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

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

July 20, 2003

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 noncongruent), 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 $\chi : \mathcal{S} \to \mathbb{N}$ 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 closed 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 finite 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. For each $i \geq 1$, let $\mathcal{S}^i$ denote the subset of disks in $\mathcal{S}$ whose radius is in the range $[2^{i-1}, 2^i)$. Let $\phi_{2^i}(\mathcal{S}^i)$ denote the maximum number of centers of disks in $\mathcal{S}^i$ that are contained in a $2^i \times 2^i$ square. We refer to $\phi_{2^i}(\mathcal{S}^i)$ as the *local density* of $\mathcal{S}^i$ (with respect to $2^i \times 2^i$ square). 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 $O\left(\min\left\{\sum_{i=1}^{\log(\rho)+1}(1+\log\phi_{2^i}(\mathcal{S}^i)), \log|\mathcal{S}|\right\}\right) = O\left(\min\left\{\log(\rho)\cdot\max_i\{\log\phi_{2^i}(\mathcal{S}^i)\}, \log|\mathcal{S}|\right\}\right)$ colors*[1]*.*

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: if we color $\mathcal{S}_r(X)$ by assigning each disk $D \in \mathcal{S}_r(X)$ the color of its center, then this is a valid 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 $\phi_1(\mathcal{S}) = n$ for which $\Omega(\log n)$ colors are necessary in every CF-coloring of $\mathcal{S}$.

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 polynomial-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]For simplicity of notation, we avoid writing $\log(x+1)$ throughout the paper, even when x may equal 1, and consider $O(0)$ to be $O(1)$.

1.1.2 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 $w(R)$ ($h(R)$, respectively) denote the width (height, respectively) 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 optimal CF-coloring of $\mathcal{R}$.*

1. *If $\mathcal{R}$ is a set of rectangles, then there exists a polynomial-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 regular hexagons, then there exists a polynomial-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-ratio 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 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 (see Figure 1). 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 $\{C_{r,O}(x)\}_{x\in X}$.

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 if we color each $c \in \mathcal{C}_{r,O}(X)$ with the color of its center, then this is a valid CF-coloring of $\mathcal{C}_{r,O}(X)$, for every scaling factor r.*

A polynomial-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 polynomial-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:

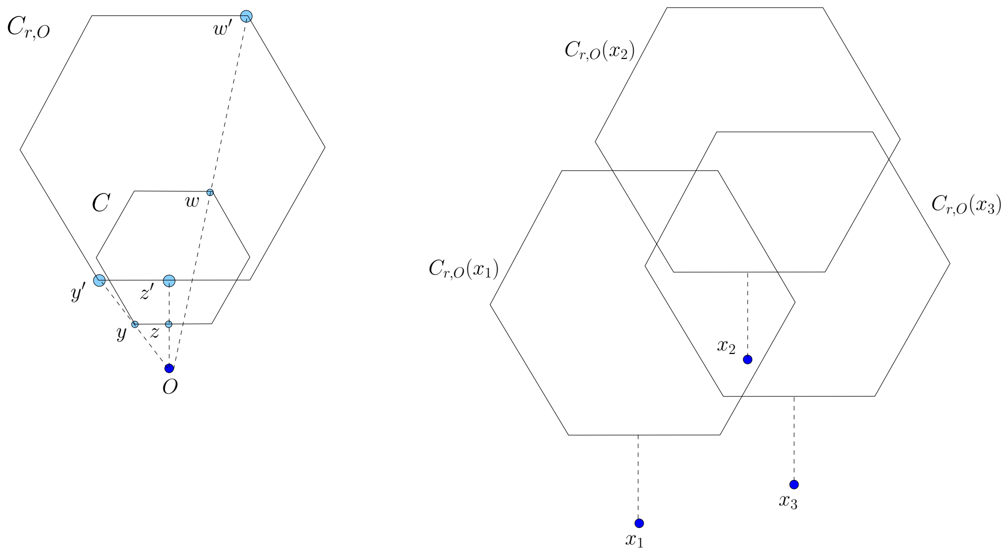

Figure 1: On the left is an example of a scaled translation $C_{r,O}$ of a regular hexagon C, with respect to the point O, where the scaling factor is $r = 2$. The points y, z and w on the small hexagon C are mapped to the points y', z', and w', respectively, on the larger hexagon $C_{r,O}$. The dashed lines correspond to the rays emanating from O towards the points y, z, and w. On the right is an additional set $X = \{x_1, x_2, x_3\}$ of three points and the corresponding set $\mathcal{C}_{r,O}(X) = \{\mathcal{C}_{r,O}(x_1), \mathcal{C}_{r,O}(x_2), \mathcal{C}_{r,O}(x_3)\}$.

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 $(X, \mathcal{R})$. When X is a set of points in the plane and $\mathcal{R}$ is the set of ranges obtained by intersections with disks, this graph is the standard Delaunay graph. 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 $[i, j]$, $1 \leq i \leq j \leq |\mathcal{S}|$, there exists a point $x \in \bigcup_{S \in \mathcal{S}} S$, such that the subcollection of subsets that contains the point x equals the subcollection 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 8. We show that chains of unit disks (respectively, unit squares and hexagons) are tight examples of Theorem 1 (respectively, 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 14 and 18) 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* $p_{v,w} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{R}$, 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 nonbinary 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 $O\left(\min\left\{\sum_{i=1}^{\log(\rho)+1}\log\phi_{2^i}(\mathcal{S}^i), \log|\mathcal{S}|\right\}\cdot\delta_{\min}\right)$ 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 notations 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 (*cells*$(\mathcal{R}), \mathcal{R})$ as a range space. To be precise, (*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 (*cells*$(\mathcal{R}), \mathcal{R})$ as the *combinatorial arrangement* induced by $\mathcal{R}$; we denote this combinatorial arrangement by $\mathcal{A}(\mathcal{R})$.

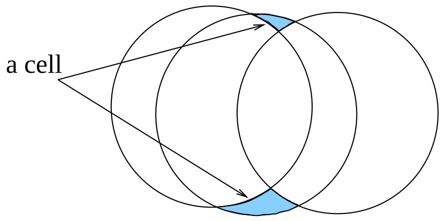

Figure 2: An arrangement of disks. The marked cell corresponds to the regions that are contained in the middle disk, and only in that disk.

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 $V \subset cells(\mathcal{R})$. 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 between 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 the 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 $\mathcal{R}$. For a point x in the plane, let $C(x)$ denote the region congruent to C that is centered at x. For a set X of points in the plane, let $\mathcal{C}(X) \triangleq \{C(x) : x \in X\}$ 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 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 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 Ranges

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 this subsection by presenting a high level description of our algorithm. We then shortly discuss how it can be applied to the special case where X is a set of n points in the plane and the ranges in $\mathcal{R}$ are intersections of X with disks.

The algorithm works in an iterative manner, where in iteration i is selects the subset of points that are colored by color number i. Let X_i denote the set of points that are colored by the color i, and let $X_{<i}$ (respectively $X_{\leq i}$) denote the set $\bigcup_{j<i} X_j$ (respectively $\bigcup_{j\leq i} X_j$). When determining X_i, the algorithm ensures that the following condition holds:

> For every range $S \in \mathcal{R}$, either (i) S can be served by a point colored $j < i$ (i.e. $\exists j < i : |S \cap X_j| = 1$), or (ii) $S \cap X_i$ contains at most one point, or (iii) S contains a point that is not colored yet (i.e. $S \nsubseteq X_{\leq i}$).

Correctness follows because if either (i) holds or (ii) holds then S can be served by a point colored $j \leq i$, while if both (i) and (ii) do not hold, then there will be a point colored by a color greater than i that can serve S. In fact, a coloring that obeys the above condition has the following property: For every range $S \in \mathcal{R}$, the highest color of a point contained in S has multiplicity 1.

Observe that we can trivially obey the above condition by selecting X_i to consist of a single point in $X \setminus X_{<i}$, so that each point is colored by a different color. However the total number of colors in this case is $|X| = n$ while we are interested in using only $O(\log n)$ colors. To obtain $O(\log n)$ colors, we show that (for the sets of ranges we consider), in each stage it is possible to select at least a constant fraction of the remaining points (i.e., $|X_i| \geq \frac{1}{4} \cdot |X \setminus X_{<i}|$).

To make the above more concrete, consider the special case of coloring a set of points X with respect to disks. First assume that the points all lie on a straight line. In such a case, the choice of X_i is simply to pick every other point of $X \setminus X_{<i}$ (see Figure 3). 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. Since $|X_i| = \left\lceil \frac{|X \setminus X_{<i}|}{2} \right\rceil$, the number of colors used is $O(\log n)$ as desired.

The choice of X_i when the points are in general position in the plane is more involved. In this case, we construct the Delaunay graph G_i of the set $X \setminus X_{<i}$: Two points p_i and p_j form an edge in G_i if and only if there is a closed disk D that contains p_i and p_j on its boundary and does not contain any other point in $X \setminus X_{<i}$. The graph G_i is planar, and hence is 4-colorable. The largest color class contains at least $\frac{1}{4}$ of the remaining points and is an independent set in G_i. In Claim 3 we prove that the largest color class is a good candidate for X_i.

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.

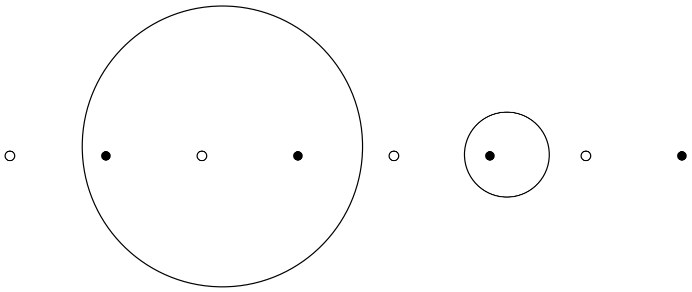

Figure 3: Selection of X_i when the points lie on a straight line. The points drawn are those in $X \setminus X_{<i}$. The points of X_i are denoted by filled dots. The unfilled dots denote points in $X \setminus X_{\leq i}$. The disk on the left contains two points in X_i and hence also an unfilled dot. The disk on the right contains only a single point in X_i.

Definition 4 *Let X be a set of points and $\mathcal{R} \subseteq 2^X$ a set of ranges. A partition (X_1, X_2) of X is* $\mathcal{R}$-useful *if $X_1 \neq \emptyset$ and*

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

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

Initialization: $i \leftarrow 1$, $X^1 \leftarrow X$, $\mathcal{R}^1 \leftarrow \mathcal{R}$. ($i$ denotes an unused color, X^i is the set of points not yet colored, and $\mathcal{R}^i$ is the set of ranges that contain more than one point in X^i and cannot be served by points colored j, for $j < i$.)
while $X^i \neq \emptyset$ **do**
  Find an $\mathcal{R}^i$-useful partition (X_1, X_2) of X^i. (See Claim 3 below.)
  Color: $\forall x \in X_1 \ : \ \chi(x) \leftarrow i$.
  Project: $X^{i+1} \leftarrow X_2$ and $\mathcal{R}^{i+1} \leftarrow \{S \cap X_2 : S \in \mathcal{R}^i,\ |S \cap X_1| \neq 1 \text{ and } |S \cap X_2| \geq 2\}$.
  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 $X^i \cap S \in \mathcal{R}^i$. In other words, in the ith iteration, the $\mathcal{R}^i$-useful partition (X_1, X_2) of X^i satisfies either $|X_1 \cap S| = 1$ or $|X_2 \cap S| = 1$. In the first case, S can be served by the single element $x \in X_1 \cap S$ (which is colored i). In the second case, S can be served by the single element $x \in X_2 \cap S$ (which is colored j, for $j > i$). Observe that if at the end of iteration i the range space $\mathcal{R}^{i+1}$ becomes empty while X^{i+1} is not empty, then the partition $(X^{i+1}, \emptyset)$ is trivially $\mathcal{R}^{i+1}$-useful, and all the points in X^{i+1} can be colored with the color $i + 1$. $\square$

Note that Algorithm CF-color computes a CF-coloring in which every range $S \in \mathcal{R}$ is served by the point with the highest color among the points in S.

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

If in every iteration i we have $|X_1| = \Omega(|X^i|)$, then Algorithm CF-color uses $O(\log |X|)$ colors. We formalize a condition guaranteeing that $|X_1|$ is a constant fraction of $|X^i|$. The condition is phrased in terms of a special graph that is a attached to the range space $(X^i, \mathcal{R}^i)$.

We refer to ranges $S \in \mathcal{R}^i$ as *minimal* if they are minimal with respect to inclusion. Recall that we initially assume that for every $S \in \mathcal{R}$, $|S| \geq 2$ (since ranges of size one are served trivially). The algorithm ensures that, in each iteration, every range in $\mathcal{R}^i$ contains at least two points. Hence, minimal ranges contain at least two points.

Definition 5 *A* Delaunay graph *of a 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* $E = \{e(S) : S \in \mathcal{R}$ *and* S *is minimal* $\}$.

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.

We now discuss how Definition 5 is an extension of the standard definition of the Delaunay graph of a set of points in the plane. The Delaunay graph of a set of points X in the plane is defined as the the dual graph of the Voronoi diagram of X [BKOS97]. Theorem 9.6 in [BKOS97] suggests an equivalent definition: Two points p_i and p_j form an edge in the Delaunay graph of X if and only if there is a closed disk D that contains p_i and p_j on its boundary and does not contain any other point in X. This equivalent definition implies that the edge set of a Delaunay graph corresponding to X equals the set of minimal ranges containing two points induced by disks. We leave it as an exercise to prove that every minimal range induced by a disk contains exactly two points. Hence, in the case of points in the plane and disks, $DG_{\mathcal{R}}(X, E)$ is the standard Delaunay graph of X.

The next claim shows how an $\mathcal{R}$-useful partition can be found by Algorithm CF-color.

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 partition 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 (i.e., $|S \cap X_1| \neq 1$) 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* $(X^i, \mathcal{R}^i)$ *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 $|X_1| = \Omega(|X^i|)$. The existence of a large independent set X_1 in the Delaunay graph $DG_{\mathcal{R}^i}(X^i, E)$ follows from the planarity of $DG_{\mathcal{R}^i}$. Planarity implies that the graph is 4-colorable [AH77a, AH77b], and therefore, the largest color-class is an independent set of size at least $|X^i|/4$. (Recall that planar graphs can be 4-colored in polynomial time [AH77a, AH77b, RSST96]. For our purposes, a coloring using 6 colors suffices. One could

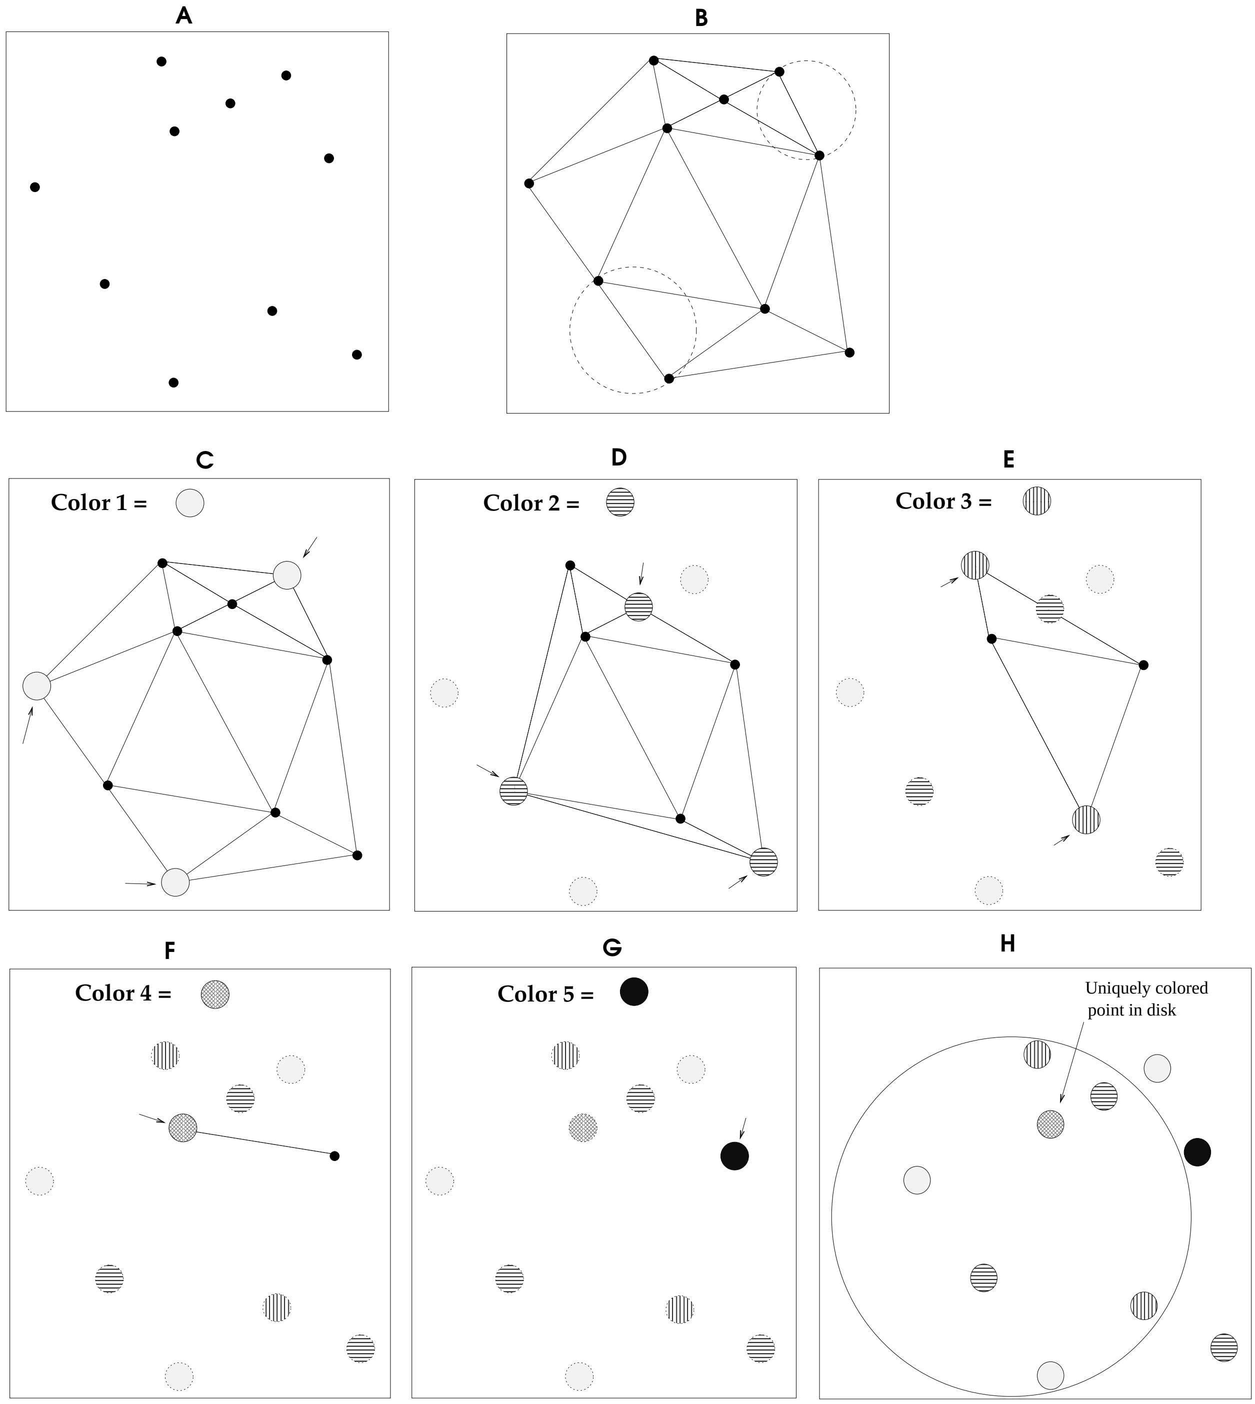

Figure 4: An illustration of the execution of Algorithm 1 in the case of disks in the plane. In Part A we see the given set of points. In Part B, the Delaunay graph is depicted. Recall that there is an edge between every pair of points p, q that are separated from the rest of the points by a disk. Two such disks are depicted for this graph. Parts C–G depict 5 steps of the algorithm. In each step we see the Delaunay graph over the remaining uncolored points, where the newly colored points are marked by arrows. (The previously colored points also appear in the figure, but they are not part of the Delaunay graph). In Part H we see the final coloring of all points, and an example of a disk and the point that can serve it (in general, there may be more than one such point).

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 $|X^i|/6$.) $\square$

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-spaces in $\mathbb{R}^3$, and scaled translations of a convex polygon, the corresponding Delaunay graphs are computable in polynomial time, which implies that Algorithm CF-color is polynomial.

3.3 Disks in the Plane

Recall that in the case of disks in the plane, the Delaunay graph that we attach to the set system $(X, \mathcal{R})$ is the standard Delaunay graph. Hence, the Delaunay graph is planar [BKOS97, Thm 9.5]. We may now apply Claim 4 to obtain the following lemma.

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 (closed) disk. Then it is possible to color X with respect to $\mathcal{R}$ using $O(\log n)$ colors.*

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.*

Let $CH(X)$ denote the convex hull of X. We make the following simplifying assumption: every point in X is an extreme point of $CH(X)$. If not, then all the points of X that are not extreme points of $CH(X)$ may be colored by a unique "passive" color. The coloring of nonextreme points by a passive color means, in effect, that these 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 $CH(X)$ 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 $X \cap H^+ = \{u, v\}$. Recall that we assumed that every point in X is an extreme point of $CH(X)$. Two points $x, y \in X$ are *adjacent* if there exists a supporting plane H of $CH(X)$ such that $H \cap X = \{u, v\}$. The *skeleton graph* $G' = (X, E')$ of $CH(X)$ is the graph over the points in X with edges between adjacent points. The skeleton graph is drawn on the boundary of $CH(X)$ using straight lines without crossings. Since the boundary of $CH(X)$ is homeomorphic to a sphere, it follows that the skeleton graph 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 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 $C' \sim C$.*

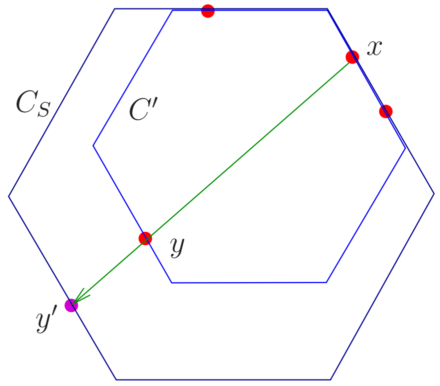

Figure 5: An illustration for the proof of Claim 8.

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 $S' \subsetneqq S$ 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 $y \in S \cap \mathring{C}_S$. The region C' is the image of C with respect to the homothecy τ that is defined as follows. 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. For an illustration, see Figure 5. $\square$

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 subsegment 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 $cells(\mathcal{S}_r(X))$. 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 (respectively, on the boundary of, outside) a disk S, if and only if the point $S^* \in \mathbb{R}^3$ lies above (respectively, on, below) the plane p^*. Indeed, assume that point $p = (a, b)$ lies inside (respectively, 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 $-2ax - 2by + a^2 + b^2 < r^2 - x^2 - y^2$ (respectively, 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 $\sum_{i=1}^{\log(\rho)+1} \log \phi_{2^i}(\mathcal{S}^i)$. Recall that ρ is the size-ratio of $\mathcal{S}$, $\mathcal{S}^i$ is the subset of disks in $\mathcal{S}$ whose radius is in the range $[2^{i-1}, 2^i)$, and and $\phi_{2^i}(\mathcal{S}^i)$ is the maximum number of disks in $\mathcal{S}^i$ whose centers reside in a common $2^i \times 2^i$ square. 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 2×2-square tiles. 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 9 different tiles. We assign a *palette* (i.e., a subset of colors) to each tile using 9 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 3×3 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 $|\mathcal{S}(T)| \leq \phi_2(\mathcal{S})$ 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 9, we get the desired upper bound of $O(\log \phi_2(\mathcal{S}))$ 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, denoted $\mathcal{S}^i$, 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 $O(\log \phi_{2^i}(\mathcal{S}^i))$ colors. By using a different (super-)palette per class, we obtain the desired bound on the number of colors, i.e., $\sum_{i=1}^{\log \rho+1} O\left(\log \phi_{2^i}(\mathcal{S}^i)\right)$.

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 $\phi_1(\mathcal{S}) \leq |\mathcal{S}|$. Since a unit square 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 $\phi_1(\mathcal{S}) = O(\frac{1}{d_{\min}(\mathcal{S}(T))^2})$. □

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. □

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 6. 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. □

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 $A_1(X)$.*

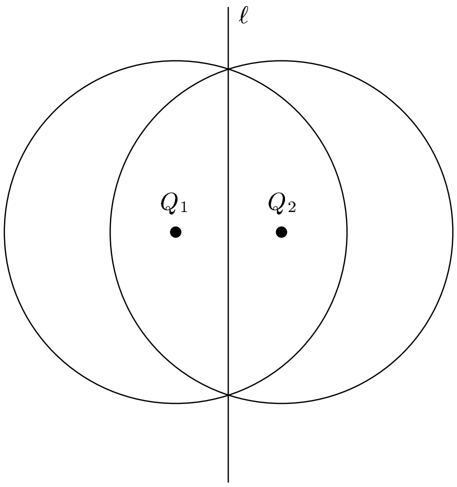

Figure 6: Proof of Lemma 11.

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 $\mathit{sect}(Q, \alpha, 1)$ of $A_1(X)$ (where $Q \in X$), such that

$$P \in \mathit{sect}(Q, \alpha, 1+\varepsilon) - \mathit{sect}(Q, \alpha, 1) \ .$$

It follows that

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

where $\mathit{sect}(Q, \alpha, 1)$ 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 $\mathit{cells}(\mathcal{C}_{r,O}(X))$. 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 regions in the plane $\mathcal{S}$. 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. We refer to a set $\{i, i+1, \ldots, j\}$ of consecutive integers as an *interval* and denote it by $[i, j]$.

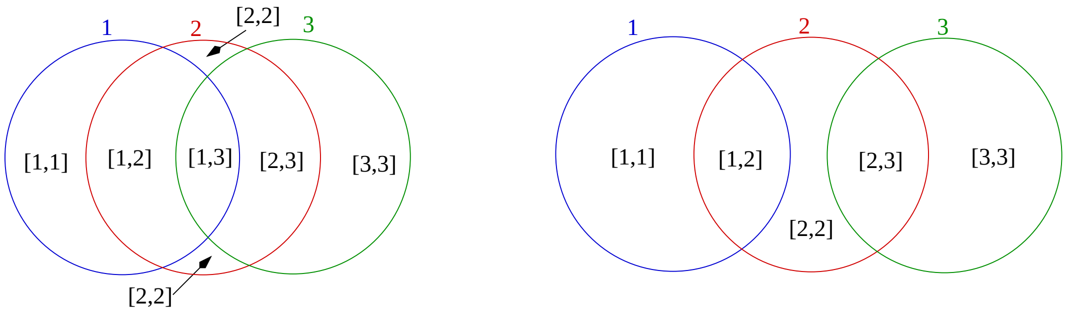

Figure 7: On the left is an arrangement of 3 disks that satisfies the full interval property, and on the right is an arrangement that satisfies the interval property but not the full interval property. In particular, in both cases the disks are unit disks and their centers reside on a line. However, in the arrangement on the right, the first and the third disk do not intersect and hence there is no cell v such that $N(v) = [1, 3]$.

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 an 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 $N(v) = [i, j]$.*

For an illustration of the interval and full interval properties, see Figure 7. 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 on the top of Figure 8. Observe that every two disks in the arrangement intersect.

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 all intervals $[i, j] \subseteq [1, n]$. It follows that the arrangement $\mathcal{A}(\mathcal{S})$ is a chain, as claimed. For an illustration see (the bottom of) Figure 8.

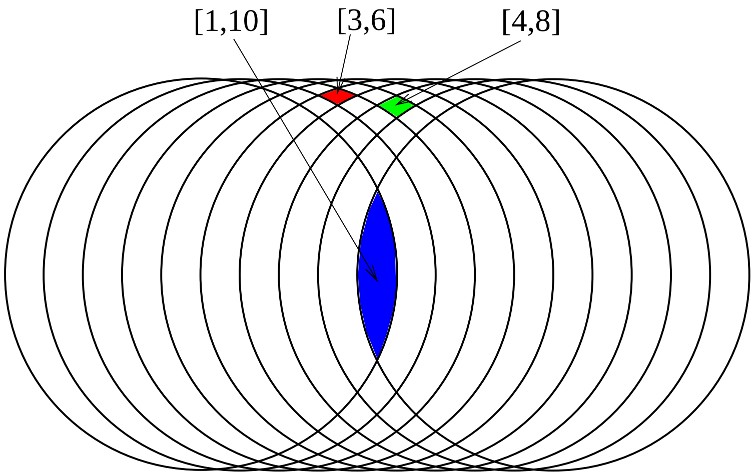

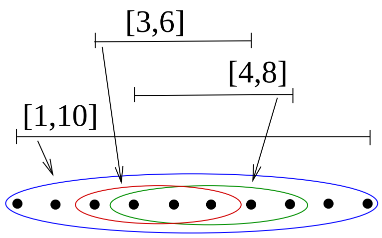

Figure 8: On the top is a chain of disks, where the disks are numbered $1, \ldots, 10$, from left to right. The three cells that are marked correspond to the three respective intervals. On the bottom is an illustration of the dual range space. In the dual space there is a point for every disk in the primal space, and a subset for every cell in the primal space. Since the cells in the primal space correspond to intervals, the subsets in the dual space correspond to intervals $[i, j] = \{i, i+1, \ldots, j\}$ as well.

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 subintervals 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 subinterval 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 subintervals of $[1, i-1]$, (ii) I' - the set of all subintervals of $[1,n]$ that contain i, and (iii) $I_{(i+1),n}$ - the set of all subintervals of $[i+1,n]$. (Observe that if $i = 1$ (respectively, $i = n$) then $I_{1,(i-1)}$ (respectively, $I_{(i+1),n}$) 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 $\chi(j)$ must 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. □

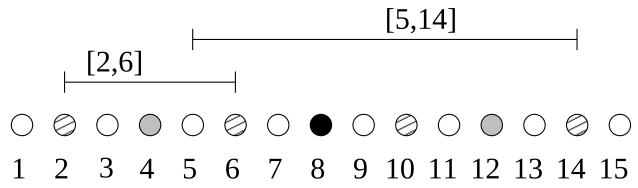

Figure 9: An illustration for Lemma 15. The 15 points in the figure are the points of the dual range space. Point number 8 is colored with the "highest color" (corresponding to the densest filling in the illustration). Points 4 and 12 are colored with the next highest color, and points 2, 6, 10 and 14 with the next one. The remaining points (all odd-numbered points) are colored with the lowest color. If we now consider for example the intervals (ranges in the dual space) $[5,14]$ and $[2,6]$, then the first can be served by point number 8, and the latter by point number 4.

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

It suffices to prove the above lemma for chains (in terms of the dual range space, this simply means that we add constraints). In fact, in Section 3.1 we already presented a proof of the above lemma in the special case of unit disks whose centers reside on a line. In Section 6.1 we showed that a chain can be obtained from an arrangement of unit disks whose centers are collinear. Hence, the lemma follows. We provide an alternative proof of the above lemma that follows the spirit of the proof of the lower bound stated in Lemma 14.

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 $([1,n], I_{1,n})$ (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 $n = 1$ 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 subintervals of $[1, n]$ that contain i. A subinterval 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. □

7 An Approximation Algorithm for Rectangles

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

We begin with a high-level description of the algorithm (for the special case of axis parallel unit squares). The algorithm starts by partitioning the plane into square tiles of side-lengths 1/2. Given a set of $\mathcal{S}$ of unit squares, we say that a square $s \in \mathcal{S}$ belongs to a tile if its center resides inside the tile. Hence the tiling induces a partition of $\mathcal{S}$. We first observe that squares that belong to sufficiently distant tiles do not intersect. Therefore, as shown for the case of disks, we may assign each tile a palette of colors so that the total number of palettes used is constant, and any two different tiles whose squares may intersect, are assigned different palettes. At this point we could simply apply Theorem 4 to separately color the squares that belong to each tile. This would give us a CF-coloring that uses $O(\log \phi(\mathcal{S}))$ colors, where $\phi(\mathcal{S})$ is the maximum number of centers of squares in $\mathcal{S}$ that are contained in a square tile of side-lengths 1/2. However, the resulting coloring may be far from optimal (recall that we are interested in a constant-ratio approximation algorithm). 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 10.

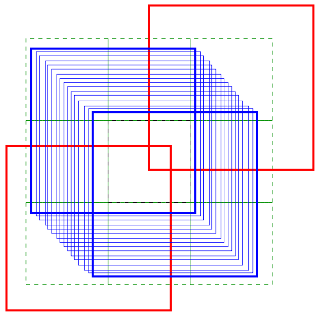

Figure 10: An example illustrating how by taking into account intersections between squares that belong to different tiles we may significantly reduce the number of colors required in a CF-coloring. Here there is a large number of squares that belong to the middle tile and constitute a chain. If we color 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.

Instead of coloring all squares as suggested above, our algorithm selects only a subset of squares, that are "essential" for serving the area covered by the union of the squares. Once this stage is over, we can "return" to each tile from which squares were requested, and color the requested squares by applying Theorem 4. The notion of being essential, is formalized later on in this section. It enables us to show that the total number of colors used is indeed necessary, up to a constant. Clearly, every tile that contains the center of at least one square, can be completely served by any one of the squares that belong to it. The main issue hence is serving tiles that lack centers of squares. We refer to such tiles as "orphan" tiles. In what follows we describe how an orphan tile selects the squares that are used to serve it.

Consider an orphan tile T (for which the set of squares that intersect it is nonempty). The squares that intersect it (and may hence serve parts of it), can be partitioned into 2 types: those that intersect it with an edge and those that intersect it with a corner. Each type can be further partitioned into 4 subtypes according to the edge-type (respectively, corner-type) with which they intersect T. We first observe that within each subtype of squares that intersect T with an edge, we can select a single square that can serve all the area within T that is covered by squares of this subtype. After selecting one square from each subtype, we are essentially left with the problem of serving a rectangular region, denoted by T', that is contained in T. For an illustration, see Figure 11.

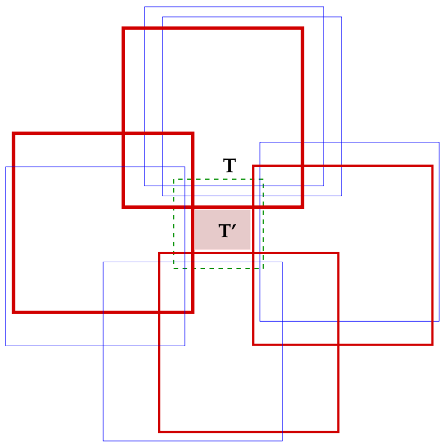

Figure 11: An illustration of the selection of squares that intersect an orphan tile T with an edge. The tile T is the dashed square, the selected squares are marked in bold, and the remaining unserved region T' is shaded.

Suppose we now consider separately each subset of squares that intersect T' with a common corner-type (e.g., top-right). For each corner type, it suffices to focus on the subset of squares that participate in the envelope of the squares that intersect T'. However, the envelopes of squares corresponding to different corner types may intersect in T'. Such intersections may help in reducing the number of squares needed to cover the intersection of T' with the union of all squares. For an illustration see Figure 12.

In order to address the issue of "intersections" between envelopes of subsets of squares corresponding to different corner types, we consider these subsets of squares in pairs. Specifically, we first deal with "adjacent" pairs whose corresponding corners have a common edge (e.g., top-right and top-left), and then with "opposite" pairs (top-right with bottom-left and bottom-right with top-left). For the first class of adjacent pairs, we show that by selecting at most two squares per

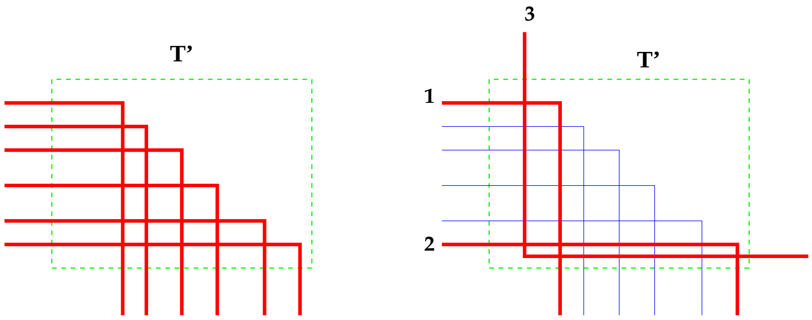

Figure 12: On the left is a subset of squares that intersect the rectangular region T' with their top-right corner. In order to cover the intersection of T' with their union it is necessary to select all squares. On the right is the same set of squares together with one more square, which intersects T' with its bottom-left corner. Now only the 3 bold squares (labeled 1, 2, and 3) are necessary.

pair, we can serve the region of the intersection of each pair. For the second class of "opposite" pairs, we describe a procedure that selects a subset of squares that is at most a constant factor larger than necessary. Further details for these, more involved, steps are given in Subsections 7.2 and 7.3.

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, respectively) in $\mathcal{R}$ ($\mathcal{S}$, respectively) 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 $\text{op} : \Gamma \to \Gamma$ denote the permutation that swaps opposite corners (i.e. $\text{op} = (\llcorner, \urcorner)(\ulcorner, \lrcorner)$). 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 (respectively, v-neighbors) if they share an edge (respectively, a corner). The v-neighbor of T that shares its γ-corner with the $\text{op}(\gamma)$ 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 determined by rectangles having the same corner 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 has at most one corner in T. Define the Q-*envelope* of $\mathcal{R}(Q,\gamma)$ to be the boundary of $\mathcal{R}(Q,\gamma)$ that is in Q (see Figure 13). The vertices of a 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 next claim shows that the corner γ determines whether the Q-envelope is nonincreasing or nondecreasing.

Claim 16 *Let Q be a rectangular region with side-lengths at most $1/2$. If $\gamma \in \{\llcorner, \urcorner\}$ then the Q-envelope of $\mathcal{R}(Q,\gamma)$ is nonincreasing, and if $\gamma \in \{\ulcorner, \lrcorner\}$ then the Q-envelope of $\mathcal{R}(Q,\gamma)$ is nondecreasing.*

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 $\widetilde{\mathcal{R}}(Q,\gamma)$. $\square$

The next definition will be useful in all that follows.

Definition 9 *Let Q be a region and let $\mathcal{S}$ be an indexed set of regions. Let $\mathcal{S}_Q$ denote the set of regions $\{Q \cap S\}_{S \in \mathcal{S}}$. Assume that each region $Q \cap S \in \mathcal{S}_Q$ inherits the index of S. We say that $\mathcal{S}$* is a chain with respect to Q *if the arrangement $\mathcal{A}(\mathcal{S}_Q)$ is a chain.*

Claim 17 *Let Q be a rectangular region with side-lengths at most $1/2$. 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*.

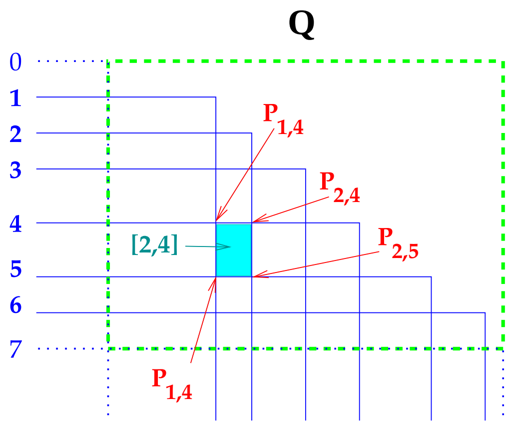

Figure 13: An illustration of a corner-chain consisting of 6 rectangles (the two dotted rectangles, numbered 0 and 7 are only for the sake of the analysis in the proof of Claim 17). The cell corresponding to the interval $[2,4]$ is filled, and its 4 defining corners are marked.

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 $\urcorner$-corner. Let R_0 and R_{m+1} be two "fictitious" rectangles, where the right side of R_0 coincides with the left side of Q, and the top side of R_{m+1} coincides with the bottom side of Q. Let $P_{i,j}$, for $0 \leq i \leq j \leq m+1$, denote the intersection of the right side of R_i and the top side of R_j. Note that the for every $1 \leq i \leq m$, $P_{i,i} = \urcorner(R_i)$, $P_{i,m+1}$ is the intersection of R_i with the bottom side of Q, and $P_{0,i}$ is the intersection of R_i with the left side of Q. By Claim 16, it follows that $P_{i,j}$ is well defined and that $P_{i,j} \in Q$, for every $1 \leq i \leq j \leq m+1$. The arrangement of $\widetilde{\mathcal{R}}(Q, \gamma)$ in Q is a set of rectangular shaped cells, the corners of which are the set of points $\{P_{i,j}\}$. Specifically, for every $1 \leq i \leq j \leq m$, the cell v for which $N(v) = [i, j]$ is the rectangle whose corners are $P_{i-1,j+1}, P_{i-1,j}, P_{i,j}$ and $P_{i,j+1}$. $\square$

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. Specifically, in this section we provide our main lemmas concerning interactions between corner-chains of opposite corners and corner-chains of adjacent corners.

7.2.1 Corner-chains of adjacent corners

Consider a rectangle Q with side lengths at most $1/2$. Let $\widetilde{\mathcal{S}}_{\urcorner} = \widetilde{\mathcal{S}}(Q, \urcorner)$ and $\widetilde{\mathcal{S}}_{\ulcorner} = \widetilde{\mathcal{S}}(Q, \ulcorner)$ denote corner-chains corresponding to adjacent corners $\urcorner$ and $\ulcorner$ (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. That is (as formalized in the next lemma), after picking one square from each chain, the squares having smaller indexes than those picked, form a chain with respect to the remaining region.

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

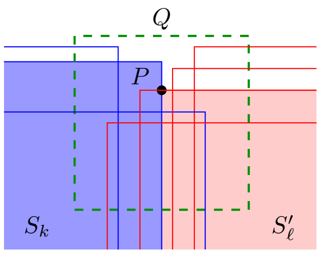

Figure 14: An illustration for Lemma 18. The region Q is depicted by a dashed rectangle. Only the corners of squares in the corner chains are depicted. The two filled squares are the selected squares S_k and S'_ℓ.

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 $\left(\bigcup_{t=k}^{m} (S_t \cap Q)\right) \cup \left(\bigcup_{t=\ell}^{m'} (S'_t \cap Q)\right)$. 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 nonincreasing (nondecreasing). 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 $B_x > P_x$. Consider first any square $S_{k'} \in \widetilde{\mathcal{S}}_{\neg}$ where $k' < k$. By our assumption that P is along the vertical right side of S_k, we have that $P_x = x_{\neg}(S_k)$. By the ordering of the squares in $\widetilde{\mathcal{S}}_{\neg}$, we have that $x_{\neg}(S_{k'}) < x_{\neg}(S_k)$, and hence for every $A \in S_{k'}$, $A_x \leq x_{\neg}(S_{k'}) < x_{\neg}(S_k)$. It directly follows that $A_x < P_x$. Next consider a square $S'_{\ell'} \in \widetilde{\mathcal{S}}_{\ulcorner}$ where $\ell' < \ell$. By our assumption that P is along the horizontal upper side of S'_ℓ, and by the ordering of the squares in $\widetilde{\mathcal{S}}_{\ulcorner}$, necessarily $x_{\ulcorner}(S'_{\ell'}) > P_x$. But, for every point $B \in S'_{\ell'}$, $B_x \geq x_{\ulcorner}(S'_{\ell'})$ and so $B_x > P_x$.

To prove Part (2) it suffices to show that (i) $(S_{k'} \setminus S_k) \cap S'_\ell = \emptyset$ if $k' < k$, and (ii) $(S'_{\ell'} \setminus S'_\ell) \cap S_k = \emptyset$ if $\ell' < \ell$. This is sufficient since $\widetilde{\mathcal{S}}_{\neg}$ (respectively, $\widetilde{\mathcal{S}}_{\ulcorner}$) is a chain with respect to Q. Hence, every cell corresponding to an interval $[i, j] \subseteq [1, k-1]$ (respectively, $[i, j] \subseteq [1, \ell-1]$) of $\widetilde{\mathcal{S}}_{\neg}$ (respectively, $\widetilde{\mathcal{S}}_{\ulcorner}$) in Q is disjoint from $S_k \cup S'_\ell$, and the full interval property is preserved. For example, consider the two squares in Figure 14 that belong to the $\ulcorner$-chain and are above S'_ℓ. If we denote them by S_1 and S_2, then we see that the regions $S_1 \setminus S'_\ell$ and $S_2 \setminus S'_\ell$ are both disjoint from S_k and that $\{S_1, S_2\}$ form a chain with respect to $Q \setminus (S_k \cap S'_\ell)$.

In order to verify (i), consider a square $S_{k'}$ for $k' < k$. To show that $(S_{k'} \setminus S_k) \cap S'_\ell = \emptyset$, consider a point $A \in (S_{k'} \setminus S_k)$. The ordering of $\widetilde{\mathcal{S}}_{\neg}$ implies that $A_y > y_{\neg}(S_k)$, and by the definition of P, $y_{\neg}(S_k) \geq P_y$. Since $P_y = y_{\ulcorner}(S'_\ell)$, we get that A is above S'_ℓ and in particular $A \notin S'_\ell$ as claimed in (i). 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, $A_y \leq y_{\neg}(S_{k'}) \leq y_{\neg}(S_k)$. Since $P_x = x_{\neg}(S_k)$, we get 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 $P_x \geq x_{\ulcorner}(S'_\ell)$, we get that $A \in S'_\ell$. Therefore, the suffix of $\widetilde{\mathcal{S}}_{\neg}$ is covered by $S_k \cup S'_\ell$. The proof for the suffix of $\widetilde{\mathcal{S}}_{\ulcorner}$ is analogous, and part (3) follows. $\square$

7.2.2 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 $\urcorner$ and $\llcorner$ (the case of the $\lrcorner$-corner and $\ulcorner$-corner is reduced to this case by "flipping or rotating 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 minimum 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 10 *A subset $\widetilde{\mathcal{S}}^{\min}_{\urcorner} \subseteq \widetilde{\mathcal{S}}_{\urcorner}$ is a* minimal cover *of $Q_{\urcorner} \setminus Q_{\llcorner}$ if (i) $\widetilde{\mathcal{S}}^{\min}_{\urcorner}$ covers $Q_{\urcorner} \setminus Q_{\llcorner}$, and (ii) no proper subset of $\widetilde{\mathcal{S}}^{\min}_{\urcorner}$ covers $Q_{\urcorner} \setminus Q_{\llcorner}$.*

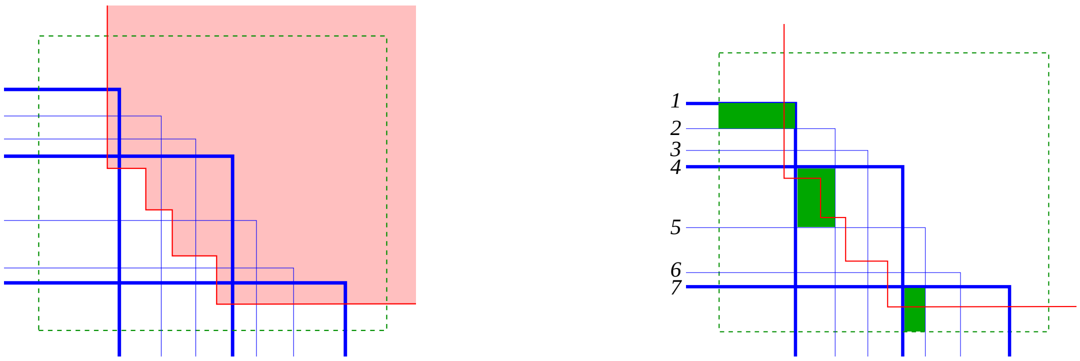

(A) (B)

Figure 15: A minimal cover of the union of opposite corner-chains. (A) The rectangle Q is depicted by a dashed rectangle. The union of the "upper" corner chain $Q_{\llcorner}$ is shaded (so that $Q_{\urcorner} \setminus Q_{\llcorner}$ is the unshaded region within Q). A minimal cover $\widetilde{\mathcal{S}}^{\min}_{\urcorner} \subseteq \widetilde{\mathcal{S}}_{\urcorner}$ is depicted by thick $\urcorner$-corners. (B) An explanation of how $\widetilde{\mathcal{S}}^{\min}_{\urcorner}$ is greedily computed. Here only the boundary of $Q_{\llcorner}$ is depicted. Cells $[1,1], [2,4]$, and $[5,7]$ are shaded. Since cell $[1,1]$ is not completely covered by $Q_{\llcorner}$, we add square 1 to the cover. Cells $[2,2]$ and $[2,3]$ are covered by $Q_{\llcorner}$, but cell $[2,4]$ is not. Therefore, square 4 is added to the cover. Now cells $[5,5]$ and $[5,6]$ are covered by $Q_{\llcorner}$, but cell $[5,7]$ is not. Therefore, square 7 is added to the minimal cover.

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

Lemma 19 *If $\widetilde{\mathcal{S}}^{\min}_{\urcorner} \subseteq \widetilde{\mathcal{S}}_{\urcorner}$ is a minimal cover of $Q_{\urcorner} \setminus Q_{\llcorner}$ and the squares in $\widetilde{\mathcal{S}}^{\min}_{\urcorner}$ are indexed according to the x-coordinate of their $\urcorner$-corners, then $\widetilde{\mathcal{S}}^{\min}_{\urcorner}$ is a chain with respect to $Q_{\urcorner} \setminus Q_{\llcorner}$.*

Proof: Let $\widetilde{\mathcal{S}}^{\min}_{\urcorner} = \{S'_1, \ldots, S'_k\}$. Since $\widetilde{\mathcal{S}}_{\urcorner}$ is a chain with respect to Q, it follows that $\widetilde{\mathcal{S}}^{\min}_{\urcorner}$ is also a chain with respect to $Q_{\urcorner}$. For simplicity, add "dummy" squares S'_0 and S'_{k+1} to $\widetilde{\mathcal{S}}^{\min}_{\urcorner}$, where $\urcorner(S'_0) = \ulcorner(Q)$ and $\urcorner(S'_{k+1}) = \lrcorner(Q)$. Note that these dummy squares do not assist in covering $Q_{\urcorner} \setminus Q_{\llcorner}$. For the sake of contradiction, assume that $\widetilde{\mathcal{S}}^{\min}_{\urcorner}$ is not a chain with respect to $Q_{\urcorner} \setminus Q_{\llcorner}$. Since $\widetilde{\mathcal{S}}^{\min}_{\urcorner}$ is not a chain, it is not empty. Consider an interval $[i,j]$, for $0 < i \le j < k+1$, such that the corresponding cell in $\mathcal{A}(\widetilde{\mathcal{S}}^{\min}_{\urcorner})$ is contained in $Q_{\llcorner}$. (See Figure 16 for an illustration of 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 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 $\urcorner(S'_{i-1})$, $\urcorner(S'_i)$, $\urcorner(S'_j)$, and $\urcorner(S'_{j+1})$ are in S_B. It is easy to see that we may omit both S'_i and S'_j from $\widetilde{\mathcal{S}}^{\min}_{\urcorner}$ while still covering $Q_{\urcorner} \setminus Q_{\llcorner}$, contradicting the assumption that $\widetilde{\mathcal{S}}^{\min}_{\urcorner}$ is a minimal cover. $\square$

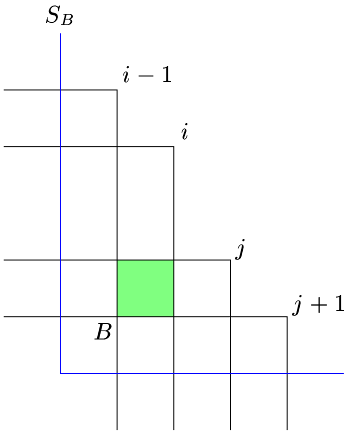

Figure 16: The construction in the proof of Lemma 19.

An algorithm for finding a minimal cover. We now describe a greedy algorithm for finding a subset $\widetilde{\mathcal{S}}_{\neg}^{\min} \subseteq \widetilde{\mathcal{S}}_{\neg}$ 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 $1 \leq a \leq b \leq m$, 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 $\widetilde{\mathcal{S}}_{\neg}^{\min} = \emptyset$). Consider all cells $\widetilde{\mathcal{S}}_{\neg}[k+1,\ell]$ where $(k+1) \leq \ell \leq m$ 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 $\widetilde{\mathcal{S}}_{\neg}^{\min}$. For an example, see Figure 15 (B).

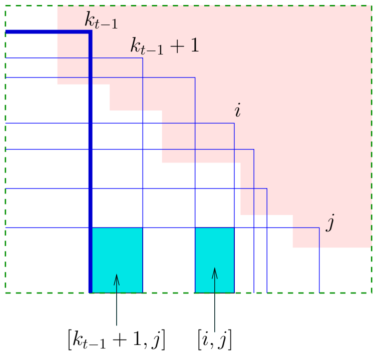

Figure 17: An illustration for the case $i > k_{t-1}$ in the proof of Claim 20.

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

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

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

$$(Q_\neg \setminus Q_\llcorner) \cap \bigcup_{i=1}^{k_t} S_i \subseteq \bigcup_{j=1}^{t} S_{k_j}.$$

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 in $Q_\neg \setminus Q_\llcorner$ then the cell $\widetilde{\mathcal{S}}_\neg[k_{t-1}+1,j]$ is also in $Q_\neg \setminus Q_\llcorner$ but in such a case, the greedy algorithm would have chosen S_{k_t} such that $k_{t-1}+1 \le k_t \le j$, a contradiction. For an illustration of this case, see Figure 17.

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

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

Let $m_\neg = |\widetilde{\mathcal{S}}_\neg^{\min}|$ and let $m_\llcorner = |\widetilde{\mathcal{S}}_\llcorner^{\min}|$. 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 $Q_\neg \cup Q_\llcorner$.*

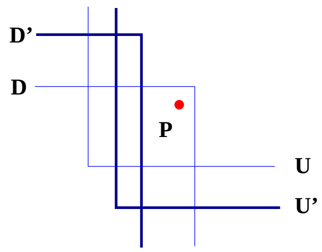

Figure 18: An illustration for Lemma 21. The point $P \in Q_\neg \cap Q_\llcorner$ is in $D \cap U$ where neither $D \in \widetilde{\mathcal{S}}_\neg^{\min}$, nor $U \in \widetilde{\mathcal{S}}_\llcorner^{\min}$. Here $D' \in \widetilde{\mathcal{S}}_\neg^{\min}$ (the square that covers the cell that contains the point slightly to the left of $\llcorner(U)$) is such that $y_\neg(D') > P_y$. Finally, $U' = B_\llcorner(D')$, that is, it is the last square from $\widetilde{\mathcal{S}}_\llcorner$ that intersects D'.

Proof: Since $\widetilde{\mathcal{S}}_\neg^{\min}$ (respectively, $\widetilde{\mathcal{S}}_\llcorner^{\min}$) covers $Q_\neg \setminus Q_\llcorner$ (respectively, $Q_\llcorner \setminus Q_\neg$) and $|\widetilde{\mathcal{S}}_\neg^{\min} \cup \widetilde{\mathcal{S}}_\llcorner^{\min}| \le 2m$, the remaining problem is to cover $Q_\neg \cap Q_\llcorner$ using $O(m)$ squares. For every square $S \in \widetilde{\mathcal{S}}_\neg^{\min}$ consider the set $\widetilde{\mathcal{S}}_\llcorner(S)$ of squares in $\widetilde{\mathcal{S}}_\llcorner$ that intersect S. Define $A_\llcorner(S)$ ($B_\llcorner(S)$, respectively) to be the first (last, respectively) square in $\widetilde{\mathcal{S}}_\llcorner(S)$ when sorted according to the y-coordinate of their $\llcorner$-corner. We claim that $\bigcup_{S\in\widetilde{\mathcal{S}}_\neg^{\min}} (A_\llcorner(S) \cup B_\llcorner(S))$ covers $(Q_\neg \cap Q_\llcorner) \setminus (\widetilde{\mathcal{S}}_\neg^{\min} \cup \widetilde{\mathcal{S}}_\llcorner^{\min})$. Hence, we need to add at most two squares from $\widetilde{\mathcal{S}}_\llcorner$ per square in $\widetilde{\mathcal{S}}_\neg^{\min}$ to cover $(Q_\neg \cap Q_\llcorner) \setminus (\widetilde{\mathcal{S}}_\neg^{\min} \cup \widetilde{\mathcal{S}}_\llcorner^{\min})$.

Consider a point $P \in Q_\neg \cap Q_\llcorner$. Let $D \in \widetilde{\mathcal{S}}_\neg$ ($U \in \widetilde{\mathcal{S}}_\llcorner$, respectively) denote a square that contains P. If $D \in \widetilde{\mathcal{S}}_\neg^{\min}$ or $U \in \widetilde{\mathcal{S}}_\llcorner^{\min}$ 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 $Q_{\urcorner} \setminus Q_{\llcorner}$, and therefore, there exists a square $D' \in \widetilde{\mathcal{S}}_{\urcorner}^{\min}$ that covers this cell. If $P \in D'$, we are done. Otherwise, we consider two cases: $y_{\urcorner}(D') \geq P_y$ and $y_{\urcorner}(D') < P_y$. In the first case 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) $x_{\llcorner}(U') \leq x_{\urcorner}(D') < P_x$ (the first inequality holds because U' and D' intersect, and the second inequality holds because $y_{\urcorner}(D') \geq P_y$ while $P \notin D'$); (ii) $y_{\llcorner}(U') \leq y_{\llcorner}(U) \leq P_y$ (the first inequality holds because $U' = B_{\llcorner}(D')$, and the second by the premise of this case). Thus $P \in U'$. The second case in which $y_{\urcorner}(D') < P_y$ is treated analogously, where here we let $U' = A_{\llcorner}(D')$. The claim follows. □

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

As we discuss in more detail shortly, 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 Claim 20 and Lemma 21 for the covering of the remaining region in the tile.

Remark 2 After the separation of adjacent corner-chains in a tile, it is not possible for both pairs of opposite corner-chain to intersect. Namely, at most one pair of opposite corner-chains may intersect. We do not rely on this property, the proof of which is easy.

7.3 Coloring Arrangements of Squares

In this section we prove Theorem 3 for unit squares. 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. As noted at the start of this section, 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 the intersection of T with their union.

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^e of T. For brevity, assume that T^e is to the left of T and that $\mathcal{S}(T^e) \neq \emptyset$. Every square $S \in \mathcal{S}(T^e)$ covers a vertical strip of T. If we select the rightmost square S in $\mathcal{S}(T^e)$, then we get that for every $S' \in \mathcal{S}(T^e)$, $S' \cap T \subseteq S \cap T$. 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 substep covers all the points in $T \cap \bigcup_{T^e \in e\text{-neighbors}(T)} \mathcal{S}(T^e)$.

After this step, the region within the tile T that still needs to be served is a rectangle. Let us denote this rectangle by T'. Note that the union of squares in $\mathcal{S}$ may either fully cover or partly cover the rectangle T'. In any case, only squares that belong to v-neighbors of T intersect T'. An illustration of this step was provided in Figure 11.

Selecting squares from v-neighbors: adjacent corners. Consider the rectangle $T' \subseteq T$ and a corner γ. The squares of $\mathcal{S}(T',\gamma)$ that participate in the T'-envelope are denoted by $\widetilde{\mathcal{S}}(T',\gamma)$. By Claim 17, $\widetilde{\mathcal{S}}(T',\gamma)$ is a chain with respect to T' 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 T' 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. Let T'' denote the subregion consisting of T' minus the union of the (at most 8) selected squares. By our notational convention, $\widetilde{\mathcal{S}}(T'',\gamma)$ denotes the subset of squares in $\mathcal{S}(T'',\gamma)$ that participate in the T'' envelope. Note that by Lemma 18, $\widetilde{\mathcal{S}}(T'',\gamma)$ is a chain with respect to T''.

Selecting squares from v-neighbors: opposite corners. In the third step we apply Claim 20 and Lemma 21 to each pair of subsets $\widetilde{\mathcal{S}}(T'',\gamma)$ and $\widetilde{\mathcal{S}}(T'',\mathrm{op}(\gamma))$. This application determines the subsets of $\widetilde{\mathcal{S}}(T'',\gamma)$ (and $\widetilde{\mathcal{S}}(T'',\mathrm{op}(\gamma))$) that suffice to serve the intersection of T'' with the union of each pair of chains. Note that, due to the separation of adjacent corner chains (see Lemma 18), at most one pair of opposite corner chains may intersect.

A subtle issue to be addressed is whether the remaining region $T'' \subseteq T' \subseteq T$ in the beginning of this step satisfies the premises of Remark 1 for each pair of opposite corner chains. Consider, for example, the subset $\widetilde{\mathcal{S}}(T'',\urcorner)$. This subset is a consecutive block of squares from $\widetilde{\mathcal{S}}(T',\urcorner)$. Let $S^f_\urcorner$ and $S^\ell_\urcorner$ denote the squares in $\widetilde{\mathcal{S}}(T',\urcorner) \setminus \widetilde{\mathcal{S}}(T'',\urcorner)$ that "hug" this block (i.e., $S^f_\urcorner$ and $S^\ell_\urcorner$ were selected in the adjacent corner-chain stage). The designated point $C_\llcorner \in T''$ is the intersection of the right side of $S^f_\urcorner$ and the top side of $S^\ell_\urcorner$. One can define in this fashion all four designated points C_γ, for $\gamma \in \Gamma$. We can now apply Lemma 21 to each pair of subsets $\widetilde{\mathcal{S}}(T'',\gamma)$ and $\widetilde{\mathcal{S}}(T'',\mathrm{op}(\gamma))$, where the corresponding designated points that satisfy the premise of Remark 1, are $C_{\mathrm{op}(\gamma))}$ and C_γ. For an illustration see Figure 19.

7.3.2 Coloring the Essential Squares

In the previous steps, each orphan tile T_o selected a subset of squares that are used to serve the points in $T_o \cap \bigcup \mathcal{S}$. Given a non-orphan tile T, let $sel(T) \subseteq \mathcal{S}(T)$ denote the subset of squares whose centers reside in T that are selected by some tile in order to participate in its cover. 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, and let $sel(T)$ contain only this square. 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.

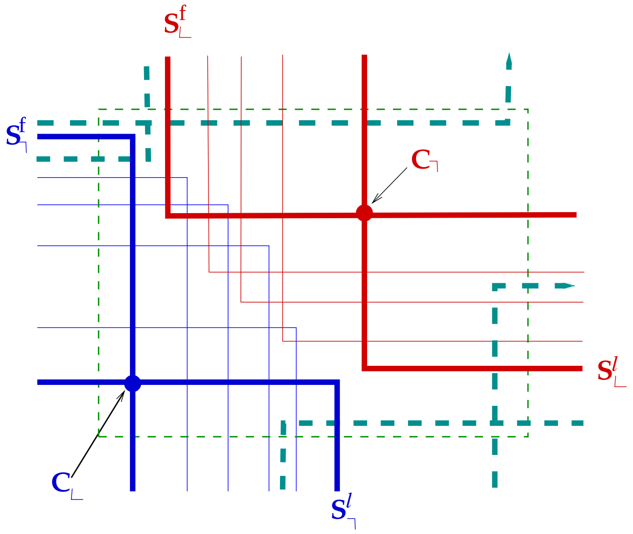

Figure 19: An illustration for the choice of points C_γ and $C_{\mathrm{op}(\gamma)}$ that satisfy the premise of Remark 1. The selected squares, which determine the two points, are labeled and marked in bold. The dashed bold corners of squares correspond to the four other selected squares that belong to the adjacent chains $\widetilde{\mathcal{S}}(T', \ulcorner)$ and $\widetilde{\mathcal{S}}(T', \lrcorner)$. The thin dashed rectangle is T', and T'' is the region obtained by removing the four bold and four dashed-bold squares from 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 were requested by v-neighbors of T in the last selection step. That is, in the step the deals with opposite corner chains.

For each corner-type γ, let $sel_\gamma(T)$ denote the subset of squares in $sel(T)$ that were selected by $T_{\mathrm{op}(\gamma)}$ in the opposite-corners selection step. The γ-corners of squares in $sel_\gamma(T)$ are contained in the tile $T_{\mathrm{op}(\gamma)}$. Let $m_\gamma(T) = |sel_\gamma(T)|$. Since $|sel(T)| = O(\max_{\gamma\in\Gamma}\{m_\gamma(T)\})$, and since there are 9 palettes, the next corollary follows:

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

7.3.3 A Lower Bound for Optimal CF-Coloring

In this section we lower bound the number of colors required by an optimal CF-coloring. Recall that for a tile T and corner $\gamma \in \Gamma$, 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.

Let T be any (orphan) tile, and let γ be a corner. The following lemma states a lower bound on $\chi_{\mathrm{opt}}(\mathcal{S})$ in terms of the size of a subset $\mathcal{S}' \subseteq \widetilde{\mathcal{S}}(T, \gamma)$ that is a chain with respect to the region $Q = T \setminus \{S : S \notin \mathcal{S}(T, \gamma)\}$. That is, the region Q is what remains of T after we remove all squares that intersect T with the exception of squares in $\mathcal{S}(T, \gamma)$.

Lemma 22 *Let T denote a tile, γ a corner type, and let $Q = T\setminus\{S : S \notin \mathcal{S}(T, \gamma)\}$. Let $\mathcal{S}' \subseteq \widetilde{\mathcal{S}}(T, \gamma)$ be a chain with respect to 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. In fact, the same lower bound holds also for CF-multi-coloring, implying Theorem 10 (see Section 9). The only difference is that here we need to take into account that squares from $\mathcal{S}(T,\gamma)\setminus\mathcal{S}'$ may cover points in Q, and hence can potentially serve cells in the chain $\mathcal{S}'$. For an illustration of the proof of the lemma, see Figure 21.

Proof: We consider the case $\gamma = \urcorner$. All other cases are proved analogously. Let $S_1,\ldots,S_m$ be an ordering of the squares in $\mathcal{S}'$ so that $x_\urcorner(S_1) < \ldots < x_\urcorner(S_m)$. For every $1 \le i \le j \le m$, let $v_{i,j}$ denote the cell in the arrangement $\mathcal{A}(\mathcal{S}')$ for which $N(v_{i,j}) = \{S_i,\ldots,S_j\}$ (recall that for a cell v, $N(v)$ denotes the subset of regions that contain v). Let $v^Q_{i,j} = v_{i,j}\cap Q$, where we know that $v^Q_{i,j}$ is nonempty for every $1 \le i \le j \le m$ because $\mathcal{S}'$ is a chain with respect to Q. Since we are dealing with squares (or more generally, rectangles), we know that each $v_{i,j}$ is a rectangle. However, it may be the case that $v^Q_{i,j}$ is not a rectangle. The exact structure of $v^Q_{i,j}$ is actually immaterial to the proof. What will be needed is the following subclaim (where we assume that $m > 3$ or else Lemma 22 holds trivially):

Observation: For every $1 \le i,j \le m$ such that $i \le j-2$, the top-right corner of $v_{i,j}$ is contained also in $v^Q_{i,j}$.

Proof of observation: Assume that the observation does not hold. This means that there exists some square $S' \notin \mathcal{S}(T,\urcorner)$ that contains the top-right corner of $v_{i,j}$. But in such a case either $S' \supseteq v_{i+1,j}$ or $S' \supseteq v_{i+1,j-1}$ or $S' \supseteq v_{i,j-1}$, contradicting the premise of the lemma that $\mathcal{S}'$ is a chain with respect to Q. For an illustration, see Figure 20.

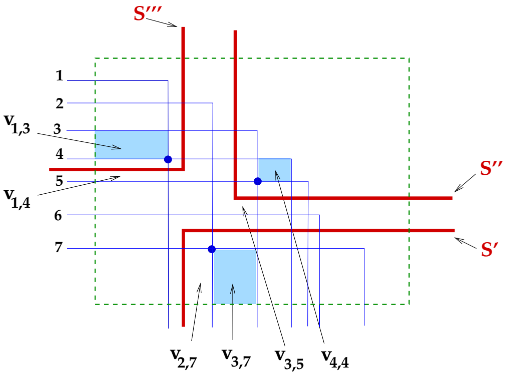

Figure 20: An illustration for the proof of the observation: (i) $i=2$, $j=7$ and S' covers $v_{3,7}$, (ii) $i=3$, $j=5$, and S'' covers $v_{4,4}$, and (iii) $i=1$, $j=4$, and S''' covers $v_{1,3}$.

Now consider an optimal CF-coloring $\chi_{\rm opt}$ of $\mathcal{S}$. Let P be the top-right corner of $v_{1,m}$. By the above claim (assuming $m > 3$), $P \in v^Q_{1,m}$. 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 $\chi_{\rm opt}$. 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 subchains. The second observation is that every square $S' \in \mathcal{S}(T,\urcorner)$ that serves the top-right corner P' of a cell in one of these chains (that corresponds to an interval of size at least 3), must also contain P. This is true because every top-right corner of such a cell 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 that contains P must be different from $\chi_{\text{opt}}(S)$.

If $\urcorner(S)$, the top-right corner of S, is not contained in any $S_k \in \mathcal{S}' = \{S_1, \ldots, S_m\}$, then define k as follows: $k = \max\{i : x_{\urcorner}(S_i) < x_{\urcorner}(S)\}$. Since $\mathcal{S}'$ is a chain with respect to Q, 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 $S' \in \mathcal{S}(T, \urcorner)$ that can serve the top-right corner of a cell in one of these chains (that corresponds to an interval of size at least 3) $\chi_{\text{opt}}(S') \neq \chi_{\text{opt}}(S)$. In either case we get the recurrence relation

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

where for any subset $\mathcal{S}''$ of less than 3 squares, $|\chi_{\text{opt}}(\mathcal{S}'')| \geq 1$. Hence $|\chi_{\text{opt}}(\mathcal{S})| = \Omega(\log |\mathcal{S}'|)$, and the lemma follows. $\square$

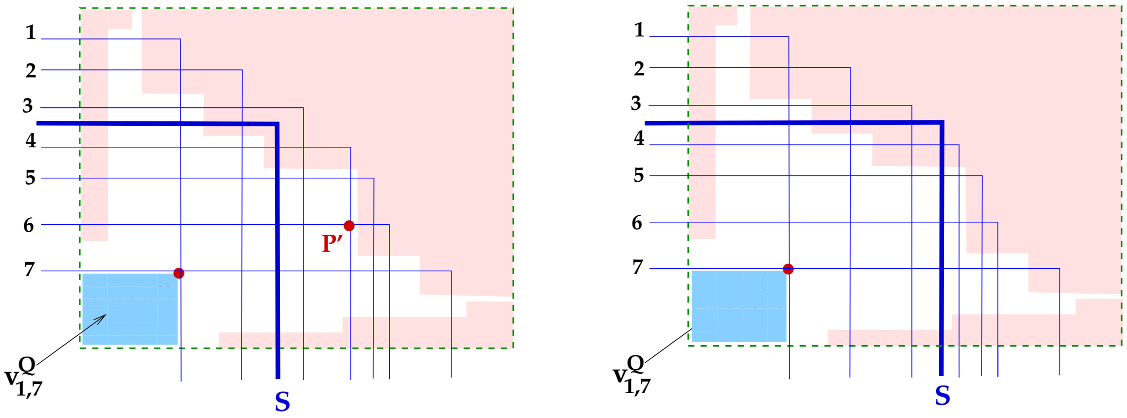

Figure 21: Illustration for the proof of Lemma 22. The squares of $\mathcal{S}'$ are labeled from 1 to 7. The region Q is the non-shaded region within the dashed rectangle. The point P is the top-right corner of $v_{1,7}$ (which here equals $v^Q_{1,7}$). On the left is an illustration of the case in which P is served (in an optimal CF coloring) by a square S such that $S \cap Q \subseteq S_3 \cap Q$. The point P' in this case is the top-right corner in a cell of the chain determined by $S_4, \ldots, S_7$. On the right is an illustration of the case in which P is served by a square S whose top-right corner does not reside in any square of $\mathcal{S}'$.

For any non-orphan tile T and corner γ, let $sel_\gamma(T)$ and $m_\gamma(T)$ $(= |sel_\gamma(T)|)$ be as defined preceding Corollary 8. Using Lemma 22 we establish the following lower bound.

Lemma 23 *For any given set of unit squares $\mathcal{S}$, $|\chi_{\text{opt}}(\mathcal{S})| = \Omega\left(\log\left(\max_{T,\gamma}\{m_\gamma(T)\}\right)\right)$.*

Proof: Consider a tile T and a corner type γ such that $m_\gamma(T)$ is maximal. Let W denote the tile $T_{\text{op}(\gamma)}$, which selected the squares in $sel_\gamma(T)$. Let Q denote the region remaining in W when the γ-corner chain and op(γ)-corner chain are considered in W. Assume, without loss of generality, that $\gamma = \urcorner$. The set $sel_{\urcorner}(T)$ consists of two kinds of squares: (i) squares that belong to the a minimal cover $\widetilde{\mathcal{S}}^{\min}_{\urcorner}$ of $Q_{\urcorner} \setminus Q_{\llcorner}$ (these squares are selected by the greedy algorithm), and (ii) pairs of squares that were selected according to the procedure defined in Lemma 21. Let m' denote the number of squares of the first kind, and let m'' denote the number of squares of the second kind.

We consider first the case that $m' \geq m''$. The separation procedure of adjacent corner chains combined with Lemma 19 implies that $\widetilde{\mathcal{S}}^{\min}_{\urcorner}$ is a chain with respect to $W \setminus \{S : S \notin \mathcal{S}(W, \urcorner)\}$. By Lemma 22 we get that the number of required colors is $\Omega(\log m') = \Omega(m_{\urcorner}(T))$.

We next consider the case that $m'' > m'$. Let V denote the tile $W_{\llcorner}$. The squares in $sel_{\llcorner}(V)$ were also selected by the tile W. The number of squares in $sel_{\llcorner}(V)$ that belong to a minimal cover $\widetilde{\mathcal{S}}^{\min}_{\llcorner}$ of $Q_{\llcorner} \setminus Q_{\urcorner}$ is at least $m''/2$. By applying the same argument now on the tile V and the $\llcorner$-corner, the lemma follows. $\square$

7.3.4 Wrapping-up the Proof of Theorem 3 for Unit Squares

Combining Corollary 8 and Lemma 23, and noting that the computational complexity of the algorithm is only due to sorting squares according their 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 algorithm for CF-coloring $\mathcal{R}$ that uses $O((\log \rho)^2) \cdot |\chi_{\text{opt}}(\mathcal{R})|$ colors. This means that if the size-ratio ρ is constant, then the algorithm is a constant-ratio approximation algorithm.

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 $1 \leq i, j < \log \rho + 1$, the class $\mathcal{R}^{i,j}$ consists of rectangles whose width is in the interval $[2^{i-1}, 2^i)$, and whose height is in the interval $[2^{j-1}, 2^j)$. Each class is colored using a distinct palette, to obtain a CF-coloring that uses $O((\log \rho + 1)^2) \cdot |\chi_{\text{opt}}(\mathcal{R})|$ 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. For a corner-type γ, and a non-bare orphan tile W, let $\mathcal{R}'(W,\gamma)$ be the subset of rectangles selected by W whose γ-corner resides in W. Finally, let $m = \max_{W,\gamma}\{|\mathcal{R}'(W,\gamma)|\}$ denote the maximum (over all tiles W and corner types γ) of the number of rectangles selected by an orphan tile W due to their participation in a γ-corner-chain within the tile. In this section we show that for every non-orphan tile T: (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 $\bigcup_{S\in sel(T)} S$ intersects at most 25 tiles. Therefore, $|sel(T)| = O(m)$, for every tile T.

Similarly to Lemma 23, $|\chi_{\text{opt}}(\mathcal{R})| = \Omega(\log m)$. 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.

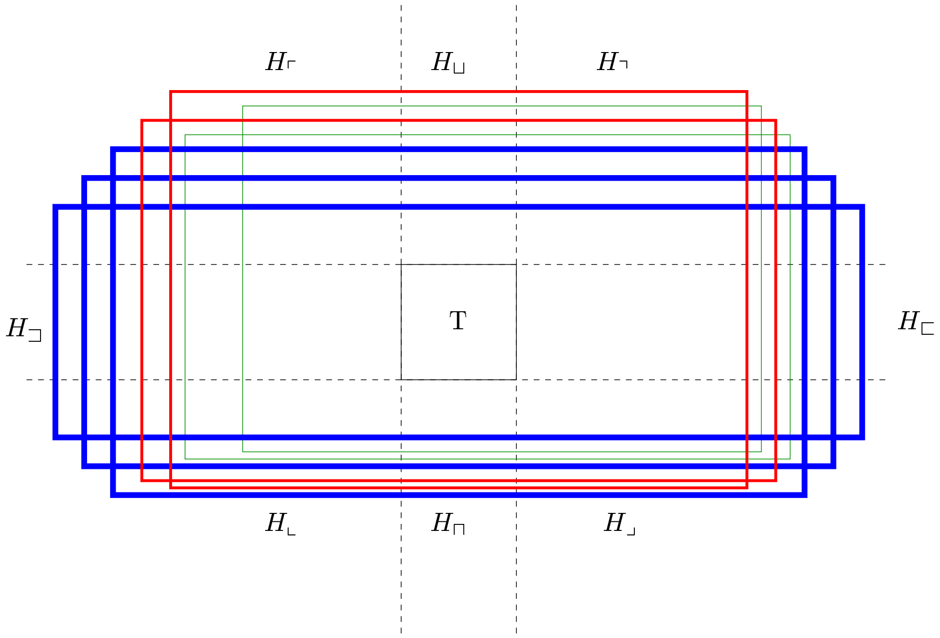

Figure 22: An illustration for the proof of Lemma 24. The thickest rectangles belong to all 4 chains. The second thickest rectangles belong to the top-left and top-right chains, and the thinnest rectangles belong only to the top-right chain.

Lemma 24 *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.*

Proof: Let T be the $1/2 \times 1/2$ tile that contains the centers of the rectangles 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. 22.

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 $\widetilde{\mathcal{R}}'_{\llcorner}$, $\widetilde{\mathcal{R}}'_{\lrcorner}$, $\widetilde{\mathcal{R}}'_{\ulcorner}$ and $\widetilde{\mathcal{R}}'_{\urcorner}$, 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 ($\llcorner$ in $H_{\llcorner}$, $\lrcorner$ in $H_{\lrcorner}$ and in general γ in H_γ) 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 $\widetilde{\mathcal{R}}'_\gamma$ 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 to the same nonempty 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 $a \neq b$).

As a consequence, if two rectangles 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. 23.) In the same fashion, we define the type of an edge of the hexagon.

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 24). 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.

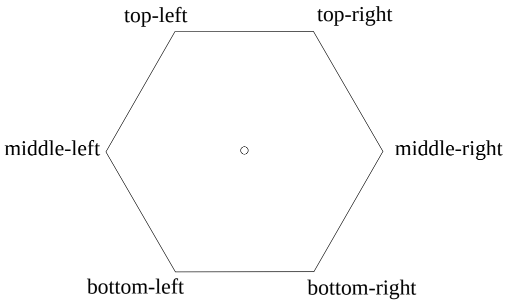

Figure 23: A hexagon and its vertices.

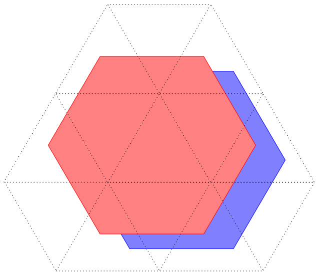

Figure 24: The triangular tiling and a pair of hexagons. The tile borders are depicted by dotted lines. Both centers of the hexagons are in the middle tile. Both hexagons completely cover the middle tile. Larger hexagons may intersect more than the 12 neighboring tiles.

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 $e_1 \cap e_2$.*

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. 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 non-orphan tile 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. nonempty 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 hexagons 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 hexagon 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. For an illustration, see Figure 25.

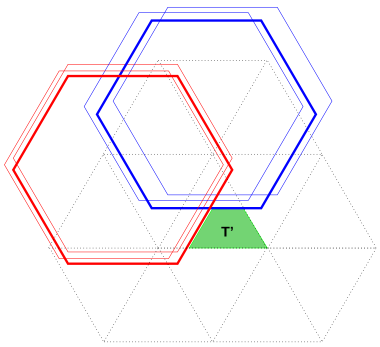

Figure 25: Let T be the central triangular tile in the figure, which is an orphan tile. The figure illustrates the choice of hexagons that intersect T with an edge. The selected hexagons are the two thick hexagons, and the region $T' \subset T$ that remains after their selection is filled.

We denote by T' the region contained in T that remains after this choice of at most 6 hexagons (on per edge type). Note that if T' is nonempty, then T' a polygon with at least 3 edges and at most 6 edges. The edges of Q are parallel to those of the hexagons in $\mathcal{H}$.

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 T'. Similarly to the analysis in the case of rectangles, the latter cover all of the intersection of T' with the former, and furthermore, they constitute a (corner) chain with respect to T'. We refer to the chain in terms of the vertex type (e.g., top-right chain). For an illustration see Figure 26.

Thus, there are at most 6 corner chains intersecting T', 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).

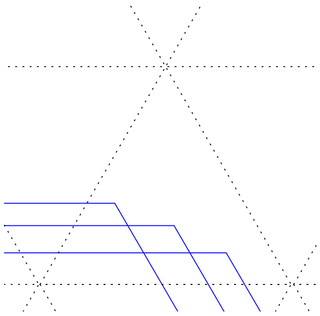

Figure 26: An example of a top-right chain. For simplicity, in this figure, $T' = T$. That is, the dotted triangle is a tile T.

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 T' minus the two hexagons. For an illustration see Figure 27.

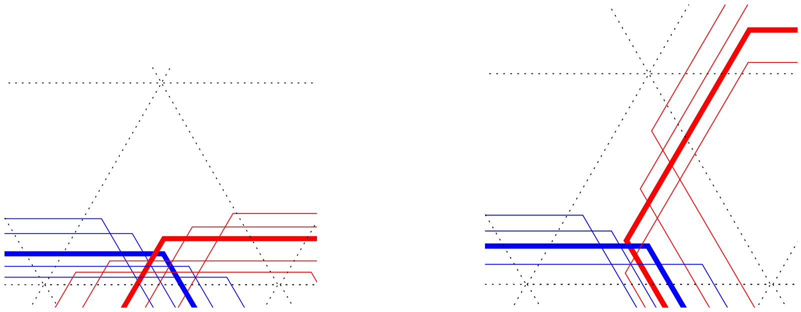

Figure 27: The two "benign" interactions between corner chains in a tile T. On the left side is a distance-one interaction between a top-right chain and a top-left chain. On the right is a distance-two interaction between a top-right chain and a middle-left chain. In each figure, the two bold hexagons are those selected from the two chains.

It follows that by selecting at most 4 hexagons from each of the 6 chains that intersect T', it is possible to service all areas of intersections between such pairs of subsets of hexagons. Let $T'' \subseteq T'$ denote the remaining region in $T' \subseteq T$ that is not covered by these selected 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 minimum subset of hexagons from the two chains so as to serve all the area in their union (within the region T''). For an illustration see Figure 28.

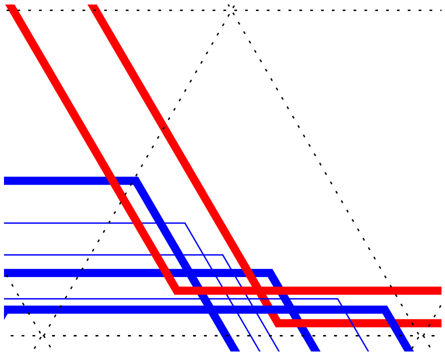

Figure 28: An interaction between the distance-three (opposite) chains top-right and bottom-left. The selected hexagons are bold.

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 25 *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.*

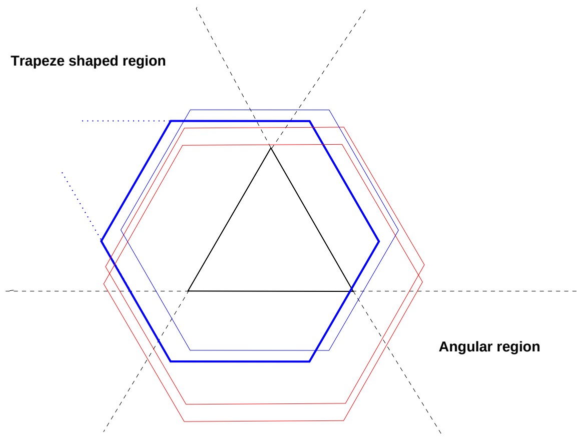

Figure 29: The partitioning of the area covered by hexagons that belong to the same tile. The 6 subregions H_R outside the tile are determined by the dashed lines that are extensions of sides of the tile. There are three angular regions, which can each be served by a single hexagon, and 3 trapeze-shaped regions. The two dotted lines within the top-left trapeze-shaped region are determined by the hexagon selected as in the proof (sketch) of Lemma 25. Each of the two dotted lines, paired with one of the dashed lines bounding the trapeze-shaped region, define the (angular) subregion that contains a (disjoint) corner-chain.

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 $\bigcup_{H\in\widetilde{\mathcal{H}}} H$).

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 29). 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 $(R \setminus H_R) \cap \left(\bigcup_{H\in\widetilde{\mathcal{H}}} H\right)$ 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 24, 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 noncongruent rectangles and hexagons

As a corollary of Lemma 24 we also obtain a universal bound for the case of rectangles that is analogous to the case of disks (Part 1 of Theorem 1).

Specifically, for each pair of integers $i, j \geq 1$, let $\mathcal{R}^{i,j}$ denote the subset of rectangles in $\mathcal{R}$ whose width is in the range $[2^{i-1}, 2^i)$, and whose height is in the range $[2^{j-1}, 2^j)$. Let $\phi_{2^i,2^j}(\mathcal{S}^{i,j})$ denote the maximum number of centers of rectangles in $\mathcal{R}^{i,j}$ that are contained in a rectangular tile of width 2^i and height 2^j. We refer to $\phi_{2^i,2^j}(\mathcal{R}^{i,j})$ as the *local density* of $\mathcal{R}^{i,j}$ (with respect to rectangular tiles of width 2^i and height 2^j).

Theorem 9 *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 $O\left(\min\left\{\sum_{i=1}^{\log(\rho)+1}\sum_{j=1}^{\log(\rho)+1}(1+\log\phi_{2^i,2^j}(\mathcal{R}^{i,j})), \log|\mathcal{R}|\right\}\right)$ colors.*

Lemma 25 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 10 *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 the reviewers for careful reading and many helpful remarks. 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.

[AH77a] K. Appel and W. Haken, "Every planar map is 4-colorable - 1: Discharging", Illinois Journal of Mathematics, 21:421-490, 1977.

[AH77b] K. Appel and W. Haken, "Every planar map is 4-colorable - 2: Reducibility", Illinois Journal of Mathematics, 21:491-567, 1977.

[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 and simple regions in the plane", in Proc. 19th Annu. ACM Sympos. Comput. Geom., 2003, to appear.

[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, "Conflict free colorings", *Discrete and Comput. Geom.*, The Goodman-Pollack Festschrift, 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, 2003.

[RSST96] R. Robertson, D. P. Sanders, P. Seymour, and R. Thomas, "Efficiently four-coloring planar graphs", In Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, 571-575, 1996.