. We denote the curve $\partial C_i \setminus \mathring{C}_{1-i}$ by γ_i, and we denote the curve $\partial C_i \cap C_{1-i}$ by γ_i'. Since the interior $\mathring{C}_{1-i}$ lacks points of X, it follows that x_i and y_i are in γ_i.

In order to prove that the segments x_0y_0 and x_1y_1 do not cross each other, it suffices to show that the line pq separates $\gamma_0 \setminus \{p, q\}$ and $\gamma_1 \setminus \{p, q\}$ (intersection of two edges means that the edges share an interior point, which cannot be p or q). Assume, for the sake of contradiction, that $\gamma_0 \setminus \{p, q\}$ and $\gamma_1 \setminus \{p, q\}$ are on the same side of the line pq. These curves do not intersect, and together with the segment pq, one must contain the other, contradicting their definition.

The case in which ∂C contains a straight side (and so $\partial C_0 \cap \partial C_1$ may contain a sub-segment of such a straight side), is dealt with similarly to the case that ∂C does not contain a straight side. It is not hard to verify that in such a case $\partial C_0 \cap \partial C_1$ consists of at most two connected components (each either straight line or a single point). By picking p to be any point from one component and q to be any point from the other component, we can apply essentially the same argument used above.

This concludes the proof of Theorem 5.

4 CF-Colorings of Arrangements of Disks

In this section we prove Theorems 1 and 2 stated in the Introduction.

4.1 Proof of Theorem 1

Part 2 of Theorem 1 is proved as follows. The disk centers $X \subset \mathbb{R}^2$ are given. Consider a radius r (which is not given to the algorithm!), and apply Corollary 6 to the arrangement $\mathcal{A}(\mathcal{S}_r(X))$. Let Y denote the set consisting of representatives from every cell in $\mathit{cells}(\mathcal{S}_r(X))$. The dual range space is isomorphic to a range space with (i) a ground set X and (ii) ranges induced by $\mathcal{S}_r(Y)$. We extend the range space to ranges induced by all the disks (of all radiuses). A CF-coloring of the points in X with respect to the set of all disks is also a CF-coloring of every arrangement $\mathcal{A}(\mathcal{S}_r(X))$. Part 2 of Theorem 1 follows now directly from Lemma 5.

We now turn to proving Part 1 of Theorem 1.

A Transformation to Points and Half-Spaces. In what follows, we show that the problem of CF-coloring of n arbitrary disks in the plane reduces to CF-coloring of a set of points X in $\mathbb{R}^3$ with respect to the set of ranges $\mathcal{H}^+(X)$ determined by all positive half-spaces containing at least two points from X.

We use a fairly standard dual transformation that transforms a point $p = (a, b)$ in $\mathbb{R}^2$ to a plane p^* in $\mathbb{R}^3$, with the parameterization: $z = -2ax - 2by + a^2 + b^2$, and transforms a disk S in $\mathbb{R}^2$, with center (x, y) and radius $r \geq 0$, to a point S^* in $\mathbb{R}^3$, with coordinates $(x, y, r^2 - x^2 - y^2)$.

It is easily seen that in this transformation, a point $p \in \mathbb{R}^2$ lies inside (resp., on the boundary of, outside) a disk S, if and only if the point $S^* \in \mathbb{R}^3$ lies above (resp., on, below) the plane p^*. Indeed, assume that point $p = (a, b)$ lies inside (resp., on the boundary of, outside) the disk S with center (x, y) and radius r. This can be formulated by the inequality: $(a - x)^2 + (b - y)^2 < r^2$ or $-2ax - 2by + a^2 + b^2 < r^2 - x^2 - y^2$ (resp, an equality $=$, or inequality with $>$). which is equivalent to that of, the point $(x, y, r^2 - x^2 - y^2) = S^*$ lies above (resp, on or below) the plane $z = -2ax - 2by + a^2 + b^2$ (which is the dual p^* of p), as asserted.

Given a collection $\mathcal{S} = \{S_1, \ldots, S_n\}$ of n distinct disks in the plane, one can use the above transformation to obtain a collection $\mathcal{S}^* = \{S_1^*, \ldots, S_n^*\}$ of n points in $\mathbb{R}^3$, such that any CF-coloring of $\mathcal{S}^*$ with respect to $\mathcal{H}^+(\mathcal{S}^*)$ with k colors, induces a CF-coloring of the disks of $\mathcal{S}$ with the same set of k colors.

As shown in Subsection 3.4 (Lemma 6), it is possible to apply Algorithm 1 to obtain a CF-coloring of the points in $\mathcal{S}^*$ with respect to $\mathcal{H}^+(\mathcal{S}^*)$ using $O(\log n)$ colors. Recall that Part 1 of Theorem 1 states that the number of colors is of the order of the minimum between $\log n$ and $(\log \rho) \cdot (\log \phi(\mathcal{S}))$ (where ρ is the size-ratio of $\mathcal{S}$ and $\phi(\mathcal{S})$ its local density). To obtain the latter bound we proceed in two steps: first we assume that the size-ratio is at most 2, and then we deal with the more general case.

The Tiling. Assume that the size ratio ρ is at most 2. By scaling, we may assume that every radius is in the interval $[1, 2]$. We partition the plane into square tiles having diameter 1. We say that a disk S *belongs* to tile T if the center of S is in T. We denote the subset of disks in $\mathcal{S}$ that belong to T by $\mathcal{S}(T)$. Note that the union of the disks in any given tile intersects at most 16 different tiles. We assign a *palette* (i.e., a subset of colors) to each tile using 16 different palettes, where the disks belonging to a particular tile are assigned colors from the tile's palette. Palettes are assigned to tiles by following a periodic 4×4 assignment. This assignment has the property that any two disks that belong to different tiles either do not intersect or their tiles are given different palettes (so that necessarily the two disks are assigned different colors). By the definition of local density we have that $|\mathcal{S}(T)| \leq \phi(\mathcal{S})$ for every tile T. Since we can color the set of disks $\mathcal{S}(T)$ belonging to tile T using $O(\log |\mathcal{S}(T)|)$ colors, and the total number of palettes is 16, we get the desired upper bound of $O(\log \phi(\mathcal{S}))$ colors.

The general case of arbitrary size-ratio is dealt with by first partitioning the set of disks into classes according to their radius. The ith class consists of disks, the radius of which is in the interval $[2^i, 2^{i+1})$. Within each class, the size-ratio is bounded by 2, hence we can CF-color each class using $O(\log \phi(\mathcal{S}))$ colors. By using a different (super-)palette per class, we obtain the desired bound on the number of colors, i.e., $O((\log \rho) \cdot (\log \phi(\mathcal{S})))$.

4.2 Bi-Criteria CF-coloring Algorithms

In this section we prove Theorem 2. The first part of the theorem reveals a tradeoff between the number of colors used and the fraction of the area that is served. The second part of the theorem reveals a tradeoff between the number of colors used to serve the union of the unit disks and the radius of the serving disks.

We first derive the following corollary from Theorem 1.

Corollary 7 *Let $\mathcal{S}$ be a set of unit disks, and let $d_{\min}(\mathcal{S})$ be the minimum distance between centers of disks in $\mathcal{S}$. If $d_{\min}(\mathcal{S}) \leq 2$ then every arrangement $\mathcal{A}(\mathcal{S})$ of unit disks can be CF-colored using $O\left(\log\left(\min\left\{|\mathcal{S}|, \frac{1}{d_{\min}(\mathcal{S})}\right\}\right)\right)$ colors.*

Observe that if $d_{\min}(\mathcal{S}) > 2$ then a single color suffices since the disks are disjoint.
Proof: Obviously $\phi(\mathcal{S}) \leq |\mathcal{S}|$. Since a square of diameter 1 can be packed with at most $O(\frac{1}{d_{\min}(\mathcal{S}(T))^2})$ many disks of radius $d_{\min}(\mathcal{S}(T))$, it follows that that $\phi(\mathcal{S}) = O(\frac{1}{d_{\min}(\mathcal{S}(T))^2})$. $\square$

Let $X \subset \mathbb{R}^2$ denote a finite set of centers of disks. Recall that $\mathcal{S}_r(X) = \{B(x,r) \mid x \in X\}$, where $B(x,r)$ denotes a disk of radius r centered at x. Let $A_r(X) = \bigcup_{x\in X} B(x,r)$. The area of a region A in the plane is denoted by $|A|$. Let $L_r(X)$ denote the length of the boundary of $A_r(X)$. In order to prove Theorem 2 we shall need the following two lemmas which are proved subsequently.

Lemma 9 *For every finite set X of points in the plane,*

$$|A_1(X)| \geq \frac{1}{2} \cdot L_1(X)$$

Lemma 10 *For every finite set X of points in the plane and every $\varepsilon > 0$,*

$$|A_{1+\varepsilon}(X) - A_1(X)| \leq (2\varepsilon + \varepsilon^2) \cdot L_1(X)$$

Proof of Theorem 2: We start with the second part. Let $X' \subseteq X$ denote a maximal subset with respect to inclusion such that $||x_1 - x_2|| \geq \varepsilon$, for every $x_1, x_2 \in X'$. Observe that $\bigcup \mathcal{S}_1(X) \subseteq \bigcup \mathcal{S}_{1+\varepsilon}(X')$. Corollary 7 implies that $\mathcal{S}_{1+\varepsilon}(X')$ can be CF-colored using $O\left(\log \frac{1+\varepsilon}{\varepsilon}\right)$ colors. The second part follows.

We now turn to the first part. Let $\varepsilon_1 = \varepsilon/6$ and X' as above. Corollary 7 implies that there exists a CF-coloring χ of $\mathcal{S}_1(X')$ using $O\left(\log \frac{1}{\varepsilon}\right)$ colors. To complete the proof we need to show that

$$\frac{|A_1(X) - A_1(X')|}{|A_1(X)|} \leq \varepsilon.$$

Since $A_1(X) \subseteq A_{1+\varepsilon_1}(X')$ and $A_1(X') \subseteq A_1(X)$. It suffices to prove that

$$\frac{|A_{1+\varepsilon_1}(X') - A_1(X')|}{|A_1(X')|} \leq \varepsilon.$$

By Lemmas 9 and 10 it follows that

$$\frac{|A_{1+\varepsilon_1}(X') - A_1(X')|}{|A_1(X')|} \leq \frac{(2\varepsilon_1 + \varepsilon_1^2) \cdot L_1(X')}{\frac{1}{2} \cdot L_1(X')} = 4 \cdot \varepsilon_1 + 2\varepsilon_1^2.$$

Since $\varepsilon < 1$, it follows that $4 \cdot \varepsilon_1 + 2\varepsilon_1^2 \leq 6 \cdot \varepsilon_1 = \varepsilon$, and the corollary follows. $\square$

4.2.1 Proving Lemmas 9 and 10

We denote a sector by *sect*(Q, α, r), where Q is its center, α is its angle, and r is its radius. A boundary sector of $A_1(X)$ is a sector *sect*$(Q, \alpha, 1)$ such that $Q \in X$ and its arc is on the boundary of $A_1(X)$. A boundary sector is maximal if it is not contained in another boundary sector. We measure angles in radians. Therefore, in a unit disk, (1) the angle of a sector equals the length of its arc, and (2) the area of a sector equals half its angle.

Lemma 11 *The intersection of every two different maximal boundary sectors in $A_1(X)$ has zero area.*

Proof: The lemma is obvious if the boundary sectors belong to the same disk. Let $Q_1, Q_2 \in X$, and let D_i denote the circles centered at Q_i, for $i = 1, 2$, as depicted in Figure 3. Let $sect_i$ denote a boundary sector that belongs to circle D_i, for $i = 1, 2$. Let ℓ denote the line defined by the intersection points of the circles D_1 and D_2. The line ℓ separates the centers Q_1 and Q_2 so that they belong to different half-planes. The sector $sect_i$ is contained in the half-plane that contains Q_i, and hence $sect_1 \cap sect_2$ contains at most two points. The lemma follows. □

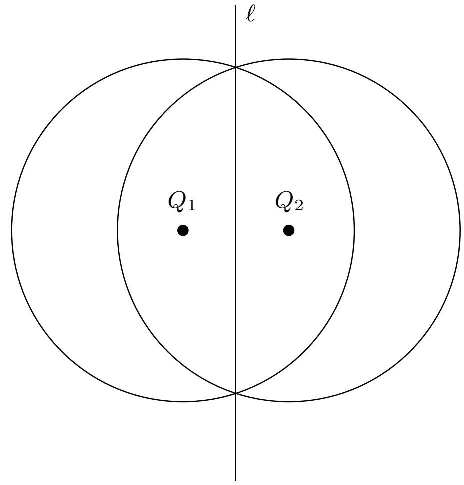

Figure 3: Proof of Lemma 11.

Proof of Lemma 9: The sum of the angles of the maximal boundary sectors of $A_1(X)$ equals $L_1(X)$. By Lemma 11, the maximal boundary sectors are disjoint, and hence the sum of their areas is bounded by $|A_1(X)|$. But, the area of a sector of radius 1 whose angle equals α is $\alpha/2$. □

Lemma 12 *Let X denote a finite set of points in the plane. For every $P \in A_{1+\varepsilon}(X) - A_1(X)$, there exists a point $Q \in X$, such that (1) $P \in B(Q, 1+\varepsilon)$ and (2) the segment $\overline{PQ}$ contains a boundary point of $A_1(X)$.*

Proof: Let Q denote a closest point in X to P. Since $P \in A_{1+\varepsilon}(X) - A_1(X)$, it follows that $P \in B(Q, 1+\varepsilon)$. Let Y denote the point at distance 1 from Q along the segment $\overline{QP}$. All we need to show is that Y is on the boundary of $A_1(X)$. If not, then Y is in the interior of a disk $B(Q', 1)$, for $Q' \in X - \{Q\}$. The triangle inequality implies that Q' is closer to P than Q, a contradiction. The lemma follows. □

Proof of Lemma 10: Lemma 12 implies that, for every point $P \in A_{1+\varepsilon}(X) - A_1(X)$, there exists boundary sector $sect(Q, \alpha, 1)$ of $A_1(X)$ (where $Q \in X$), such that

$$P \in sect(Q, \alpha, 1+\varepsilon) - sect(Q, \alpha, 1)$$

It follows that

$$|A_{1+\varepsilon}(X) - A_1(X)| \leq \sum_{sect(Q,\alpha,1)} |sect(Q, \alpha, 1+\varepsilon) - sect(Q, \alpha, 1)| = \sum_{sect(Q,\alpha,1)} \alpha \cdot (2\varepsilon + \varepsilon^2),$$

where $sect(Q, \alpha, 1)$ ranges over all maximal boundary sectors of $A_1(X)$. The claim follows by observing that the sum of the angles of the boundary sectors of $A_1(X)$ equals $L_1(X)$. □

5 Proof of Theorem 4

Theorem 4 follows from Theorem 5 similarly to the way Part 2 of Theorem 1 was shown to follow from Lemma 5.

Specifically, let C be a centrally-symmetric convex region with a center point O, and X the set of centers we are given. Consider a particular scaling factor r, and apply Corollary 6 to the arrangement $\mathcal{A}(\mathcal{C}_{r,O}(X))$. Let Y denote the set consisting of representatives from every cell in $cells(\mathcal{C}_{r,O}(X))$. The dual range space is isomorphic to a range space with (i) a ground set X (ii) ranges induced by $\mathcal{C}_{r,O}(Y)$. We extend the range space to ranges induced by all scaled translations of C. A CF-coloring of the points in X with respect to all scaled translations of C is also a CF-coloring of every arrangement $\mathcal{A}(\mathcal{C}_{r,O}(X))$. Theorem 4 follows now directly from Theorem 5.

6 Chains and CF-Coloring of Chains

In this section we introduce a combinatorial structure that we call a *chain*. Chains are used to establish the tightness of Theorem 1. They are also central to our $O(1)$ approximation algorithms for rectangles and hexagons.

6.1 Combinatorial Structure

Consider an arrangement $\mathcal{A}(\mathcal{S})$ of a collection of sets $\mathcal{S}$. We associate with every cell $v \in cells(\mathcal{S})$ the subset $N(v) \subseteq \mathcal{S}$ of regions that contain the cell, namely, $N(v) = \{S \in \mathcal{S} : v \subseteq S\}$.

A set $\mathcal{S}$ of regions in the plane is said to be *indexed* if the regions are given indexes from 1 to $|\mathcal{S}|$. In the following definition we identify a region with its index.

Definition 8 *Let $\mathcal{S}$ denote an indexed set of n regions. The arrangement $\mathcal{A}(\mathcal{S})$ satisfies the* interval property *if $N(v)$ is a (discrete) interval $[i, j] \subseteq [1, n]$, for every cell $v \in cells(\mathcal{S})$.*

The arrangement $\mathcal{A}(\mathcal{S})$ satisfies the full interval property *if it satisfies the interval property and, in addition, for every interval $[i, j] \subseteq [1, n]$, there exists a cell $v \in cells(\mathcal{S})$ such that $N(v) = [i, j]$.*

The definition of the (full) interval property is sensitive to the indexing. Indexes of regions are usually based on the order of appearance of the regions along the boundary of the union of the regions. We refer, in short, to an arrangement of an indexed set of regions that satisfies the full interval property as a *chain*.

The definition of a chain implies that an arrangement $\mathcal{A}(\mathcal{S})$ is a chain if and only if the dual range space is isomorphic to $(\{1, \ldots, n\}, \{[i, j] : 1 \leq i \leq j \leq n\})$, where $n = |\mathcal{S}|$. The next lemma, which follows directly from this observation, shows that the chain property is hereditary.

Lemma 13 *Let $\mathcal{S}$ denote an indexed set of regions. Let $\mathcal{S}' \subseteq \mathcal{S}$, and let the indexes of regions $\mathcal{S}'$ agree with their order in $\mathcal{S}$. If $\mathcal{A}(\mathcal{S})$ is a chain, then $\mathcal{A}(\mathcal{S}')$ is also a chain.*

Before discussing colorings of chains we observe that it is easy to construct chains. Consider a set $\mathcal{S}$ of n unit disks with centers positioned along a straight line distance $\frac{1}{n+1}$ apart. Index the disks from 1 to n according to the position of their center from left to right. The arrangement $\mathcal{A}(\mathcal{S})$ is depicted in Figure 4. Observe that every two disks in the arrangement intersect.

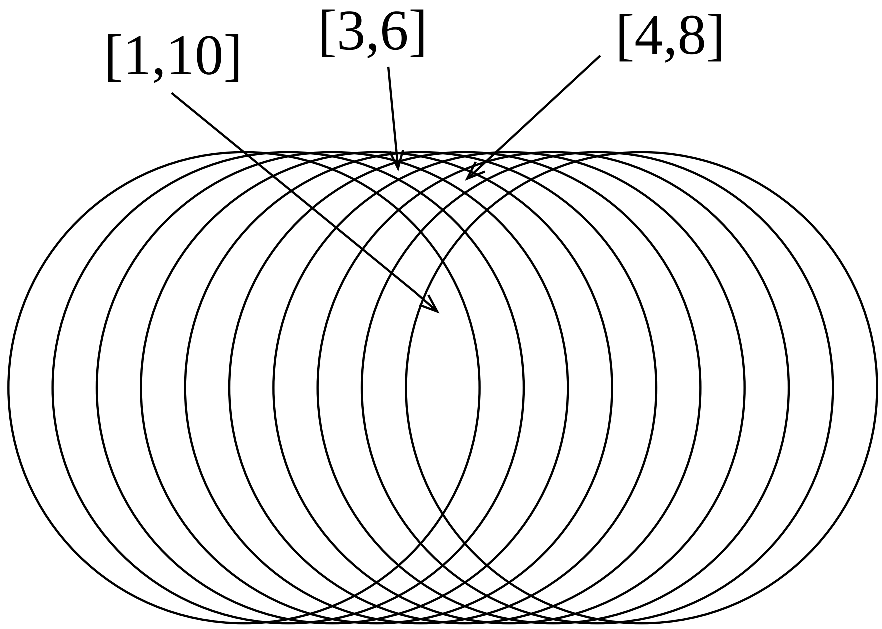

Figure 4: A chain of disks, where the disks are numbered $1, \ldots, 10$, from left to right. The three cells that are marked correspond to the three respective intervals.

We apply duality to prove that the arrangement $\mathcal{A}(\mathcal{S})$ is a chain. The arrangement is the range space $(cells(\mathcal{S}), \mathcal{S})$. Let X denote a set of representatives of cells in $cells(\mathcal{S})$ and let Y denote the centers of unit disks in $\mathcal{S}$. The dual range space is the pair $(Y, \{N(x)\}_{x \in X})$. Since the disks are unit disks it follows that $N(x)$ is the intersection of Y with a unit disk centered at x. The set Y is indexed and its points are located along a line sufficiently close so that they are included in a unit disk. Hence the collection of sets $\{N(x)\}_{x \in X}$ is simply the set of of all intervals $[i, j] \subseteq [1, n]$. It follows that the arrangement $\mathcal{A}(\mathcal{S})$ is a chain, as claimed.

6.2 CF-Colorings of Chains

In this subsection we show that the number of colors both necessary and sufficient for CF-coloring a chain of n regions is $\Theta(\log n)$.

Lemma 14 *Every CF-coloring of a chain of n regions uses $\Omega(\log n)$ colors.*

Proof: Let $I_{a,b}$ denote the set $\{[i, j] : a \leq i \leq j \leq b\}$, namely, the set of all sub-intervals of $[a, b]$. By definition, the dual range space of a chain is isomorphic to the range space $([1, n], I_{1,n})$. Therefore, CF-coloring a chain is equivalent to CF-coloring $[1, n]$ with respect to $I_{1,n}$. We hence focus on the latter problem. Let $f(n)$ denote the minimum number of colors required for such a coloring.

Consider an optimal CF-coloring χ_n of $[1, n]$ with respect to $I_{1,n}$. Let i denote the index that serves the interval $[1, n]$. It follows that for every index $j \neq i$, $\chi(j) \neq \chi(i)$. Since $\chi(i)$ is unique, it follows that every sub-interval that contains i can be served by i.

We partition $I_{1,n}$ into three sets as follows: (i) $I_{1,(i-1)}$ - the set of all sub-intervals of $[1, i-1]$, (ii) I' - the set of all sub-intervals of $[1, n]$ that contain i, and (iii) $I_{(i+1),n}$ - the set of all sub-intervals of $[i+1, n]$. (Observe that if $i = 1$ ($i = n$) then $I_{1,(i-1)}$ ($I_{(i+1),n}$) is empty.)

Since i can only serve intervals in I', we are left with two range spaces that are the dual of (shorter) chains. Namely, the range space $([1, (i-i)], I_{1,(i-1)})$ and the range space $([(i+i), n], I_{(i+1),n})$.

Since $\chi(j)$ just differ from $\chi(i)$ for every $j \neq i$, it follows that $f(n)$ satisfies the following recurrence equation:

$$f(n) \geq 1 + \max_i \{f(i-1), f(n-i)\}.$$

Therefore, $f(n) = \Omega(\log n)$, and the lemma follows. $\square$

Lemma 15 *Every indexed arrangement of n regions that satisfies the interval property can be CF-colored with $O(\log n)$ colors.*

Since every chain satisfies the interval property, the above lemma holds in particular for chains. Recall that in Section 3.1 we already presented a proof of the above lemma in the special case of unit disks.

Proof: We use the same notation as in the proof of the previous lemma. Without loss of generality the dual range space is isomorphic to $([1, n], I_{i,n})$ (adding ranges does not make CF-coloring a set of points with respect to a set of ranges any easier). Hence, we focus on CF-coloring of such a dual range space.

We show by induction that $f(n) \leq \lceil \log n \rceil + 1$. The induction basis is trivial. For $n > 1$, let $i = \lceil n/2 \rceil$ and color it with the color $\lceil \log n \rceil$. The index i serves all the sub-intervals of $[1, n]$ that contain i. A sub-interval of $[1, n]$ that does not contain i is either in $I_{1,(i-1)}$ or in $I_{(i+1),n}$. The induction hypothesis implies that the range spaces $([1, (i-1)], I_{1,(i-1)})$ and $([(i+1), n], I_{(i+1),n})$ can each be colored by $1 + \lceil \log(n/2) \rceil = \lceil \log n \rceil$ colors. Since the ground sets of these range spaces are disjoint, we may use the same set of colors for these two range spaces. It follows that at most $\lceil \log n \rceil + 1$ colors are used, as required. $\square$

7 An Approximation Algorithm for Rectangles

In this section we prove Theorem 3 for the case of axis-parallel rectangles. Most of the proof deals with the special case of axis-parallel unit squares. In Section 7.4 we point out modifications required for rectangles.

7.1 Preliminaries

Let $\mathcal{R}$ be a set of axis-parallel rectangles of side length at least 1. We denote a set of axis-parallel unit-squares by $\mathcal{S}$. For simplicity, we assume that the rectangles (squares, resp.) in $\mathcal{R}$ ($\mathcal{S}$, resp.) are arranged in general position. Let $\Gamma = \{\urcorner, \ulcorner, \lrcorner, \llcorner\}$ denote the set of corner *types*. We denote the top-right corner of a rectangle R by $\urcorner(R)$. In general, for a corner $\gamma \in \Gamma$, we denote the γ-corner of R by $\gamma(R)$. The x-coordinate (y-coordinate) of a γ-corner of a rectangle R is denoted by $x_\gamma(R)$ ($y_\gamma(R)$). Let $\mathrm{op} : \Gamma \to \Gamma$ denote the permutation that swaps opposite corners (i.e. $\mathrm{op} = (\llcorner, \urcorner)(\ulcorner, \lrcorner)$). The *center* of a rectangle R is the intersection point of its two main diagonals.

The Tiling. We partition the plane into "half-open" square tiles having side-lengths 1/2, namely $T_{i,j} = [i/2, (i+1)/2) \times [j/2, (j+1)/2)$. We say that a rectangle R belongs to tile T, if the center of R is in T. We denote the set of rectangles in $\mathcal{R}$ that belong to tile T by $\mathcal{R}(T)$. A tile T is an orphan if $\mathcal{R}(T) = \emptyset$. A tile is bare if no rectangle in $\mathcal{R}$ intersects it. We say that two tiles are e-neighbors

(v-neighbors) if they share an edge (a corner). The v-neighbor T' of T which shares its γ-corner with the op(γ) corner of T is denoted T_γ.

Tiles are half-open and their side length is defined to be half the minimum side length of a rectangle so that: (i) if a rectangle R belongs to a tile T, then rectangle R covers the tile T; and (ii) a tile can contain at most one corner of a rectangle.

In the case of a set $\mathcal{S}$ of unit-squares, squares belonging to T_γ intersect T with their γ-corner. Moreover, the corners of a unit-square $S \in \mathcal{S}(T)$ reside in v-neighbors of T. Hence, a square S only intersects the tile it belongs to and the neighbors of that tile.

Corner Chains. We next consider chains obtained by rectangles having the same corner-type in a common region. Let T be a fixed tile, and let $Q \subseteq T$ denote a rectangle. Let $\gamma \in \Gamma$ denote a corner type. Let $\mathcal{R}(Q,\gamma)$ denote the set of rectangles $R \in \mathcal{R}$ that satisfy $\gamma(R) \in Q$. The size of the tile T implies that every rectangle of side length at least 1 contains at most one corner in T. Define the Q-*envelope* of $\mathcal{R}(Q,\gamma)$ to be the boundary $\bigcup \mathcal{R}(Q,\gamma)$ that is in Q (see Figure 5). The vertices of an Q-envelope are either corners $\gamma(R)$, for $R \in \mathcal{R}(Q,\gamma)$, or intersections of sides of two rectangles. Let $\widetilde{\mathcal{R}}(Q,\gamma)$ denote the subset of rectangles in $\mathcal{R}(Q,\gamma)$ that participate in the Q-envelope of $\mathcal{R}(Q,\gamma)$.

The following claim shows that the corner-type γ determines whether the Q-envelope is non-increasing or non-decreasing.

Claim 16 *The Q-envelope of $\mathcal{R}(Q,\gamma)$ is non-decreasing (non-increasing) if $\gamma \in \{\ulcorner, \lrcorner\}$ ($\gamma \in \{\llcorner, \urcorner\}$).*

Proof: We prove the claim for $\gamma = \urcorner$. An analogous argument holds for the other cases. Let $R_1, \ldots, R_m$ ($m = |\widetilde{\mathcal{R}}(Q,\gamma)|$) be an ordering of $\widetilde{\mathcal{R}}(Q,\gamma)$ which satisfies $x_\urcorner(R_1) < x_\urcorner(R_2) < \ldots < x_\urcorner(R_m)$. We show that $y_\urcorner(R_1) > y_\urcorner(R_2) > \ldots > y_\urcorner(R_m)$. Assume in contradiction that for some pair of squares $R_k, R_\ell \in \widetilde{\mathcal{R}}(Q,\gamma)$ where $k < \ell$ (so that $x_\urcorner(R_k) < x_\urcorner(R_\ell)$), we have that $y_\urcorner(R_k) < y_\urcorner(R_\ell)$. But in such a case we would have that $(x_\urcorner(R_k), y_\urcorner(R_k)) \in R_\ell$, contradicting the fact that R_k belongs to the envelope $\widetilde{\mathcal{R}}(Q,\gamma)$. $\square$

In the next claim we show that the set of cells of the arrangement of rectangles $\widetilde{\mathcal{R}}(Q,\gamma)$ that are contained in Q form a chain. Formally, this means that the arrangement corresponding to the set of rectangles $\{R \cap Q\}_{R \in \widetilde{\mathcal{R}}(Q,\gamma)}$ is a chain. To simplify notation we state the claim as follows.

Claim 17 *Index the rectangles of $\widetilde{\mathcal{R}}(Q,\gamma)$ according to the x-coordinate of their γ-corner. Then $\widetilde{\mathcal{R}}(Q,\gamma)$ is a chain with respect to Q.*

Claim 17 justifies referring to $\widetilde{\mathcal{R}}(Q,\gamma)$ as a *corner-chain*.

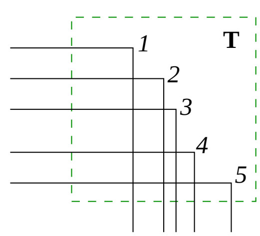

Figure 5: An illustration of a corner-chain; indexes appear next to the corners of the rectangles.

Proof: We prove the claim for $\gamma = \urcorner$. The other 3 cases can be reduced to this case by "turning the picture". Let $R_1, \ldots, R_m$ ($m = |\widetilde{\mathcal{R}}(Q,\gamma)|$) be an ordering of $\widetilde{\mathcal{R}}(Q,\gamma)$ according to the x coordinates of their centers. Let $P_{i,j}$, for $i < j$, denote the intersection of the right side of R_i and the topside of R_j. By Claim 16, it follows that $P_{i,j}$ is well defined and that $P_{i,j} \in Q$, for every $1 \le i < j \le m$. The arrangement of $\widetilde{\mathcal{R}}(Q,\gamma)$ in Q is a set of rectangle shaped cells, the corners of which are the set of points $\{P_{i,j}\}_{i,j}$ plus intersections of the rectangles R_i with the sides of Q. The cell v whose corners are $P_{i-1,j+1}, P_{i-1,j}, P_{i,j}$ and $P_{i,j+1}$ satisfies $N(v) = [i,j]$, and the claim follows. $\square$

Disjoint Palettes. In the case of unit squares we assign a palette (i.e., a subset of colors) to each tile using in total 9 disjoint palettes. Palette distribution is such that neighboring tiles are assigned different palettes (i.e., periodically assign 9 different palettes to blocks of 3×3 tiles). The tile size implies that if two squares belong to different tiles that are assigned the same palette, then the squares have an empty intersection.

7.2 Main Lemmas

In this section we lay the ground for our algorithm and its analysis by presenting our main lemmas. For simplicity we focus on a collection $\mathcal{S}$ of unit squares. In Subsection 7.4 we discuss how to perform the extension to general rectangles.

Tiling combined with coloring of squares in each tile using $O(\log \phi(\mathcal{S}))$ colors may lead to a CF-coloring that is far from optimal. The reason is that squares whose centers reside in different, but neighboring tiles, may interact with each other in a manner that allows us to save in the number of colors used. For an illustration see Figure 6.

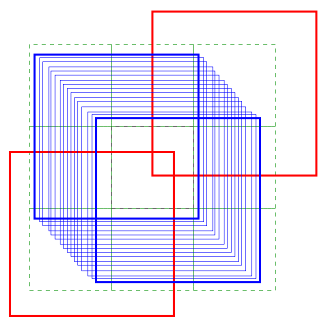

Figure 6: An example illustrating how by taking into account intersections between squares that belong to different tiles we may significantly reduce the number of colors required in a CF-coloring. Here there is a large number of squares that belong to the middle tile and constitute a chain. If we color the squares of each tile separately, the number of colors used is logarithmic in the size of the chain. However, there is a CF-coloring that uses only 5 colors: simply color each of the thick squares by a distinct color and use the 5th color for the remaining squares.

In the rest of this section we provide our main lemmas concerning interactions between corner-chains of opposite corners and corner-chains of adjacent corners.

Corner-chains of adjacent corners. Consider a rectangle Q with side lengths at most $1/2$. Let $\widetilde{\mathcal{S}}_{\ulcorner} = \widetilde{\mathcal{S}}(Q, \ulcorner)$ and $\widetilde{\mathcal{S}}_{\urcorner} = \widetilde{\mathcal{S}}(Q, \urcorner)$ denote corner-chains corresponding to adjacent corners (the other 3 cases of pairs of adjacent corners can be reduced to this case by "turning the picture"). We show that by picking at most one square from each corner-chain, it is possible to "separate" between the chains.

Let $\{S_i\}_{i=1}^{m}$ ($\{S'_i\}_{i=1}^{m'}$) denote the ordering of the squares in $\widetilde{\mathcal{S}}_{\neg}$ ($\widetilde{\mathcal{S}}_{\ulcorner}$) in increasing (decreasing) order of x-coordinate of their centers (or corners in Q). By Claim 17 both indexed sets $\widetilde{\mathcal{S}}_{\neg}$ and $\widetilde{\mathcal{S}}_{\ulcorner}$ are chains with respect to Q.

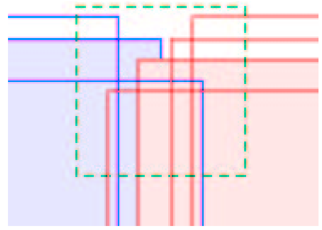

Figure 7: An illustration for Lemma 18. The tile is depicted by a green dashed square. Only the corners of squares in the corner chains are depicted. The two filled squares are the selected squares S_k and S'_ℓ.

Lemma 18 *There exist two squares, $S_k \in \widetilde{\mathcal{S}}_{\neg}$ and $S'_\ell \in \widetilde{\mathcal{S}}_{\ulcorner}$ such that*

1. *The prefixes $\{S_1, \ldots, S_{k-1}\}$ and $\{S'_1, \ldots, S'_{\ell-1}\}$ are disjoint, namely, for every $S_{k'}$ and $S'_{\ell'}$ such that $k' < k$, and $\ell' < \ell$, we have $S_{k'} \cap S'_{\ell'} = \emptyset$;*
2. *Each of the prefixes $\{S_1, \ldots, S_{k-1}\}$ and $\{S'_1, \ldots, S'_{\ell-1}\}$ is a chain with respect to $Q \setminus (S_k \cup S'_\ell)$.*
3. *The union of S_k and S'_ℓ covers every point in Q that is covered by a square in one of the suffixes. Namely, $\left(\bigcup_{t=k+1}^{m}(S_t \cap Q)\right) \bigcup \left(\bigcup_{t=\ell+1}^{m'}(S'_t \cap Q)\right) \subseteq S_k \cup S'_\ell$.*

The implication of this lemma is that it is possible to select two squares S_k and S'_ℓ so as to serve all cells that are contained in the union $(\bigcup_{t=k}^{m}(S_t \cap Q)) \cup \left(\bigcup_{t=\ell}^{m'}(S'_t \cap Q)\right)$. Furthermore, each of the prefixes is a chain with respect to the remaining region.

Proof: Consider the Q-envelopes of the two corner-chains. Both envelopes are "stairs"-curves. By Claim 16, the Q-envelope of $\widetilde{\mathcal{S}}_{\neg}$ ($\widetilde{\mathcal{S}}_{\ulcorner}$) is non-increasing (non-decreasing). Hence the Q-envelopes intersect at most once. If they do not intersect, then the claim is trivial (pick the last square from each chain). Otherwise, let P denote the intersection point. Let the selected squares S_k and S'_ℓ be the squares that intersect in point P. We assume that P is along the horizontal upper side of S'_ℓ (i.e., $P_y = y_{\ulcorner}(S'_\ell)$) and along the vertical right side of S_k (i.e., $P_x = x_{\neg}(S_k)$) (the reverse case is reduced to this case by "flipping the picture").

Part (1) of the claim follows by showing that the vertical line passing through P separates the prefixes. Namely, if $A \in S_{k'}$, $k' < k$, then $A_x < P_x$ (i.e., the x-coordinate of point A is less than the x-coordinate of point P). Similarly, if $B \in S'_{\ell'}$, $\ell' < \ell$, then $B_x > P_x$. To show that $A_x < P_x$, assume (for the sake of contradiction) that $A_x \geq P_x$. It follows that $x_{\neg}(S_{k'}) \geq A_x \geq P_x$. The ordering of squares in $\widetilde{\mathcal{S}}_{\neg}$ implies that $x_{\neg}(S_{k'}) < x_{\neg}(S_k) = P_x$, a contradiction. To show that $B_x > P_x$, assume that $B_x \leq P_x$. Since the Q-envelope of $\widetilde{\mathcal{S}}_{\ulcorner}$ is non-decreasing and squares in $\widetilde{\mathcal{S}}_{\ulcorner}$ are indexed from right to left, it follows that $y_{\ulcorner}(S_{\ell'}) > y_{\ulcorner}(S_\ell) = P_y$. It follows that P is in the interior of $S'_{\ell'}$, a contradiction.

To prove Part (2) it suffices to show that (i) $S_{k'} \cap S'_\ell = \emptyset$ if $k' < k$, and (ii) $S'_{\ell'} \cap S_k = \emptyset$ if $\ell' < \ell$. This is sufficient since $\widetilde{\mathcal{S}}_{\neg}$ ($\widetilde{\mathcal{S}}_{\ulcorner}$) is a chain with respect to Q. Hence, every cell corresponding to an interval $[i,j] \subseteq [1, k-1]$ ($\subseteq [1, \ell-1]$) of $\widetilde{\mathcal{S}}_\urcorner$ ($\widetilde{\mathcal{S}}_\ulcorner$) in Q is disjoint from $S_k \cap S'_\ell$, and the full interval property is preserved.

In order to verify (i), consider a square $S_{k'}$ for $k' < k$. To show that $S_{k'} \cap S'_\ell = \emptyset$, we prove that $x_\urcorner(S_{k'}) < x_\ulcorner(S'_\ell)$. The ordering of $\widetilde{\mathcal{S}}_\urcorner$ implies that $y_\urcorner(S_{k'}) > y_\urcorner(S_k) > P_y$. If $x_\urcorner(S_{k'}) \geq x_\ulcorner(S'_\ell) = P_y$, then $P \in S_{k'}$, a contradiction. It follows that $x_\urcorner(S_{k'}) < x_\ulcorner(S'_\ell)$, and part (i) follows. Part (ii) is proved analogously, and Part (2) follows.

It remains to prove Part (3). Consider a point $A \in S_{k'} \cap Q$, for $k' > k$. There are two possibilities: (i) $A_x \leq P_x$. In this case, $A_y \leq y_\urcorner(S_{k'}) \leq y_\urcorner(S_k)$. Since $P_x = x_\urcorner(S_k)$, it follows that $A \in S_k$. (ii) $A_x > P_x$. If $A_y \geq P_y$, then $\urcorner(S_{k'})$ is above and to the right of P, hence $P \in S_{k'}$, a contradiction. It follows that $A_y < P_y = y_\ulcorner(S'_\ell)$. Since $P_x \geq x_\ulcorner(S'_\ell)$, it follows that $A \in S'_\ell$. It follows that the suffix of $\widetilde{\mathcal{S}}_\urcorner$ is covered by $S_k \cap S'_\ell$. The proof for the suffix of $\widetilde{\mathcal{S}}_\ulcorner$ is analogous, and part (3) follows. $\square$

Corner-chains of opposite corners. Consider a rectangle Q with side lengths at most 1/2. Let $\widetilde{\mathcal{S}}_\urcorner = \widetilde{\mathcal{S}}(Q, \urcorner)$ and $\widetilde{\mathcal{S}}_\llcorner = \widetilde{\mathcal{S}}(Q, \llcorner)$ denote corner-chains corresponding to opposite corners (the case of the $\lrcorner$-corner and $\ulcorner$-corner is reduced to this case by "flipping the picture"). Let $Q_\urcorner = Q \cap \bigcup_{S \in \widetilde{\mathcal{S}}_\urcorner} S$ and $Q_\llcorner = Q \cap \bigcup_{S \in \widetilde{\mathcal{S}}_\llcorner} S$. Our goal is to select an approximately minimal subset from each corner-chain so as to cover $Q_\urcorner \cup Q_\llcorner$. To this end we find minimal covers of $(Q_\urcorner \setminus Q_\llcorner)$, $(Q_\llcorner \setminus Q_\urcorner)$, and $Q_\urcorner \cap Q_\llcorner$.

Definition 9 *A subset $\widetilde{\mathcal{S}}^m_\urcorner \subseteq \widetilde{\mathcal{S}}_\urcorner$ is a* minimal cover *of $(Q_\urcorner \setminus Q_\llcorner)$ if (i) $\widetilde{\mathcal{S}}^m_\urcorner$ covers $(Q_\urcorner \setminus Q_\llcorner)$, and (ii) no proper subset of $\widetilde{\mathcal{S}}^m_\urcorner$ covers $(Q_\urcorner \setminus Q_\llcorner)$.*

The following lemma shows that minimal covers of $(Q_\urcorner \setminus Q_\llcorner)$ are chains with respect to $(Q_\urcorner \setminus Q_\llcorner)$.

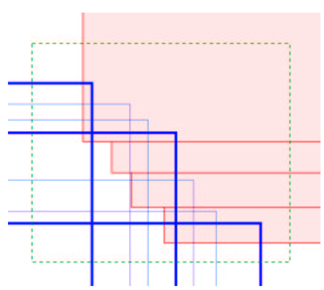

Figure 8: An minimal cover of $(Q_\urcorner \setminus Q_\llcorner)$. The squares of $\widetilde{\mathcal{S}}_\llcorner$ are depicted by filled $\llcorner$-corners. A minimal cover $\widetilde{\mathcal{S}}^m_\urcorner \subseteq \widetilde{\mathcal{S}}_\urcorner$ is depicted by thick $\urcorner$-corners.

Lemma 19 *If $\widetilde{\mathcal{S}}^m_\urcorner \subseteq \widetilde{\mathcal{S}}_\urcorner$ is a minimal cover of $(Q_\urcorner \setminus Q_\llcorner)$, then $\widetilde{\mathcal{S}}^m_\urcorner$ is a chain with respect to $(Q_\urcorner \setminus Q_\llcorner)$.*

Proof: Index the squares in $\widetilde{\mathcal{S}}^m_\urcorner$ according to the x-coordinate of their $\urcorner$-corners. Let $\widetilde{\mathcal{S}}^m_\urcorner = \{S'_1, \ldots, S'_k\}$. Since $\widetilde{\mathcal{S}}_\urcorner$ is a chain with respect to Q, it follows that $\widetilde{\mathcal{S}}^m_\urcorner$ is also a chain with respect to $Q_\urcorner$. For the sake of contradiction, assume that $\widetilde{\mathcal{S}}^m_\urcorner$ is not a chain with respect to $Q_\urcorner \setminus Q_\llcorner$. It follows that there is an interval $[i,j]$ such that the corresponding cell in $\mathcal{A}(\widetilde{\mathcal{S}}^m_\urcorner)$ is contained in $Q_\llcorner$. Assume that $1 < i < j < k$. See Fig. 9 for this case. Consider the corner B of the cell $[i,j]$ in Q defined by the intersection of the sides of S'_{i-1} and S'_{j+1} in Q. Since the cell $[i,j]$ is in $Q_\llcorner$, so is the

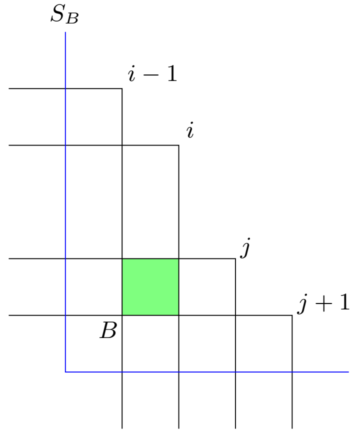

Figure 9: The construction in the proof of Lemma 19.

point B. Let $S_B \in \widetilde{\mathcal{S}}_{\llcorner}$ denote a square that contains B. It follows that the whole cell $[i,j]$ as well as $\neg(S'_{i-1})$, $\neg(S'_i)$, $\neg(S'_j)$, and $\neg(S'_{j+1})$ are in S_B. It is easy to see that we may omit both S'_i and S'_j from $\widetilde{\mathcal{S}}^m_{\neg}$ while still covering $Q_{\neg} \setminus Q_{\llcorner}$, contradicting the assumption that $\widetilde{\mathcal{S}}^m_{\neg}$ is a minimal cover. The argument shows, in fact, that since $\widetilde{\mathcal{S}}^m_{\neg}$ is a minimal cover, then every square in $\widetilde{\mathcal{S}}_{\llcorner}$ contains at most two corners of squares in $\widetilde{\mathcal{S}}^m_{\neg}$. Hence $i = 1$ and $j = k$. This implies that cell $[i,j]$ contains the bottom left corner of Q, hence $\widetilde{\mathcal{S}}_{\llcorner}$ covers $Q_{\neg}$ (with a single square). This leads to a contradiction since in this case $Q_{\neg} \setminus Q_{\llcorner}$ is empty and so is $\widetilde{\mathcal{S}}^m_{\neg}$. □

We now describe a greedy algorithm for finding a subset $\widetilde{\mathcal{S}}^m_{\neg} \subseteq \widetilde{\mathcal{S}}_{\neg}$ that is a minimal cover of $Q_{\neg} \setminus Q_{\llcorner}$. Let $S_1, \ldots, S_m$ be an ordering of the squares in $\widetilde{\mathcal{S}}_{\neg}$ according to increasing value of $x_{\neg}(S_i)$. Recall that $\widetilde{\mathcal{S}}_{\neg}$ is a chain with respect to Q, and therefore every subset of $\widetilde{\mathcal{S}}_{\neg}$ is a chain with respect to Q. For any two indexes $1 \le a \le b \le m$, let $\widetilde{\mathcal{S}}_{\neg}[a,b]$ denote the cell v in the arrangement $\mathcal{A}(\widetilde{\mathcal{S}}_{\neg})$ such that $N(v) = \{S_a, \ldots, S_b\}$.

The greedy algorithm works in an iterative fashion. Let k be the index of the square selected in the last iteration (where initially $k = 0$ and $\widetilde{\mathcal{S}}^m_{\neg} = \emptyset$). Consider all cells $\widetilde{\mathcal{S}}_{\neg}[k+1,\ell]$ where $(k+1) \le \ell \le m$ such that $\widetilde{\mathcal{S}}_{\neg}[k+1,\ell] \cap Q$ is not fully contained in $Q_{\llcorner}$. If there is no such cell, then the algorithm terminates. Otherwise, let ℓ be the minimum index such that $\widetilde{\mathcal{S}}_{\neg}[k+1,\ell]$ is not fully contained in $Q_{\llcorner}$, and add S_ℓ to $\widetilde{\mathcal{S}}^m_{\neg}$.

Claim 20 *The greedy algorithm computes a minimal cover $\widetilde{\mathcal{S}}^m_{\neg} \subseteq \widetilde{\mathcal{S}}_{\neg}$ of $(Q_{\neg} \setminus Q_{\llcorner})$.*

By "rotating the picture" we can obtain an analogous claim concerning a minimal cover $\widetilde{\mathcal{S}}^m_{\llcorner} \subseteq \widetilde{\mathcal{S}}_{\llcorner}$ of $Q_{\llcorner} \setminus Q_{\neg}$.

Proof: Let $k_1 < k_2 < \cdots < k_r$ denote the sequence of squares added to $\widetilde{\mathcal{S}}^m_{\neg}$ by the greedy algorithm. We show that the algorithm computes a cover $\widetilde{\mathcal{S}}^m_{\neg}$ of $Q_{\neg} \setminus Q_{\llcorner}$ by showing that the following invariant holds throughout the algorithm:

$$(Q_{\neg} \setminus Q_{\llcorner}) \cap (S_1 \cup S_2 \cdots \cup S_{k_t}) \subseteq \bigcup_{j \le t} S_{k_j}.$$

The invariant holds trivially when the algorithms starts (as $k_t = 0$). Assume, for the sake of contradiction, that a cell $\widetilde{\mathcal{S}}_{\neg}[i,j]$ (for $i \le j < k_t$) in $Q_{\neg} \setminus Q_{\llcorner}$ is not covered by $\bigcup_{j \le t} S_{k_j}$. If $i \le k_{t-1}$, then there are two cases: (i) $j \le k_{t-1}$, in which case the induction hypothesis already implies that cell $\widetilde{\mathcal{S}}_\neg[i,j]$ is contained in $\bigcup_{j<t} S_{k_j}$. (ii) $j > k_{t-1}$, in which case cell $\widetilde{\mathcal{S}}_\neg[i,j]$ is contained in $S_{k_{t-1}}$. Both cases lead to a contradiction, so we assume that $i > k_{t-1}$. It can be verified that if the cell $\widetilde{\mathcal{S}}_\neg[i,j]$ is not covered by $\bigcup_{j\le t} S_{k_j}$, then the cell $\widetilde{\mathcal{S}}_\neg[k_{t-1}+1,j]$ is also not covered by $\bigcup_{j\le t} S_{k_j}$. In such a case, the greedy algorithm would have chosen $k_t \le j$, a contradiction.

The stopping condition of the algorithm combined with the invariant guarantees that, when the algorithm terminates, $\widetilde{\mathcal{S}}^m_\neg$ covers $Q_\neg \setminus Q_\llcorner$.

Minimality of $\widetilde{\mathcal{S}}^m_\neg$ is proved as follows. Consider a square $S_j \in \widetilde{\mathcal{S}}^m_\neg$. When S_j was added to $\widetilde{\mathcal{S}}^m_\neg$, it was added due to a cell $[i,j]$, with i greater than the index of the square added to $\widetilde{\mathcal{S}}^m_\neg$ just before S_j. The cell $[i,j]$ is covered only by S_j, and hence minimality follows. $\square$

Let $m_\neg = |\widetilde{\mathcal{S}}^m_\neg|$ and let $m_\llcorner = |\widetilde{\mathcal{S}}^m_\llcorner|$. Let $m = \max\{m_\neg, m_\llcorner\}$. In the next lemma we show that it is possible to cover $Q_\neg \cup Q_\llcorner$ by $O(m)$ squares from $\widetilde{\mathcal{S}}_\neg \cup \widetilde{\mathcal{S}}_\llcorner$.

Lemma 21 *There exists a subset $\mathcal{S}' \subseteq \widetilde{\mathcal{S}}_\neg \cup \widetilde{\mathcal{S}}_\llcorner$ of $O(m)$ squares that covers $Q_\neg \cup Q_\llcorner$.*

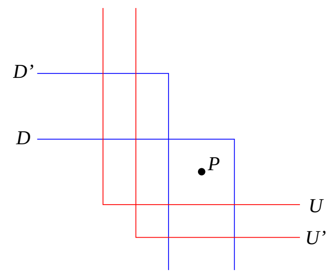

Figure 10: An illustration for Lemma 21. The point P is in $D \cap U$.

Proof: Since $\widetilde{\mathcal{S}}^m_\neg$ ($\widetilde{\mathcal{S}}^m_\llcorner$) covers $Q_\neg \setminus Q_\llcorner$ ($Q_\llcorner \setminus Q_\neg$) and $|\widetilde{\mathcal{S}}^m_\neg \cup \widetilde{\mathcal{S}}^m_\llcorner| \le 2m$, the remaining problem is to cover $Q_\neg \cap Q_\llcorner$ using $O(m)$ squares. For every square $S \in \widetilde{\mathcal{S}}^m_\neg$ consider the set $\widetilde{\mathcal{S}}_\llcorner(S)$ of squares in $\widetilde{\mathcal{S}}_\llcorner$ that intersect S. Define $A_\llcorner(S)$ ($B_\llcorner(S)$, resp.) to be the first (last, resp.) square in $\widetilde{\mathcal{S}}_\llcorner(S)$ when sorted according to their y and/or x-coordinate. We claim that $\bigcup_{S\in\widetilde{\mathcal{S}}^m_\neg}(A_\llcorner(S) \cup B_\llcorner(S))$ covers $Q_\neg \cap Q_\llcorner$.

Consider a point $P \in Q_\neg \cap Q_\llcorner$. Let $D \in \widetilde{\mathcal{S}}_\neg$ ($U \in \widetilde{\mathcal{S}}_\llcorner$, resp.) denote a square that contains P. If $D \in \widetilde{\mathcal{S}}^m_\neg$ or $U \in \widetilde{\mathcal{S}}^m_\llcorner$ then we are done. Otherwise, consider the cell in $\mathcal{A}(\widetilde{\mathcal{S}}_\neg)$ that contains a point slightly to the left of $\llcorner(U)$. This cell is in $Q_\neg \setminus Q_\llcorner$, and therefore, there exists a square $D' \in \widetilde{\mathcal{S}}^m_\neg$ that covers this cell. If $P \in D'$, we are done. Otherwise, consider the square $U' = B_\llcorner(D')$. Such a square exists since U intersects D'. We can now bound the coordinates of $\llcorner(U')$ to show that $P \in U'$ as follows: (i) $x_\llcorner(U') < x_\neg(D') < P_x$, and (ii) $y_\llcorner(U') \le y_\llcorner(U) < P_y$. The claim follows. $\square$

Remark 1 Lemmas 19 and 21 and Claim 20 regarding opposite corner-chains were stated with respect to a rectangle Q that is contained in a tile. The same lemmas and claim hold with respect to a region $Q \subseteq T$ that satisfies the following properties:

The region Q contains two designated points $C_\llcorner$ and $C_\neg$. (When Q is a rectangle then $C_\llcorner$ is the bottom left corner and $C_\neg$ is the top right corner.) The point $C_\llcorner$ is contained in every square in $\widetilde{\mathcal{S}}_\neg$, and the point $C_\neg$ is contained in every square in $\widetilde{\mathcal{S}}_\llcorner$. Moreover, if a square $S \in \widetilde{\mathcal{S}}_\neg$ ($S \in \widetilde{\mathcal{S}}_\llcorner$, resp.) contains the point $C_\neg$ ($C_\llcorner$, resp.), then $Q_\llcorner \setminus Q_\neg = \emptyset$ ($Q_\llcorner \setminus Q_\neg = \emptyset$, resp.).

Note that if Lemma 18 is applied to separate corner-chains of adjacent corners, then the remaining uncovered region in a tile is a region that satisfies the above condition. Hence, after separating corner-chains of adjacent corners, we may apply Lemmas 19 and 21 and Claim 20 for the covering of the remaining region in the tile.

7.3 Coloring Arrangements of Squares

In this section we prove Theorem 3 for unit squares. We first provide a short overview of the proof. The tiling by squares of side length $1/2$ partitions the unit squares into subsets according to the tile that their center belongs to. The assignment of palettes to tiles guarantees that squares from different tiles do not conflict (only tiles at least 3 tiles apart may be assigned the same palette). The goal of the algorithm is to pick an "essential" subset of squares per tile whose union must be served. The coloring of the essential squares per tile is done according to Theorem 4. Recall that a tile is an orphan tile if it does not contain a center of a square. By picking an arbitrary square from each non-orphan tile, all non-orphan tiles are served. The main thrust of the algorithm and its analysis is in serving the covered regions in orphan tiles (i.e., the union of the squares minus the union of non-orphan tiles). The task of selecting a subset of squares that serves the covered parts of orphan tiles is "done independently" by the orphan tiles. The set of essential squares per non-orphan tile is the set of squares that belong to the tile that have been selected by one of the neighboring orphan tiles.

7.3.1 Selection of Squares by Non-Bare Orphan Tiles

Consider a non-bare orphan tile T. In this section we describe how squares from neighboring tiles are selected by T so that these squares serve the area that is covered in T.

Selection of squares consists of three steps: (1) Selection of at most one square from each e-neighbor. This step maximizes service from e-neighbors. (2) Selection of at most two squares from each v-neighbor. This step resolves all interactions between chains of squares corresponding to adjacent corners. (3) Final selection of squares from the remaining chains corresponding to corners. This step takes into account interactions between chains corresponding to opposite corners.

Selecting squares from e-neighbors. Consider the tile T and the set of squares the belong to an e-neighbor T' of T. For brevity, assume that T' is to the left of T and that $\mathcal{S}(T') \neq \emptyset$. Every square $S \in \mathcal{S}(T')$ covers a vertical strip of T. If we select the rightmost square S in $\mathcal{S}(T')$, then it follows that, for every $S' \in \mathcal{S}(T')$, $S' \cap T \subseteq S \cap T$. In the same fashion, we select the closest square to T from each e-neighbor of T. By selecting at most one square from each e-neighbor of T, the first sub-step covers all the points in $T \cap \bigcup_{T' \in e\text{-neighbors}(T)} \mathcal{S}(T')$.

After this step, the region within the tile T that still needs to be served is a rectangle. Let us denote this rectangle by Q. Note that the union of squares in $\mathcal{S}$ may either fully cover or partly cover the rectangle Q. In any case, only squares that belong to v-neighbors of T intersect Q.

Selecting squares from v-neighbors: adjacent corners. Consider the rectangle Q and a corner γ. The squares of $\mathcal{S}(Q,\gamma)$ that participate in the Q-envelope are denoted by $\widetilde{\mathcal{S}}(Q,\gamma)$. By Claim 17, $\widetilde{\mathcal{S}}(Q,\gamma)$ is a chain with respect to Q when indexed according the the x-coordinate of its centers (or γ-corners). By applying Lemma 18 to the 4 appropriate pairs of chains corresponding to adjacent corners, we obtain at most 8 squares that serve as "separators" between the pairs of chains. The selected squares cover all points in Q that are covered by squares in the tails of the chains. Each corner-chain is reduced to a consecutive block of squares between the two selected squares in that chain. The remaining portions of adjacent corner-chains are disjoint. We denote the remaining sub-chain of $\widetilde{\mathcal{S}}(Q,\gamma)$ by $\widehat{\mathcal{S}}_\gamma$. Let Q' denote the sub-region consisting of Q minus the union of the (at most 8) selected squares.

Selecting squares from v-neighbors: opposite corners. In the third step we apply Lemma 21 to each pair of opposite chains $\widehat{\mathcal{S}}_\gamma$ and $\widehat{\mathcal{S}}_{\mathrm{op}(\gamma)}$. This application determines the subsets $\check{\mathcal{S}}_\gamma \subseteq \widehat{\mathcal{S}}_\gamma$ that suffice to serve the intersection of Q' with the union of each pair of chains.

A subtle issue to be addressed is whether the remaining region $Q' \subseteq T$ in the beginning of this step satisfies the premises of Remark 1. Consider for example the chain $\widehat{\mathcal{S}}_\neg$. This chain is a consecutive block of squares from $\widetilde{\mathcal{S}}(T,\neg)$. Let S' and S'' denote the squares in $\widetilde{\mathcal{S}}(T,\neg) \setminus \widehat{\mathcal{S}}_\neg$ that "hug" this block (i.e., S' and S'' were selected in the adjacent corner-chain stage). The designated point $C_\llcorner \in Q'$ is the intersection of the right side of S' and the top side of S''. One can define in this fashion all four designated points C_γ, for $\gamma \in \Gamma$, to show that indeed Q' satisfies the premises of Remark 1.

7.3.2 Coloring the Essential Squares

In the previous steps, each tile T selected a subset of squares that are used to serve the points in $T \cap \bigcup \mathcal{S}$. Given a non-orphan tile T', let $sel(T') \subseteq \mathcal{S}(T')$ denote the subset of squares that are selected by some tile T. If no square in $\mathcal{S}(T')$ is requested from orphan tiles then we select an arbitrary square in $\mathcal{S}(T')$ to serve T'. At this stage we apply Theorem 4 and color each subset $sel(T')$ by $O(\log |sel(T')|)$ colors; these colors are taken from the palette assigned to the tile T'.

Recall that at most one square from $\mathcal{S}(T')$ was requested from each of its 4 e-neighbors. Each of its 4 v-neighbors initially requested at most 2 squares (as "separators" between chains). These requests amount to at most 12 squares. The main contribution to $sel(T')$ is due to the subsets of squares that constitute chains in the v-neighbors of T' and were requested by them.

Let us denote the 4 chains by $\check{\mathcal{S}}_\gamma(T)$ where $\gamma \in \Gamma$ (according to the corner-type they have in the requesting v-neighbor), and let $m_\gamma(T) = |\check{\mathcal{S}}_\gamma(T)|$. Since $|sel(T')| = O(\max_{\gamma\in\Gamma}\{m_\gamma(T)\})$, and since there are 9 palettes, the following corollary follows:

Corollary 8 *For any given set of unit squares $\mathcal{S}$, it is possible to CF-color $\mathcal{S}$ using $O(\log(\max_{T,\gamma}\{m_\gamma(T)\}))$ colors.*

7.3.3 A Lower Bound for Optimal CF-Coloring

In this section we state the lemma that lower bounds the number of colors required by an optimal CF-coloring. Recall that for a tile T and corner-type $\gamma \in \Gamma$, the set of squares that intersect T with corner type γ is denoted by $\mathcal{S}(T,\gamma)$. Recall that $\widetilde{\mathcal{S}}(T,\gamma)$ denotes the subset of squares from $\mathcal{S}(T,\gamma)$ that appear in the T-envelope. By Lemma 17, $\widetilde{\mathcal{S}}(T,\gamma)$ is a (corner) chain.

The following lemma states a lower bound on $\chi_{\mathrm{opt}}(\mathcal{S})$ in terms of the size of a chain $\mathcal{S}'$ with respect to to a region Q in a tile T. The lemma requires two conditions: (i) the chain $\mathcal{S}'$ is a subset of $\widetilde{\mathcal{S}}(T,\gamma)$, for a tile T, and (ii) only rectangles in $\mathcal{S}(T,\gamma)$ contain points in Q, namely, rectangles not in $\mathcal{S}(T,\gamma)$ do not intersect Q. The proof of lemma is similar in structure to the proof of Lemma 14, but differs in one aspect: Rectangles in $\mathcal{S}(T,\gamma) \setminus \mathcal{S}'$ may contains points in Q, and hence can potentially serve cells in the chain $\mathcal{S}'$.

Lemma 22 *Let T denote a tile and let $Q \subseteq T$ denote a closed region. Let $\mathcal{S}' \subseteq \widetilde{\mathcal{S}}(T,\gamma)$ denote a chain with respect to Q. If every square in $\mathcal{S} \setminus \mathcal{S}(T,\gamma)$ does not intersect the region Q, then every CF-coloring of $\mathcal{A}(\mathcal{S})$ requires $\Omega(\log |\mathcal{S}'|)$ colors.*

The proof of Lemma 22 follows the same outline as the proof of Lemma 14 but is provided for completeness. In fact, the same lower bound holds also for CF-multi-coloring, implying Theorem 11 (see Sec. 9).

Proof: For simplicity we consider the case $\gamma = \urcorner$. Let $S_1, \ldots, S_m$ be an ordering of the squares in $\mathcal{S}'$ so that $x_\urcorner(S_1) < \ldots < x_\urcorner(S_m)$. Consider an optimal CF-coloring χ_{opt} of $\mathcal{S}$. Let $I_{1,m}$ denote the set $\bigcap_{i=1}^m S_i \cap Q$. Since $\mathcal{S}'$ is a chain with respect to Q, it follows that $I_{1,m}$ is not empty. Note that when Q is a rectangle, $I_{1,m}$ is a rectangle as well. However, we interested also in non-rectangular regions Q. Loosely speaking, let P denote an "upper right corner" of $I_{1,m}$. More formally, denote the coordinates of P by (P_x, P_y). Select P so that for every $(P'_x, P'_y) \in I_{1,m}$, if $P'_x > P_x$ then $P'_y < P_y$. (Note that an upper right corner is not uniquely defined if Q is not a rectangle).

Since $P \in Q$, only squares in $\mathcal{S}(T, \urcorner)$ contain P. Let $S \in \mathcal{S}(T, \urcorner)$ be a square that serves P in the coloring χ_{opt}. We first consider the case that $\urcorner(S)$, the top-right corner of S, is contained in some $S_k \in \mathcal{S}'$, (in particular, S_k may equal S). In this case, $S \cap Q \subseteq S_k \cap Q$.

We make two observations. The first is that both $\{S_1, \ldots, S_{k-1}\}$ and $\{S_{k+1}, \ldots, S_m\}$ are chains with respect to $Q \setminus S$. This is true since $\{S_1, \ldots, S_m\}$ is a chain with respect to Q, (hence both subsets must be chains with respect to $Q \setminus S_k$), and $Q \cap S \subseteq Q \cap S_k$. Thus, S cannot fully serve any of the cells in these two sub-chains. The second observation is that every square $S' \in \widetilde{\mathcal{S}}(T, \urcorner)$ that serves the top-right corner P' of a cell in one of these sub-chains must also contain P. This is true because every top-right corner of a cell in the arrangement induced by $\{Q \cap S_i\}_{i=1}^k$ dominates P (i.e., the x and y coordinates of such corners are not smaller than P_x and P_y, respectively). Since S serves P, it follows that the color of every square $S_i \neq S$ must be different from $\chi_{\text{opt}}(S)$.

If $\urcorner(S)$, the top-right corner of S, is not contained in any $S_k \in \widetilde{\mathcal{S}}(T, \urcorner)$, then define k as follows: $k = \max\{i : x_\urcorner(S_i) < x_\urcorner(S)\}$. Since $\mathcal{S}'$ is a subset of a corner chain, it follows that $\{S_1, \ldots, S_k\}$ and $\{S_{k+1}, \ldots, S_m\}$ are chains with respect to $Q \setminus S$. Furthermore, similarly to what was shown above, for any square S' that can serve (the top-right corner of) a cell in one of these chains, $\chi_{\text{opt}}(S') \neq \chi_{\text{opt}}(S)$. In either case we get the recurrence relation

$$|\chi_{\text{opt}}(\{S_1, \ldots, S_m\}| \geq 1 + \min_{1 \leq k \leq m} \{\max\left(|\chi_{\text{opt}}(\{S_1, \ldots, S_{k-1})\}|, |\chi_{\text{opt}}(\{S_{k+1}, \ldots, S_m\})|\right\}\}$$

And so $|\chi_{\text{opt}}(\mathcal{S})| = \Omega(\log |\mathcal{S}'|)$, and the lemma follows. □

Let T be a non-bare orphan tile. For each of its v-neighbors, T_γ, $\gamma \in \Gamma$, let $\check{\mathcal{S}}_\gamma(T_\gamma)$ and $m_\gamma(T_\gamma)$ be as defined preceding Corollary 8. The selection of sub-chains of corner-chains according to Lemmas 19 and 21 implies the existence of a chain $\check{\mathcal{S}}_\gamma(T_\gamma)$ with respect to a region $Q \subseteq T$ whose length is $\Omega(\max_{T,\gamma}\{m_\gamma(T_\gamma)\})$. We apply Lemma 22 to $\check{\mathcal{S}}_\gamma(T_\gamma)$ and Q to obtain the following corollary:

Corollary 9 $|\chi_{\text{opt}}(\mathcal{S})| = \Omega\left(\log\left(\max_{T,\gamma}\{m_\gamma(T_\gamma)\}\right)\right)$.

Wrapping-up the Proof of Theorem 3 for unit-squares. Combining Corollary 8 and Corollary 9, and noting that the computational complexity of the algorithm is only due to sorting squares according their x coordinates, Theorem 3 for unit-squares directly follows.

7.4 General Rectangles

Consider a collection $\mathcal{R}$ of rectangles with size-ratio ρ. Our goal is to prove the existence of an efficient approximation algorithm for CF-coloring $\mathcal{R}$. The number of colors required by the algorithm is $O((\log \rho)^2 \cdot |\chi_{\text{opt}}(\mathcal{R})|)$, hence the approximation ratio is constant if ρ is constant.

By scaling separately the x-axis and the y-axis, we may assume that the minimum width and height of rectangles in $\mathcal{R}$ are equal to to 1. Hence, all side-lengths are in the range $[1, \rho]$.

The algorithm proceeds in two steps (as in the proof of Theorem 1). First, consider the case that $\rho \leq 2$. For this case we show that $O(|\chi_{\text{opt}}(\mathcal{R})|)$ colors suffice. For the more general case of $\rho > 2$, we partition the set of rectangles into $\log^2 \rho$ classes. For $0 \leq i, j < \log \rho$, class (i, j) consists of rectangles whose width is in the interval $[2^i, 2^{i+1})$, and whose height is in the interval $[2^j, 2^{j+1})$. Each class is colored using a distinct palette, to obtain a CF-coloring that uses $O((\log \rho)^2 \cdot |\chi_{\text{opt}}(\mathcal{R})|)$ colors, as required.

7.4.1 Rectangles with $\rho \leq 2$

We outline the algorithm for the case $\rho \leq 2$ below.

1. The tiling is the same as in the case of unit squares. The tiles are assigned 25 different palettes (instead of 9).

2. An orphan tile may now be completely covered by a rectangle. An orphan tile that is completely covered by a rectangle selects such a rectangle (this type of selection does not exist in the case of unit squares).

3. Instead of selecting closest rectangles from e-neighbors, every non-bare orphan tile that is not completely covered by a single rectangle selects the rightmost rectangle (if any) whose right edge intersects both the bottom and top side of the tile. The same selection takes place in the other 3 axis-parallel directions. In this stage an orphan cell selects at most 4 rectangles.

4. A non-bare orphan tile that still contains a region covered by $\mathcal{R}$ but not by the rectangles selected so far, selects rectangles from the corner-chains as in the algorithm for unit squares. The reason that the same techniques apply is that the intersection of a rectangle with a tile contains at most one corner.

5. The essential (selected) rectangles from each tile are colored as described in the following paragraph.

Coloring the Essential Rectangles. Given a non-orphan tile T, let $sel(T)$ denote the set of rectangles that belong to T and were selected in the previous stages. Let $m = \max_{T',\gamma}\{m_\gamma(T')\}$ denote the maximum (over all tiles T' and corner types γ) of the number of rectangles selected by an orphan tile T' due to their participation in a γ-corner-chain. In this section we show that: (i) $|sel(T)| \leq O(m)$, and (ii) $sel(T)$ can be CF-colored using $O(\log |sel(T)|)$ colors.

We begin with counting the number of rectangles in $sel(T)$. Since the side length of every rectangle is in the range $[1, 2]$, and all the rectangles in $sel(T)$ are centered in T, it follows that $\bigcup sel(T)$ intersects at most 25 tiles. Therefore, $|sel(T)| = O(m)$, for every tile T.

Similarly to Corollary 9, $|\chi_{\text{opt}}(\mathcal{R})| = \Omega(m)$. To obtain the constant-ratio approximation algorithm, we next show that $sel(T)$ can be CF-colored using $O(\log |sel(T)|)$ colors. Note that Theorem 4 is not applicable in this case since the rectangles are not congruent.

Lemma 23 *Let $\mathcal{R}'$ be a set of axis parallel rectangles with minimum width (height) at least* 1*. Assume that all centers of rectangles in $\mathcal{R}'$ reside in a square tile of side-length* $1/2$*. Then it is possible to CF-color $\mathcal{R}'$ using $O(\log(|\mathcal{R}'|)$ colors.*

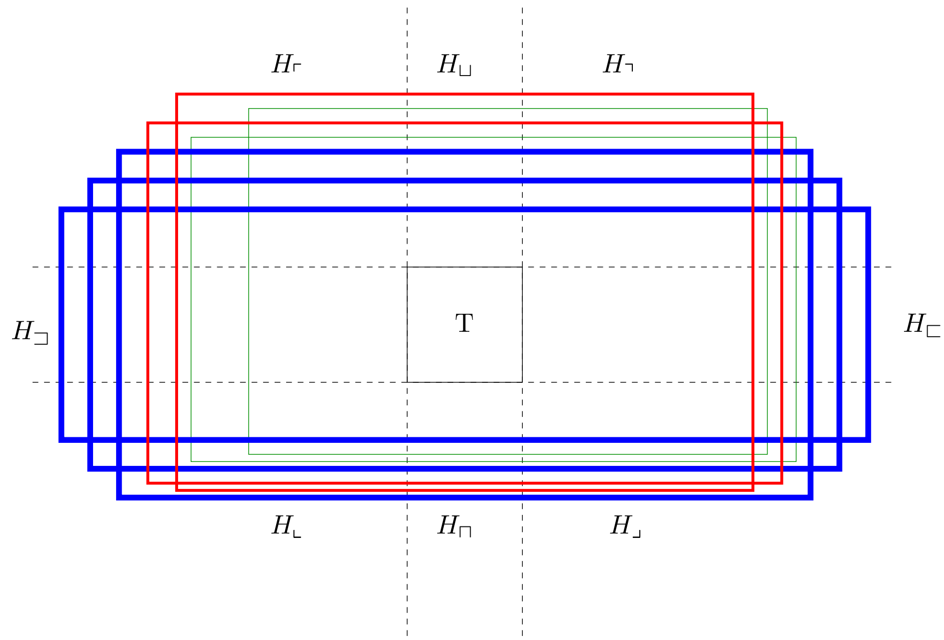

Figure 11: An illustration for the proof of Lemma 23. The thickest (blue) rectangles belong to all 4 chains. The second thickest (red) rectangles belong to the top-left and top-right chains, and the thinnest (green) rectangles belong only to the top-right chain.

Proof: Let T be the $1/2 \times 1/2$ tile that contains the centers of the unit squares in $\mathcal{R}'$. Extend the sides of T into lines, and consider the subdivision of the plane into 9 regions by these 4 lines. The subdivision consists of (i) the tile itself T, (ii) 4 corner regions denoted by $H_\llcorner, H_\lrcorner, H_\urcorner$, and $H_\ulcorner$, and (iii) 4 remaining regions denoted by $H_\sqcup, H_\sqcap, H_\sqsubset$, and $H_\sqsupset$. These regions are depicted in Fig. 11.

Since each of the 4 regions $H_\sqcup$, $H_\sqcap$, $H_\sqsubset$ and $H_\sqsupset$ is of height/width $1/2$, it suffices to select one rectangle for each and give it a unique color, in order to serve the intersection of $\mathcal{R}'$ with each of them. In particular, for $H_\sqcup$ we take the rectangle whose top edge has the largest y coordinate, for $H_\sqcap$ the rectangle whose bottom edge has the smallest y coordinate, and similarly for $H_\sqsubset$ and $H_\sqsupset$. Any one of these (at most) 4 rectangles can serve all of T as well.

Next we observe that in order to serve each of the 4 corner regions $H_\llcorner$, $H_\lrcorner$, $H_\ulcorner$ and $H_\urcorner$, it suffices to focus on 4 corner chains. Let $\widetilde{\mathcal{R}}'_\urcorner$, $\widetilde{\mathcal{R}}'_\ulcorner$, $\widetilde{\mathcal{R}}'_\lrcorner$ and $\widetilde{\mathcal{R}}'_\llcorner$, respectively, denote the set of rectangles that appear in the *envelope* of $\mathcal{R}'$ in each of the 4 corner parts. That is, those rectangles in $\mathcal{R}'$ whose corresponding corners ($\urcorner$ in $H_\llcorner$, $\ulcorner$ in $H_\lrcorner$ and in general γ in $H_{\mathrm{op}(\gamma)}$) are not contained in any other rectangle in $\mathcal{R}'$. The intersection of any other rectangle in $\mathcal{R}'$ with each H_γ, $\gamma \in \Gamma$, is contained in the intersection of the corresponding subset $\widetilde{\mathcal{R}}'_{\mathrm{op}(\gamma)}$ with H_γ. Note that the 4 subsets are not necessarily disjoint.

By a slight variant of Claim 17, each corner chain $\widetilde{\mathcal{R}}'_{\mathrm{op}(\gamma)}$ is indeed a chain with respect to H_γ. While it is possible to apply Lemma 15 to each of these chains, we do not directly obtain a single consistent coloring because the different chains are not necessarily disjoint. Instead, we partition the rectangles into $2^4 - 1 = 15$ disjoint subsets, where each subset consists of rectangles that belong a non-empty subset of corner chains (e.g., $\widetilde{\mathcal{R}}'_\urcorner$ and $\widetilde{\mathcal{R}}'_\lrcorner$ but not $\widetilde{\mathcal{R}}'_\ulcorner$ and $\widetilde{\mathcal{R}}'_\llcorner$).

The important observation regarding the envelope of $\mathcal{R}'$ is that if every boundary segment is given a symbol that corresponds to the rectangle it belongs to, then the sequence of symbols is a

Davenport-Schinzel sequence $DS(n, 2)$ [SA95]. Namely, no two consecutive symbols are equal and there is no alternating subsequence of length 4 (i.e., no "$\ldots a \ldots b \ldots a \ldots b \ldots$", for every pair of symbols $a \neq b$).

It follows that if two squares belong to more than 1 chain (that is to 2, 3, or even all 4 chains), then they appear *in the same order* (up to reversal) in all chains they belong to. Hence we can color each of the 15 subsets separately in a consistent manner (using 15 different palettes). The total number of colors used is hence $O(\log(|\widetilde{\mathcal{R}}|))$, as required. $\square$

8 Coloring Arrangements of Regular Hexagons

In this section we prove Theorem 4 for the case of regular hexagons. The proof follows the ideas used in the proof for the case of rectangles. We therefore we only sketch the details (with accompanying illustrations).

8.1 Preliminaries

The sets of regular hexagons that we consider are axis-parallel, namely, two of the sides of the hexagons are parallel to the x axis. The *type* of a vertex is determined by the slope of the segment connecting the center of the hexagon with the vertex. (See Fig. 12.) In the same fashion, we define the type of an edge of the hexagon.

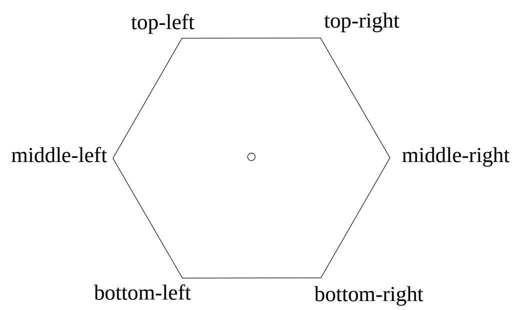

Figure 12: A hexagon and its vertices.

The Tiling. In the case of hexagons we consider a tiling of the plane by equilateral triangles with unit side-lengths; one side of each triangular tile is horizontal (see Figure 13). Triangular tiles have two possible orientations: in the *up* orientation, the vertex opposite the horizontal edge is above that edge, and in the *down* orientation that vertex is below the horizontal edge.

We adapt the notation of Section 7 as follows. The set of hexagons is denoted by $\mathcal{H}$. We assume that the side length of every hexagon is in the range $[1, \rho]$. For a tile T, we let $\mathcal{H}(T)$ denote the set of hexagons in $\mathcal{H}$ that belong to T (that is, whose center resides in T). A tile T is an *orphan* if $\mathcal{H}(T) = \emptyset$, and it is *bare* if no hexagon in $\mathcal{H}$ intersects it.

Since the tiles are equilateral triangles of side length 1, the following holds for any set of hexagons $\mathcal{H}$ with side-lengths at least 1.

Observation 1 *For every tile T and hexagon $H \in \mathcal{H}$: (i) if $H \in \mathcal{H}(T)$ then $T \subset H$, (ii) T contains at most one vertex of H. (iii) If T intersects two edges e_1, e_2 of a hexagon H, then these edges are adjacent and T contains also the vertex $e_1 \cap e_2$.*

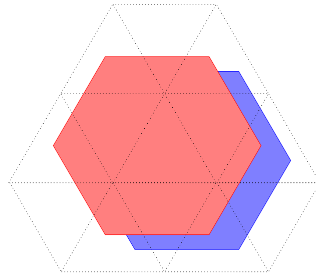

Figure 13: The triangular tiling and a pair of hexagons. The tile borders are depicted by dotted lines. Both centers of the hexagons are in the middle tile. Both hexagons completely cover the middle tile. Larger hexagons may intersect more than the 12 neighboring tiles.

Disjoint palettes. As in the case of disks (c.f. Theorem 1) we reduce the problem to the case of size-ratio 2 by paying a factor of $\log \rho$. Henceforth, we assume that $\rho \leq 2$. We assign a palette to every tile T. The colors assigned to hexagons in $\mathcal{H}(T)$ belong to the palette assigned to T. Palettes are disjoint and the distribution of palettes is such that intersecting hexagons from different tiles are assigned different colors. This requires only a constant number of palettes. For example, consider a tiling of the plane with hexagonal super-tiles that contain a constant number of triangular tiles. Assign every triangular tile within a hexagonal super-tile a different palette, and extend this coloring periodically according to the hexagonal super-tiles. The resulting assignment of palettes is as required.

8.2 Coloring arrangements of hexagons

As in the case of rectangles, the algorithm has two stages. In the first stage, each non-bare orphan tile T selects a subset of hexagons whose union serves the covered regions in T. This stage is somewhat more involved in the case of hexagons than in the case of rectangles. For each non-orphan tile T, let $sel(T)$ denote the the subset of hexagons in $\mathcal{H}(T)$ that were selected by orphan tiles in the first stage. In the second stage, the hexagons in $sel(T)$ are CF-colored, for every tile non-orphan T, using colors from the palette assigned to T.

8.2.1 Selection of hexagons by non-bare orphan tiles

Consider a non-bare orphan tile T. If there exists a hexagon that covers all of T, then we simply select one of these hexagons to serve it and no more selections are required. We now consider non-bare orphan tiles that are not covered by a single hexagon.

For an edge type e, let $\mathcal{H}(T, e)$ denote the set of hexagons that intersect T with an edge that is of type e (i.e. non-empty intersection, but no vertex of the hexagon is contained in T). We claim that a single hexagon covers the intersection of T with hexagons from $\mathcal{H}(T, e)$. For example, let e denote the top horizontal edge. The set of tiles that intersect T with their top horizontal edge is denoted by $\mathcal{H}(T, e)$. Among these hexagons, pick the hexagon H with the highest center. The hexagons H covers the intersection of T with every hexagon in $\mathcal{H}(T, e)$. This completes the discussion of the selection of hexagons that intersect T with edge.

We denote by Q the region contained in T that remains after this choice of at most 6 hexagons (on per edge type). Note that if Q is non-empty, then Q a polygon with at least 3 edges and at most 6 edges. The edges of Q are parallel to those of the hexagons in $\mathcal{H}$.

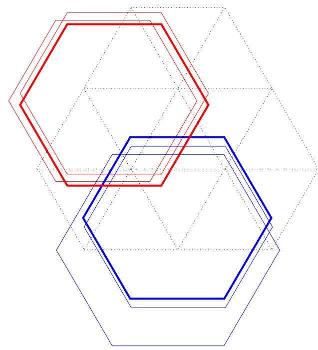

Figure 14: Let T be the central triangular tile in the figure, which is an orphan tile. The figure illustrates the choice of hexagons that intersect T with an edge. The selected hexagons are the two thick hexagons.

We now consider the selection of hexagons that intersect T with a vertex. Let γ denote a vertex type (e.g. top-right), and let $\mathcal{H}(T,\gamma)$ denote the set of hexagons whose γ-vertex is in T. Amongst the hexagons in $\mathcal{H}(T,\gamma)$ let $\widetilde{\mathcal{H}}(T,\gamma)$ denote the hexagons that participate in the envelope of $\mathcal{H}(T,\gamma)$ in Q. Similarly to the analysis in the case of rectangles, the latter cover all of the intersection of Q with the former, and furthermore, they constitute a (corner) chain with respect to Q. We refer to the chain in terms of the vertex type (e.g. top-right chain). For an illustration see Figure 15.

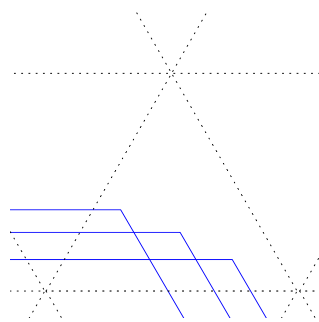

Figure 15: An example of a top-right chain.

Thus, there are at most 6 corner chains intersecting Q, one for each vertex type. Here we have three types of "interactions" between chains depending on the distance between the corresponding vertex types on the hexagons - that is: distance-one (e.g. top-right and top-left), distance-two (e.g. top-right and middle-left), and distance 3 (e.g. top-right and bottom-left).

Interactions between distance-one and distance-two chains. Interactions between distance-one chains and distance-two chains are analogous to the interactions between corner chains of adjacent corners in the case of rectangles. Specifically, for each such pair of chains, we can select a single hexagon from each chain so that: (1) The union of the two selected hexagons covers the intersection between the chains; (2) The remaining hexagons (not covered by the two selected hexagons) constitute disjoint chains with respect to Q minus the two hexagons. For an illustration see Figure 16.

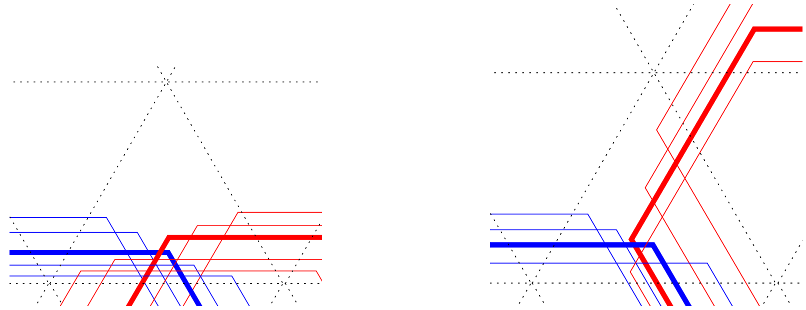

Figure 16: The two "benign" interactions between chains. On the left side is a distance-one interaction between a top-right chain and a top-left chain. On the right is a distance-two interaction between a top-right chain and a middle-left chain. In each figure, the two bold hexagons are those selected from the two chains.

It follows that by selecting at most 4 hexagons from each of the 6 chains that intersect Q, it is possible to service all areas of intersections between such pairs of subsets of hexagons.

Interactions between pairs of distance-three (opposite) chains. The case of interactions between distance-three chains is analogous to the interaction between opposite chains in the case of rectangles. In particular it is possible to select an approximately minimal subset of hexagons from the two chains so as to serve all the area in their union (within the remaining region in T). For an illustration see Figure 17.

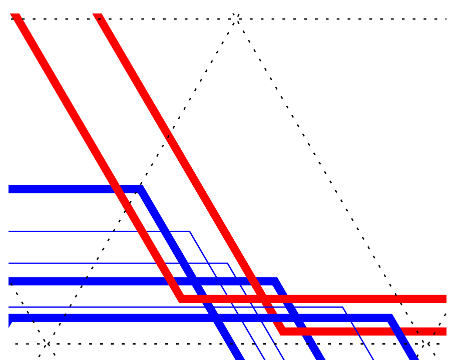

Figure 17: An interaction between the distance-three (opposite) chains top-right and bottom-left. The selected hexagons are bold.

8.3 Coloring the selected hexagons

We now return to each non-orphan tile T, and assign colors to the hexagons requested from it. Note that Theorem 4 is not applicable since the hexagons are not congruent.

Lemma 24 *Let $\widetilde{\mathcal{H}}$ be a subset of axis aligned hexagons with side-lengths at least 1, that all belong to the same tile T. Then it is possible to CF-color $\widetilde{\mathcal{H}}$ using $O(\log(|\widetilde{\mathcal{H}}|)$ colors.*

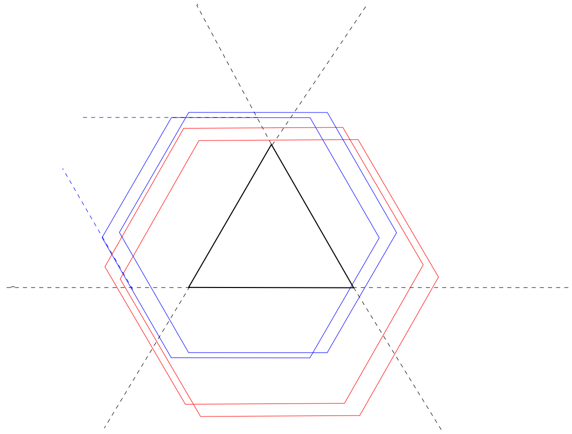

Figure 18: The partitioning of the area covered by hexagons that belong to the same tile.

Proof sketch: First, we assume, without loss of generality, that every hexagon in $\widetilde{\mathcal{H}}$ participated in the envelope (i.e. contains a vertex in $\bigcup\widetilde{\mathcal{H}}$).

Similarly to the proof of Theorem 3 for rectangles, we extend the sides of a tile to partition the area covered by $\widetilde{\mathcal{H}}$ into several subregions (see Figure 18). The number of resulting regions is 7 (including the tile T itself which is covered by every hexagon in $\widetilde{\mathcal{H}}$). Three of these subregions have a common vertex with T (and are referred to as the "angular" subregions) and three have a common edge (and are referred to as the "trapeze" subregions). The vertices of every hexagon in $\widetilde{\mathcal{H}}$ are in the trapeze subregions. Hence, it is possible to select at most 3 hexagons to serve the angular subregions. Each of these hexagons are assigned a unique color (and thus T itself is also served).

We now deal with serving points in the trapeze regions. We wish to identify two chains in each trapeze region. Fix a trapeze region R . Every hexagon has two adjacent vertices in the trapeze region (as well as the edge connecting these vertices). Let u and v denote the vertex types that appear in R. Pick the hexagon H_R whose edge is farthest away from the corresponding edge of the triangular tile. Consider the sequence of vertices along the envelope of $\widetilde{\mathcal{H}}$ in R. This sequence starts with a block of vertices of type u and ends with a block of vertices of type v. The two vertices of H_R in R appear consecutively in this envelope. By picking H_R and assigning it a unique color, the envelope in R is separated in two parts. Moreover, the region $(R \setminus H_R) \cap \bigcup\widetilde{\mathcal{H}}$ consists of two disjoint connected parts. The hexagons whose vertices appear in the envelope in each part are chains with respect to $R \setminus H_R$. Thus, by picking at most 6 hexagons and assigning them unique colors, we have identified 6 disjoint chains.

As in the proof of Lemma 23, hexagons that belong to multiple chains appear in the same order (up to reversal) in these chains. Hence we partition the hexagons that appear in chains into at most $2^6 - 1$ subsets, where within each subset all hexagons belong to the same chains. (A finer counting argument is based on showing that for every 3 or more chains, there can be at most one hexagon that belongs to all these chains. Hence we actually focus on subsets of hexagons that belong to one or two chains.) Each such subset is provided with a disjoint palette and can be colored using a logarithmic (in its size) number of colors. $\square$

Finally, the proof of Theorem 3 for regular hexagons follows by combining the above lemma with a lower bound analogous to Lemma 22, the basic properties of the tiling, and the requesting process from orphan tiles.

9 Consequences

9.1 Universal bounds for non-congruent rectangles and hexagons

As a corollary of Lemma 23 we also obtain a universal bound for the case of rectangles that is analogous to the case of disks (Part 1 of Theorem 1).

Theorem 10 *There exists an algorithm that given a set $\mathcal{R}$ of axis-parallel rectangles with side lengths in the interval $[1, \rho]$, finds a CF-coloring χ of $\mathcal{R}$ using $O((\log \rho)^2 \cdot \log(\phi(\mathcal{R})))$ colors. Here $\phi(\mathcal{R})$ is the maximum number of rectangles whose centers (i.e., intersection of their diagonals) all reside in a square of side-lengths $1/2$.*

Lemma 24 implies an analogous theorem also for hexagons.

9.2 CF-Multi-coloring

An interesting byproduct of Theorem 3 and its analysis has to do with minimum *CF-Multi-coloring*. A CF-multi-coloring of a collection $\mathcal{S}$ is a mapping χ from $\mathcal{S}$ to *subsets* of colors. The requirement is that for every point $x \in \bigcup_{S \in \mathcal{S}} S$, there exist a color i such that $\{S : x \in S, i \in \chi(S)\}$ contains a single subset. It has been observed by Bar-Yehuda ([B01], based on [BGI92]), that every set-system $(X, \mathcal{S})$ can be CF-multi-colored using $O(\log |X| \cdot \log |\mathcal{S}|)$ colors. Since the problem of minimum graph coloring can be reduced to CF-coloring of set-systems, it follows that there exist set-systems for which there is an exponential gap between the minimum number of colors required in a CF-coloring and the minimum number of colors required in a CF-multi-coloring. In particular, this is true when the set system $(X, \mathcal{S})$ corresponds to a clique $G = (V, E)$ as follows: there is a set S_v for every vertex $v \in V$, and there is a point $x_e \in X$ for every edge $e \in E$. The set S_v contains the point x_e if and only if v is an endpoint of e. The number of colors required to CF-color this set system is $|\mathcal{S}| = |V|$ in contrast to $O(\log^2 |\mathcal{S}|)$ colors that are sufficient for CF-multi-coloring.

A natural question is whether in the geometric setting that we study, the number of colors required for CF-multi-coloring is significantly smaller than that required for CF-coloring. An example in which CF-multi-coloring saves colors is a “circle” of 5 congruent squares, such that every adjacent pair of squares intersect, and no 3 squares intersect. Since the number of squares is odd, 3 colors are needed for CF-coloring. However, CF-multi-coloring requires only 2 colors: color the first square with 2 colors, and then color the rest of the squares with alternating colors. The lower bound proved in Lemma 14 also applies to CF-multi-coloring, hence CF-multi-coloring does not save colors in chains. Furthermore, it follows from our analysis (c.f., Lemma 22) that CF-multi-coloring reduces the number of colors by at most a constant in the case of congruent squares (or hexagons).

Theorem 11 *Let $\mathcal{S}$ denote a set of congruent axis-parallel squares, and let $\chi_{\text{opt}}^{\text{multi}}(\mathcal{S})$ denote an optimal CF-multi-coloring of $\mathcal{S}$. Then $|\chi_{\text{opt}}^{\text{multi}}(\mathcal{S})| = \Theta(|\chi_{\text{opt}}(\mathcal{S})|)$.*

Acknowledgments

We thank an anonymous FOCS reviewer for suggesting that the NP-completeness of coloring intersection graphs of unit disks could be used to prove the NP-completeness of CF-coloring arrangements of unit disks. We thank Micha Sharir for helpful discussions.

Guy Even and Dana Ron were partly supported by the LSRT (Large Scale Rural Telephony) Consortium.

References

[AHK+01] K. Aardal, S. van Hoesel, A. Koster, C. Mannino and A. Sassano, "Models and Solution Techniques for the Frequency Assignment Problem", ZIB Report 01-40, December 2001. `http://www.zib.de/PaperWeb/abstracts/ZR-01-40/`.

[AKM+01] Z. Abrams, J. Könemann, A. Meyerson, K. Munagala and S. Plotkin, "Facility location with interference", manuscript 2001.

[BGI92] R. Bar-Yehuda, O. Goldreich and A. Itai, "On the time-complexity of broadcast in multi-hop radio networks: An exponential gap between determinism and randomization", Journal of Computer and System Sciences, 45:104–126, 1992.

[B01] Reuven Bar-Yehuda, personal communication, 2001.

[BKOS97] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwartzkopf, "Computational Geometry - Algorithms and Applications", Springer Verlag, 1997.

[CCJ90] B. N. Clark, C. J. Colburn and D. S. Johnson, "Unit Disks Graphs", Discrete Mathematics 86:165–177, 1990.

[C69] Coxeter, H. S. M., "Introduction to Geometry", 2nd ed. New York: Wiley, 1969.

[DBJC98] N. W. Dunkin, J. E. Bater, P. G. Jeavons and D. A. Cohen, "Towards High Order Constraint Representations for the Frequency Assignment Problem", Technical Report CSD-TR-98-05, June, 1998. `http://www.dcs.rhbnc.ac.uk/research/constraints/publications/index.shtml`.

[FK98] U. Feige and J. Kilian, "Zero knowledge and the chromatic number", J. Comput. System Sci. 57, 187-199, 1998.

[GGRV00] M. Galota, C. Glasser, S. Reith, and H. Vollmer. "A polynomial-time approximation scheme for base station positioning in UMTS networks", Technical Report 264, Institut für Informatik, Universität Würzburg, 2000.

[HS02] S. Har-Peled and S. Smorodinsky, "On Conflict-Free Coloring of Points", in preparation, 2002.

[H01] Xuemin Huang, "Automatic cell planning for mobile network design: optimization models and algorithms", dissertation, Universität Karlsruhe, 2001.

[KMR01] S.O. Krumke, M.V. Marathe and S.S. Ravi, "Models and Approximation Algorithms for Channel Assignment in Radio Networks," invited paper in *Wireless Networks*, 7(6), pp. 575-584, 2001.

[MBH+95] M.V. Marathe, H. Breu, H.B. Hunt III, S.S. Ravi and D.J. Rosenkrantz, "'Simple Heuristics for Unit Disk Graphs," in *Networks*, (25), pp. 59–68, 1995.

[PT02] J. Pach and G. Tóth, "Relaxing planarity for topological graphs", *Discrete and Comput. Geom.* to appear.

[SA95] M. Sharir and P. Agarwal, "Davenport-Schinzel sequences and their geometric applications", Cambridge University Press, 1995.

[SM02] S. Smorodinsky, "Combinatorial Problems in Computational Geometry", Ph.D dissertation, Tel-Aviv University, in preparation, 2002.