Dynamic conflict-free colorings in the plane

Citation for published version (APA):

de Berg, M., & Markovic, A. (2019). Dynamic conflict-free colorings in the plane. *Computational Geometry*, 78, 61-73. <https://doi.org/10.1016/j.comgeo.2018.10.005>

Document license:
TAVERNE

DOI:
[10.1016/j.comgeo.2018.10.005](https://doi.org/10.1016/j.comgeo.2018.10.005)

Document status and date:
Published: 01/06/2019

Document Version:
Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Dynamic conflict-free colorings in the plane [☆]

Mark de Berg, Aleksandar Markovic ^{*}

TU Eindhoven, the Netherlands

ARTICLE INFO

Article history:

Received 26 March 2018

Accepted 3 September 2018

Available online 24 October 2018

Keywords:

Conflict-free colorings

Dynamic data structures

ABSTRACT

We study dynamic conflict-free colorings in the plane, where the goal is to maintain a conflict-free coloring (CF-coloring for short) under insertions and deletions.

- First we consider CF-colorings of a set $\mathcal{S}$ of unit squares with respect to points. Our method maintains a CF-coloring that uses $O(\log n)$ colors at any time, where n is the current number of squares in $\mathcal{S}$, at the cost of only $O(\log n)$ recolorings per insertion or deletion of a square. We generalize the method to rectangles whose sides have lengths in the range $[1, c]$, where c is a fixed constant. Here the number of colors used becomes $O(\log^2 n)$. The method also extends to arbitrary rectangles whose coordinates come from a fixed universe of size N , yielding $O(\log^2 N \log^2 n)$ colors. The number of recolorings for both methods stays in $O(\log n)$.
- We then present a general framework to maintain a CF-coloring under insertions for sets of objects that admit a unimax coloring with a small number of colors in the static case. As an application we show how to maintain a CF-coloring with $O(\log^3 n)$ colors for disks (or other objects with linear union complexity) with respect to points at the cost of $O(\log n)$ recolorings per insertion. We extend the framework to the fully-dynamic case when the static unimax coloring admits weak deletions. As an application we show how to maintain a CF-coloring with $O(\sqrt{n} \log^2 n)$ colors for points with respect to rectangles, at the cost of $O(\log n)$ recolorings per insertion and $O(1)$ recolorings per deletion.

These are the first results on fully-dynamic CF-colorings in the plane, and the first results for semi-dynamic CF-colorings for non-congruent objects.

© 2018 Published by Elsevier B.V.

1. Introduction

Consider a set of base stations in the plane that can be used for mobile communication. To ensure a good coverage, the base stations are typically positioned in such a way that the communication ranges of different base stations overlap. However, if a user is within range of several base stations using the same frequency, then interference occurs and the communication is lost. Therefore, we want to assign frequencies to the base stations such that any user within range of at least one base station, is also within range of at least one base station using a frequency where no interference occurs. The easy solution would be to give all stations a different frequency. However, this is undesirable as the set of available frequencies is limited. The question then arises: how many different frequencies are needed to ensure that any user that is

[☆] MdB and AM are supported by the Netherlands' Organisation for Scientific Research (NWO) under project no. 024.002.003.

^{*} Corresponding author.

E-mail address: a.markovic@tue.nl (A. Markovic).

within range of at least one base station has an interference-free base station at his disposal? Motivated by this and other applications, Even et al. [13] and Smorodinsky [16] introduced the notion of *conflict-free colorings* or *CF-colorings* for short. Here the ranges of the base stations are modeled as regions (disks, or other objects) in the plane, and different frequencies are represented by different colors. A CF-coloring is now defined as follows.

Let $\mathcal{S}$ be a set of objects in the plane. For a point $q \in \mathbb{R}^2$, let $S_q := \{S \in \mathcal{S} | q \in S\}$ be the subset of objects containing q . A coloring $col : \mathcal{S} \rightarrow \mathbb{N}$ of the objects in $\mathcal{S}$ —here we identify colors with non-negative integers—is said to be *conflict-free (with respect to points)* if for each point q with $S_q \neq \emptyset$ there is an object $S \in S_q$ whose color is unique among the objects in S_q . A CF-coloring is called *unimax* when the maximum color in S_q is unique.

We can also consider a dual version of planar CF-colorings. Here we are given a set $\mathcal{S}$ of points and a family $\mathcal{F}$ of geometric ranges, and the goal is to color the points in $\mathcal{S}$ such that any range from $\mathcal{F}$ containing a least one point, contains a point with a unique color. Both versions of CF-colorings in the plane—coloring objects with respect to points, and coloring points with respect to ranges—can be formulated as coloring nodes in a hypergraph such that any hyperedge has a node with a unique color. In this paper we stick to the more intuitive geometric view.

Conflict-free colorings have received a lot of attention since they were introduced by Even et al. [13] and Smorodinsky [16]; see the overview paper by Smorodinsky [17], which surveys the work up to 2010. We review the work most relevant to our results.

Even et al. proved that it is always possible to CF-color a set of disks in the plane using $O(\log n)$ colors, and that $\Omega(\log n)$ colors are needed in the worst case. The authors extended the result to sets of translates of any given centrally symmetric polygon. Later, Har-Peled and Smorodinsky [14] further generalized the result to regions with near-linear union complexity. The dual version of the problem was also studied by Even et al. [13]; they showed it is possible to CF-color points using $O(\log n)$ colors with respect to disks, or with respect to scaled translations of a centrally symmetric convex polygon. Moreover, Ajwani et al. [1] showed how to CF-color points with respect to rectangles; the bound however goes up to $O(n^{0.382})$.

Recall that CF-colorings correspond to interference-free frequency assignments in a cellular network. When a node in the network fails, the resulting assignment may no longer be interference-free. This leads to the study of *k-fault-tolerant* CF-colorings, where we want $\min(k, |S_q|)$ objects from S_q to have a unique color. In other words, a *k-fault-tolerant* CF-coloring allows the deletion of k objects without losing the conflict-free property. Cheilaris et al. [5] studied the 1-dimensional case, and presented a polynomial-time algorithm with approximation ratio $5 - \frac{2}{k}$ for the problem of finding a CF-coloring with a minimum number of colors. For $k = 1$ —that is, the regular CF-coloring—the algorithm gives a 2-approximation. Horev et al. [15] studied the 2-dimensional case and proved a $O(k \log n)$ bound for disks and, more generally, regions with near-linear union complexity.

To increase coverage or capacity in a cellular network it may be necessary to increase the number of base stations. This led Fiat et al. [7] to study *online* CF-colorings. Here the objects to be CF-colored arrive over time, and as soon as an object appears it must receive a color which cannot be changed later on. For CF-coloring points with respect to intervals, they proposed a deterministic algorithm using $O(\log^2 n)$ colors as well as two randomized algorithms, one of which uses at most $O(\log n \log \log n)$ colors in expectation and always produces a valid coloring. Later, Chen et al. [6] improved the bound with an algorithm using an expected $O(\log n)$ colors. Chen et al. [8] considered the 2-dimensional online problem in the dual setting, i.e., coloring points with respect to geometric ranges. They showed that for ranges that are half-planes, unit disks, or bounded-size rectangles—i.e. rectangles whose heights and widths all lie in the range $[1, c]$, for some fixed constant c —one can obtain an online CF-coloring using $O(\log n)$ colors with high probability. For bounded-size rectangles they also presented a deterministic result for online coloring, using $O(\log^3 n)$ colors. Bar-Noy et al. [3] provided a general strategy for online CF-coloring of hypergraphs. Their method uses $O(k \log n)$ colors with high probability, where k is the so-called *degeneracy* of the hypergraph. Their method can for instance be applied for points with respect to half-planes using $O(\log n)$ colors, which implies [8] the same result for unit disks with respect to points. They also introduced a deterministic algorithm for points with respect to intervals in $\mathbb{R}^1$ if *recolorings* are allowed. Their method uses at most $n - \log n$ recolorings in total; they did not obtain a bound on the number of recolorings for an individual insertion. Note that the results for online colorings in $\mathbb{R}^2$ are rather limited: for the primal version of the problem—online CF-coloring objects with respect to points—there are essentially only results for unit disks or unit squares (where the problem is equivalent to the dual version of coloring points with respect to unit disks and unit squares, respectively). Moreover, most of the results are randomized.

De Berg et al. [11] introduced the *fully dynamic* variant of the CF-coloring problem, which generalizes and extends the fault-tolerant and online variants. Here the goal is to maintain a CF-coloring under insertions and deletions. It is easy to see that if we allow deletions and we do not recolor objects, we may need to give each object in $\mathcal{S}$ its own color. (Indeed, any two intersecting objects must have a different color when all other objects are deleted.) Using n colors is clearly undesirable. On the other hand, recoloring all objects after each update—using then the same number of colors as in the static case—is not desirable either. Thus the main question is which trade-offs can we get between the number of colors and the number of recolorings? De Berg et al. proved a lower bound on this trade-off for the 1-dimensional problem of CF-coloring intervals with respect to points. (For this case it is straightforward to give a static CF-coloring with only three colors.) Their lower bound implies that if we insist on using $O(1/\varepsilon)$ colors, we must sometimes re-color $\Omega(\varepsilon n^\varepsilon)$ intervals, and that if we allow only $O(1)$ recolorings we must use $\Omega(\log n / \log \log n)$ colors in the worst case. They also presented a strategy that uses $O(\log n)$ colors at the cost of $O(\log n)$ recolorings. The main goal of our paper is to study fully dynamic CF-colorings for the 2-dimensional version of the problem.

Our contributions. In Section 2 we give an algorithm for CF-coloring unit squares using $O(\log n)$ colors and $O(\log n)$ recolorings per update. Note that $\Omega(\log n)$ is a lower bound on the number of colors for a CF-coloring of unit squares even in the static case, so the number of colors our fully dynamic method uses is asymptotically optimal. We also present an adaptation for bounded-size rectangles which uses $O(\log^2 n)$ colors, while still using $O(\log n)$ recolorings per update. The method also extends to arbitrary rectangles whose coordinates come from a fixed universe of size N , yielding $O(\log^2 N \log^2 n)$ colors, still at the cost of $O(\log n)$ recolorings per insertion or deletion. These constitute the first results on fully-dynamic CF-colorings in $\mathbb{R}^2$.

In Section 3, we give two general approaches that can be applied in many cases. The first uses a static coloring to solve insertions-only instances. It can be applied in settings where the static version of the problem admits a unimax coloring with a small number of colors. The method can for example be used to maintain a CF-coloring for pseudodisks with $O(\log^3 n)$ colors and $O(\log n)$ recolorings per update, or to maintain a CF-coloring for fat regions. This is the first result for the semi-dynamic CF-coloring problem for such objects: previous online results for coloring objects with respect to points in $\mathbb{R}^2$ only applied to unit disks or unit squares. We extend the method to obtain a fully-dynamic solution, when the static solution allows what we call weak deletions. We can apply this technique for instance to CF-coloring points with respect to rectangles, using $O(\sqrt{n} \log^2 n)$ colors and $O(\log n)$ recolorings per insertion and $O(1)$ recolorings per deletion.

2. Dynamic CF-colorings for unit squares and rectangles

In this section we explain how to color unit squares with $O(\log n)$ colors and $O(\log n)$ recolorings per update. We then generalise this coloring to bounded-size rectangles, and to rectangles with coordinates from a fixed universe. We first explain our basic technique on so-called anchored rectangles.

2.1. A subroutine: maintaining a CF-coloring for anchored rectangles

We say that a rectangle r is *anchored* if its bottom-left vertex lies at the origin. Let S be a set of n anchored rectangles. We denote the x - and y -coordinate of the top-right vertex of a rectangle r by r_x and r_y , respectively. Our CF-coloring of S is based on an augmented red-black tree, as explained next.

To simplify the description we assume that the x -coordinates of the top-right vertices (and, similarly, their y -coordinates) are all distinct—extending the results to degenerate cases is straightforward. We store S in a red-black tree $\mathcal{T}$ where r_x (the x -coordinate of the top-right vertex of r) serves as the key of the rectangle $r \in S$. It is convenient to work with a *leaf-oriented* red-black tree, where the keys are stored in the leaves of the tree and the internal nodes store splitting values.¹ We can assume without loss of generality that the splitting values lie strictly in between the keys.

For a node $v \in \mathcal{T}$, let $\mathcal{T}_v$ denote the subtree rooted at v and let $\mathcal{S}(v)$ denote the set of rectangles stored in the leaves of $\mathcal{T}_v$. We augment $\mathcal{T}$ by storing a rectangle $r_{\max}(v)$ at every (leaf or internal) node of v , defined as follows:

$$r_{\max}(v) := \text{the rectangle } r \in \mathcal{S}(v) \text{ that maximizes } r_y.$$

Let $\text{left}(v)$ and $\text{right}(v)$ denote the left and right child, respectively, of an internal node v . Notice that $r_{\max}(v)$ is the rectangle whose top-right vertex has maximum y -value among $r_{\max}(\text{left}(v))$ and $r_{\max}(\text{right}(v))$, so $r_{\max}(v)$ can be found in $O(1)$ time from the information at v 's children. Hence, we can maintain the extra information in $O(\log n)$ time per insertion and deletion [9].

Next we define our coloring function. To this end we define for each rectangle $r \in S$ a set $N(r)$ of nodes in $\mathcal{T}$, as follows.

$$N(r) := \{v \in \mathcal{T} : v \text{ is the leaf storing } r, \\ \text{or } v \text{ is an internal node with } r_{\max}(\text{right}(v)) = r\}.$$

Observe that $N(r)$ only contains nodes on the search path to the leaf storing r and that $N(r) \cap N(r') = \emptyset$ for any two rectangles $r, r' \in S$. Let $\text{height}(v)$ denote the height of $\mathcal{T}_v$. Thus $\text{height}(v) = 0$ when v is a leaf, and for non-leaf nodes v we have $\text{height}(v) = \max(\text{height}(\text{left}(v)), \text{height}(\text{right}(v)) + 1)$. We now define the color of a rectangle $r \in S$ as

$$\text{col}(r) := \max_{v \in N(r)} \text{height}(v).$$

Since $N(r)$ always contains at least one node, namely the leaf storing r , this is a well-defined coloring. Note that since red-black trees are updated over time, the color defined changes and requires recolorings. Fig. 1 depicts the coloring on a small instance.

Lemma 1. *The coloring defined above is conflict-free.*

¹ Such a leaf-oriented red-black tree can be seen as a regular red-black tree on a set $X'(S)$ that contains a splitting value between any two consecutive keys. Hence, all the normal operations can be done in the standard way.

Fig. 1. An instance of rectangles with the red-black tree. All rectangles of S are stored as leaves of the red-black tree $\mathcal{T}$ using their x -coordinate as key. Then each internal node of the tree is labelled with the rectangle among its two children with the highest y -coordinate. For any rectangle r , the set $N(r)$ is defined as the set of nodes of $\mathcal{T}$ that store r . Finally, each rectangle r receives color of the highest node of $N(r)$. In this example, red corresponds to height 4, blue to height 3, green to height 2, and yellow to height 1. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

Proof. Recall that $S(v)$ denotes the set of rectangles stored in the subtree rooted at v . We prove by induction on $\text{height}(v)$ that the coloring of $S(v)$ is conflict-free. Since $S = S(\text{root}(\mathcal{T}(S)))$, this proves the lemma.

When $\text{height}(v) = 0$ then $S(v)$ is a singleton, which is trivially colored conflict-free. Now assume $\text{height}(v) > 0$. Let $x(v)$ denote the splitting value stored at v , and consider any point $q := (q_x, q_y)$ in the plane. Let $S_q(v) \subseteq S(v)$ denote the set of rectangles from $S(v)$ containing q , and assume $S_q(v) \neq \emptyset$.

If $q_x > x(v)$ then q does not lie in any of the rectangles in $S(\text{left}(v))$, and so $S_q(v) = S_q(\text{right}(v))$. Since the coloring of $S(\text{right}(v))$ is conflict-free by induction, this implies that $S_q(v)$ has a rectangle with a unique color.

Now suppose $q_x \leq x(v)$. If $q_y > r_{\max}(\text{right}(v))_y$ then q does not lie in any rectangle from $S(\text{right}(v))$. Since the coloring of $S(\text{left}(v))$ is conflict-free, this implies that $S_q(v)$ has a rectangle with a unique color. Otherwise if $q_x > x(v)$, then $q \in r_{\max}(\text{right}(v))$.

Claim 1. The rectangle $r_{\max}(\text{right}(v))$ has a unique color in $S(v)$ and, hence, in $S_q(v)$.

Proof. Let u be the node that defines the color of $r_{\max}(\text{right}(v))$, that is, the node in $N(r_{\max}(\text{right}(v)))$ with maximum height. Since $v \in N(r_{\max}(\text{right}(v)))$, either $u = v$ or u is an ancestor of v . Let r be any other rectangle in $S(v)$ and let w be the node that defines the color of r . Because $r \in S(v)$, we know that w is a node in $\mathcal{T}_v$ or an ancestor of v . In the former case, since $N(r)$ is disjoint from $N(r_{\max}(\text{right}(v)))$ and hence does not contain v , we conclude that $\text{col}(r) < \text{col}(r_{\max}(\text{right}(v)))$. In the latter case we observe that u and w both are nodes on the path from the root node to v , which means that $\text{height}(u) \neq \text{height}(w)$ and so $\text{col}(r) \neq \text{col}(r_{\max}(\text{right}(v)))$. $\square$

We conclude that the coloring is conflict-free. $\square$

We obtain the following theorem.

Theorem 2. Let S be a set of anchored rectangles in the plane. Then it is possible to maintain a CF-coloring on S with $O(\log n)$ colors using $O(\log n)$ recolorings per insertions and deletion, where n is the current number of rectangles in S .

Proof. Consider the coloring method described above. Lemma 1 states that the coloring is conflict-free. Since red-black trees have height $O(\log n)$, the number of colors used is $O(\log n)$ as well.

Now consider an update on S . The augmented red-black tree can be updated in $O(\log n)$ time in a standard manner [9]. The color of a rectangle $r \in S$ can only change when (i) the set $N(r)$ changes, or (ii) the height of a node in $N(r)$ changes. We argue that this only happens for $O(\log n)$ rectangles. Consider an insertion; the argument for deletions is similar. In the first phase of the insertion algorithm for red-black trees [9] a new leaf is created for the rectangle to be inserted. This may change $\text{height}(v)$ or $r_{\max}(v)$ only for nodes v on the path to this leaf, so it affects the color of $O(\log n)$ rectangles. In the second phase the balance is restored using $O(1)$ rotations. Each rotation changes $\text{height}(v)$ or $r_{\max}(v)$ for only $O(\log n)$ nodes, so also here only $O(\log n)$ rectangles are affected. $\square$

2.2. Maintaining a CF-coloring for unit squares

Let S be a set of unit squares. We first assume that all squares in S contain the origin.

A naive way to use the result from the previous section is to partition each square $s \in S$ into four rectangular parts by cutting it along the x -axis and the y -axis. Note that the set of north-east rectangle parts (i.e., the parts to the north-east

of the origin) are all anchored rectangles, so we can use the method described above to maintain (using recolorings) a CF-coloring on them. The other part types (south-east, south-west, and north-west) can be treated similarly. Thus every square $s \in \mathcal{S}$ receives four colors. If we now assign a final color to s that is the four-tuple consisting of those four colors, then we obtain a CF-coloring with $O(\log^4 n)$ colors. (This trick of using a “product color” was also used by, among others, Ajwani et al. [1].)

It is possible to improve this by using the following fact: the ordering of the x -coordinates of the top-right corners of the squares in $\mathcal{S}$ is the same as the ordering of their bottom-right (or bottom-left, or top-left) corners. This implies that instead of working with four different trees we can use the same tree structure for all part types. Moreover, even the extra information stored in the internal nodes is the same for the north-east and north-west parts, since the y -coordinates of the top-right and top-left vertices are the same. Similarly, the extra information for the south-east and south-west parts are the same. Therefore, we can modify the augmented red-black tree to store two squares per internal node instead of one:

- $s_{\max}(v) :=$ the square $s \in \mathcal{S}(v)$ that maximizes s_y ,
- $s_{\min}(v) :=$ the square $s \in \mathcal{S}(v)$ that minimizes s_y .

Next we modify our coloring function. Therefore we first redefine the set $N(s)$ of nodes for each square $s \in \mathcal{S}$:

$$N(s) := \{\text{the leaf storing } s\} \cup N_{\text{NE}}(s) \cup N_{\text{SE}}(s) \cup N_{\text{SW}}(s) \cup N_{\text{NW}}(s),$$

where

- $N_{\text{NE}}(s) := \{v \in \mathcal{T} : v \text{ is an internal node with } s_{\max}(\text{right}(v)) = s\},$
- $N_{\text{SE}}(s) := \{v \in \mathcal{T} : v \text{ is an internal node with } s_{\min}(\text{right}(v)) = s\},$
- $N_{\text{SW}}(s) := \{v \in \mathcal{T} : v \text{ is an internal node with } s_{\min}(\text{left}(v)) = s\},$
- $N_{\text{NW}}(s) := \{v \in \mathcal{T} : v \text{ is an internal node with } s_{\max}(\text{left}(v)) = s\}.$

The coloring is as follows. We now allow four colors per height-value, namely for height-value h we give colors $4h + j$ for $j \in \{0, 1, 2, 3\}$. These colors essentially correspond to the colors we would give out for the four part types. The color of a square s is now defined as

$$\text{col}(s) := \begin{cases} 0 & \text{if } \max_{v \in N(s)} \text{height}(v) = 0 \\ & (s \text{ is only stored at a leaf}), \\ 4 \cdot \max_{v \in N(s)} \text{height}(v) + j & \text{if } \max_{v \in N(s)} \text{height}(v) > 0, \end{cases}$$

where

$$j := \begin{cases} 0 & \text{if } \text{height}(s) = \max_{v \in N_{\text{NE}}(s)} \text{height}(v), \\ 1 & \text{if the condition for } j = 0 \text{ does not apply and} \\ & \text{height}(s) = \max_{v \in N_{\text{SE}}(s)} \text{height}(v), \\ 2 & \text{if the conditions for } j = 0, 1 \text{ do not apply and} \\ & \text{height}(s) = \max_{v \in N_{\text{SW}}(s)} \text{height}(v) \\ 3 & \text{otherwise (we now must have } \text{height}(s) = \max_{v \in N_{\text{NW}}(s)} \text{height}(v)).} \end{cases}$$

Similarly as before, the color changes over time and we need to recolor the squares when it happens. The following lemma can be proven in exactly the same way as Lemma 1. The only addition is that, when considering a set $\mathcal{S}(v)$, we need to make a distinction depending on in which quadrant the query point q lies. If it lies in the north-east quadrant we can follow the proof verbatim, and the other cases are symmetric.

Lemma 3. *The coloring defined above is conflict-free.*

It remains to remove the restriction that all squares contain the origin. To this end we use a grid-based method, similar to the one used by, e.g., Chen et al. [8]. Consider the integer grid, and assign each square in $\mathcal{S}$ to the grid point it contains; if a square contains multiple grid points, we assign it to the lexicographically smallest one. Thus we create for each grid point (i, j) a set $\mathcal{S}(i, j)$ of squares that all contain the point (i, j) . We maintain a CF-coloring for each such set using the method described above. Note that a square in $\mathcal{S}(i, j)$ can only intersect squares in $\mathcal{S}(i', j')$ when (i', j') is one of the eight neighboring grid points of (i, j) . Hence, when $i' = i \bmod 2$ and $j' = j \bmod 2$ we can re-use the same color set, and so we only need four color sets of $O(\log n)$ colors each.

Theorem 4. *Let S be a set of unit squares in the plane. Then it is possible to maintain a CF-coloring on S with $O(\log n)$ colors using $O(\log n)$ recolorings per insertions and deletion, where n is the current number of squares in S .*

2.3. Maintaining a CF-coloring for bounded-size rectangles

Let S be a set of *bounded-size rectangles*: rectangles whose widths and heights are between 1 and c for some fixed constant c . Note that in practice, two different base stations have roughly the same coverage, hence it makes sense to assume the ratio is bounded by some constant c .

First consider the case where all rectangles in S contain the origin. Here, the x -ordering of the top-right corners of the rectangles may be different from the x -ordering of the top-left corners as we no longer use unit squares. Therefore the trees for the east (that is, north-east and south-east) parts no longer have the same structure. Note that the x -ordering of the top-left and bottom-left corners are the same, hence only one tree suffices for the east parts, and the same holds for the west parts. Hence, we build and maintain (using recolorings) two separate trees, one for the east parts of the rectangles and one for the west parts. In the east tree we only work with the sets $N_{NE}(S)$ and $N_{SE}(S)$, and in the west tree we only work with $N_{SW}(S)$ and $N_{NW}(S)$; for the rest of the structures and colorings are defined the same as before. We then use the product coloring to obtain our bound: we give each rectangle a pair of colors—one coming from the east tree, one coming from the west tree—resulting in $O(\log^2 n)$ different color pairs.

To remove the restriction that each rectangle contains the origin we use the same grid-based approach as for unit squares. The only difference is that a rectangle in a set $S(i, j)$ can now intersect rectangles from up to $(1 + 2c)^2 - 1$ sets $S(i', j')$, namely with $i - c \leq i' \leq i + c$ and $j - c \leq j' \leq j + c$. Since c is a fixed constant, we still need only $O(1)$ color sets.

Theorem 5. *Let S be a set of bounded-size rectangles in the plane. Then it is possible to maintain a CF-coloring on S with $O(\log^2 n)$ colors using $O(\log n)$ recolorings per insertion and deletion, where n is the current number of rectangles in S .*

2.4. Maintaining a CF-coloring for rectangles with coordinates from a fixed universe

The solution can also be extended to rectangles of arbitrary sizes, if their coordinates come from a fixed universe $U := \{0, \dots, N - 1\}$ of size N . Again, from a practical point of view it makes sense as in a city for instance, the places a base station can be created are limited.

To this end we construct a balanced tree $\mathcal{T}_x$ over the universe U , and we associate each rectangle $r = [r_x, r'_x] \times [r_y, r'_y]$ to the highest node v in $\mathcal{T}_x$ whose x -value $x(v)$ is contained in $[r_x, r'_x]$. Let $S(v)$ be the set of objects associated to v . For each node $v \in \mathcal{T}_x$ we construct a balanced tree $\mathcal{T}_y(v)$ over the universe, and we associate each rectangle $r \in S(v)$ to the highest node w in $\mathcal{T}_y(v)$ whose y -value $y(w)$ is contained in $[r_y, r'_y]$. (In other words, we are constructing a 2-level interval tree [10] on the rectangles, using the universe to provide the skeleton of the tree. The reason for using a skeleton tree is that otherwise we have to maintain balance under insertions and deletions, which is hard to do while ensuring worst-case bounds on the number of recolorings.) Let $S(w)$ be the set of objects associated to a node w in any second-level tree $\mathcal{T}_y(v)$. All rectangles in $S(w)$ have a point in common, namely the point $(x(v), y(w))$. Therefore we can proceed as in the previous section, and maintain a CF-coloring on $S(w)$ with a color set of size $O(\log^2 n)$, using $O(\log n)$ recolorings per insertions and deletion.

Note that for any two nodes w, w' at the same level in a tree $\mathcal{T}_y(v)$, any two rectangles $r \in S(w)$ and $r' \in S(w')$ are disjoint (since neither contains the value stored in their lowest common ancestor). Hence, over all nodes $w \in \mathcal{T}_x(v)$ we only need $O(\log N)$ different color sets. Similarly, for any two nodes v, v' of $\mathcal{T}_x$ at the same level, any two rectangles $r \in S(v)$ and $r' \in S(v')$ are disjoint. Hence, the total number of color sets we need is $O(\log^2 N)$. This leads to the following result.

Theorem 6. *Let S be a set of rectangles in the plane, whose coordinates come from a fixed universe of size N . Then it is possible to maintain a CF-coloring on S with $O(\log^2 N \log^2 n)$ colors using $O(\log n)$ recolorings per insertions and deletion, where n is the current number of rectangles in S .*

Remark. Instead of assuming a skeleton tree and working with a fixed skeleton for our 2-level interval tree, we can also use randomized search trees. Then, assuming the adversary doing the insertions and deletions is oblivious of our structure and coloring, the tree is expected to be balanced at any point in time. Hence, we obtain $O(\log^4 n)$ colors in expectation, at the cost of $O(\log n)$ recolorings (worst-case) per update.

3. A general technique

In this section we present a general technique to obtain a dynamic CF-coloring scheme in cases where there exists a static unimax coloring. (Recall that a unimax coloring is a CF-coloring where for any point q the object from S_q with the maximum color is unique.) Our technique results in a dynamic CF-coloring that uses $O(\gamma_{um}(n) \log^2 n)$ colors, where $\gamma_{um}(n)$ is the number of colors used in the static unimax coloring, at the cost of $O(\log n)$ recolorings per update. We first describe

our technique for the case of insertions only. Then we extend the technique to the fully-dynamic setting, for the case where the unimax coloring allows for so-called weak deletions.

We remark that even though we describe our technique in the geometric setting in the plane, the techniques provided in this section can be applied in the abstract hypergraph setting as well.

3.1. An insertion-only solution

Let S be a set of objects in the plane and assume that S can be colored in a unimax fashion using $\gamma_{\text{um}}(n)$ colors, where γ_{um} is a non-decreasing function. Here it does not matter if S is a set of geometric objects that we want to CF-color with respect to points, or a set of points that we want to CF-color with respect to a family of geometric ranges. For concreteness we refer to the elements from S as objects.

Our technique to maintain a CF-coloring under insertions of objects into S is based on the *logarithmic method* [4], which is also used to make static data structures semi-dynamic. Thus at any point in time we have $\lceil \log n \rceil + 1$ sets S_i such that each set S_i , for $i = 0, \dots, \lceil \log n \rceil$, is either empty or contains exactly 2^i objects. The idea is to give each set S_i its own color set, consisting of $\gamma_{\text{um}}(2^i)$ colors. Maintaining a CF-coloring under insertions such that the *amortized* number of recolorings is small, is easy (and it does not require the coloring to be unimax): when inserting a new object we find the first empty set S_i , and we put all objects in $S_0 \cup \dots \cup S_{i-1}$ together with the new object into S_i . The challenge is to achieve a *worst-case* bound on the number of recolorings per insertion. Note that for the maintenance of data structures, it is known how to achieve worst-case bounds using the logarithmic method. The idea is to build the new data structure for S_i “in the background” and switch to the new structure when it is ready. For us this does not work, however, since we would still need many recolorings when we switch. Hence, we need a more careful approach.

When moving all objects from $S_0 \cup \dots \cup S_{i-1}$ (together with the new object) into S_i , we do not recolor them all at once but we do so over the next 2^i insertions. As long as we still need to recolor objects from S_i , we say that S_i is *in migration*. We need to take care that the coloring of a set that is in migration, where some objects still have the color from the set S_j they came from and others have already received their new color in S_i , is valid. For this we need to recolor the objects in a specific order, which requires the static coloring to be unimax as explained below. Another complication is that, because the objects in S_j that are being moved to S_i still have their own color, we have to be careful when we create a new set S_j . To avoid any problems, we need several color sets per set. Next we describe our scheme in detail.

As already mentioned, we have sets $S_0, \dots, S_\ell$, where $\ell := \lceil \log n \rceil$. Each set can be in one of three states: *empty*, *full*, or *in migration*. For each i with $0 \leq i \leq \ell$ we have $\ell - i + 1$ color sets of size $\gamma_{\text{um}}(2^i)$ available, denoted by $\mathcal{C}(i, t)$ for $0 \leq t \leq \ell - i$. The insertion of an object s into S now proceeds as follows.

1. Let i be the smallest index such that S_i is empty. Note that i might be $\ell + 1$, in which case we introduce a new set and redefine ℓ . Note that this only happens when the number of objects reaches a power of 2.
2. Set $S_i := \{s\} \cup S_0 \cup \dots \cup S_{i-1}$. Mark $S_0, \dots, S_{i-1}$ as *empty*, and mark S_i as *in migration*.
3. Take an unused color set $\mathcal{C}(i, t)$ —we argue below that at least one color set $\mathcal{C}(i, t)$ with $0 \leq t \leq \ell - i$ is currently unused—and compute a unimax coloring of S_i using colors from $\mathcal{C}(i, t)$. We refer to the color from $\mathcal{C}(i, t)$ that an object in S_i receives as its *final color* (for the current migration). Except for the newly inserted object s , we do not recolor any objects to their final color in this step; they all keep their current colors.
4. For each set S_k in migration—this includes the set we just created in Step 3—we recolor one object whose final color is different from its current color and whose final color is maximal among such objects. When multiple objects share that property, we arbitrarily choose one of them. If all objects in S_k now have their final color, we mark S_k as *full*.

Lemma 7. Suppose that when we insert an object s into S , the first empty set is S_i . Then the sets $S_0, \dots, S_{i-1}$ are full.

Proof. Suppose for a contradiction that S_j , for some $0 \leq j < i$ is in migration. Consider the last time at which S_j was created—that is, the last time at which we inserted an object s' that caused the then-empty set S_j to be created and marked as in migration. Upon insertion of s' , we already perform one recoloring in S_j . At that point all sets $S_0, \dots, S_{j-1}$ were marked empty and it takes $\sum_{t=0}^{j-1} 2^t = 2^j - 1$ additional insertions to fill them, giving us as many recolorings in S_j . Thus before we create any set S_i with $i > j$, we have already marked S_j as full. Since s' was the last object whose insertion created S_j , by the time we create S_i the set S_j must still be full—it cannot in the mean time have become empty and later be re-created (and thus be in migration). $\square$

Next we show that in Step 3 we always have an unused color set at our disposal.

Lemma 8. When we create a new set S_i in Step 3, at least one of the color sets $\mathcal{C}(i, t)$ with $0 \leq t \leq \ell - i$ is currently unused.

Proof. Consider a color set $\mathcal{C}(i, t)$. The reason we may not be able to use $\mathcal{C}(i, t)$ when we create S_i is that there is a set $S_{i'}$ with $i' > i$ that is currently in migration: the objects from a previous instance of S_i (that were put into $S_{i'}$ when we created $S_{i'}$) may not all have been recolored yet. By Lemma 7 this previous instance was full when it was put into $S_{i'}$ and

so it only blocks a single color set, namely one for $\mathcal{S}_i$. Thus the number of color sets $\mathcal{C}(i, t)$ currently in use is at most $\ell - i$. Since we have $\ell - i + 1$ such color sets at our disposal, one must be unused. $\square$

We obtain the following result.

Theorem 9. *Let $\mathcal{F}$ be a family of objects such that any subset of n objects from $\mathcal{F}$ admits a unimax coloring with $\gamma_{\text{um}}(n)$ colors, where $\gamma_{\text{um}}(n)$ is non-decreasing. Then we can maintain a CF-coloring on a set $\mathcal{S}$ of objects from $\mathcal{F}$ under insertions, such that the number of used colors is $O(\gamma_{\text{um}}(n) \log^2 n)$ and the number of recolorings per insertion is at most $\lceil \log n \rceil$, where n is the current number of objects in $\mathcal{S}$.*

Proof. The number of colors used is $\sum_{i=0}^{\ell} (\ell - i + 1) \gamma_{\text{um}}(2^i)$, where $\ell = \lceil \log n \rceil$. Since $\gamma_{\text{um}}(n)$ is non-decreasing, this is bounded by $O(\gamma_{\text{um}}(n) \log^2 n)$. The number of recolorings per insertion is at most one per set $\mathcal{S}_i$, so at most $\lceil \log n \rceil$ in total. (The total number of sets is actually $\lceil \log n \rceil + 1$, but not all of them can be in migration.)

It remains to prove that the coloring is conflict-free. Consider a point $q \in \mathbb{R}^2$. (Here we use terminology from CF-coloring of objects with respect to points. In the dual version, q would be a range.) Let $\mathcal{S}_i$ be a set containing an object s with $q \in s$; if no such set exists there is nothing to prove.

If $\mathcal{S}_i$ is full then it has a unimax coloring using a color set $\mathcal{C}(i, t)$ not used by any other set $\mathcal{S}_j$. Hence, there is an object containing q with a unique color.

Now suppose that $\mathcal{S}_i$ is in migration. We have two cases: (i) q is contained in an object from $\mathcal{S}_i$ that has already received its final color, (ii) all objects in $\mathcal{S}_i$ containing q still have their old color.

In case (i) the object containing q with the highest final color must have a unique color, because of the following easy-to-prove fact.

Fact 1. *Consider any set A colored with a unimax coloring. Let z be an integer, and let $B \subseteq A$ be a subset that contains all objects of color greater than z , some objects of color z , and at most one other object. Then the coloring of B is unimax.*

Proof. Consider a point q contained in at least one object of B . We show that the subset of B of objects containing q has exactly one object of the highest color. If q is not contained in any object from $A \setminus B$, then the object of highest color is unique since the coloring on A is unimax. Suppose now q is also contained in some objects in $A \setminus B$ and let us assume that the object $a \ni q$ with maximum color is in $A \setminus B$ (the other case is trivial). The color of a has then to be at most z . Because the coloring on A is unimax no object with color z or higher can contain q except possibly a . Since there is at most one object $b \in B$ with color smaller than z , and since we supposed q is contained in at least one object of B , the set of objects of B containing q must be the singleton $\{b\}$. Therefore the coloring is unimax on B . $\square$

This fact proves the statement above for case (i), because we recolor the objects in decreasing order of their colors and the coloring we are migrating to is a unimax coloring. (The “at most one other object” mentioned in the fact is needed because the object that caused the migration immediately receives its color, and this color needs not be the highest color.)

In case (ii), q is contained in an object from some old set $\mathcal{S}_j$ with $j < i$. At the time we created $\mathcal{S}_i$ this set $\mathcal{S}_j$ was CF-colored, and since we did not yet recolor any object from $\mathcal{S}_j$ that contains q —otherwise we are in case (i)—we conclude that q is contained in an object with a unique color. $\square$

Application: objects with near-linear union complexity. Har-Peled and Smorodinsky [14] proved that any family of objects with linear union complexity (for example disks, or pseudodisks) can be colored in a unimax fashion using $O(\log n)$ colors. In fact, their result is more general: if the union complexity is at most $n \cdot \beta(n)$ then the number of colors is $O(\beta(n) \log n)$. Note that for disks and pseudodisks we have $\beta(n) = O(1)$, for fat triangles we have $\beta(n) = O(\log^* n)$ [2] and for locally fat objects we have $\beta(n) = O(2^{O(\log^* n)})$ [2]. This directly implies the following result.

Corollary 10. *Let $\mathcal{F}$ be a family of objects such that the union complexity of any subset of n objects from $\mathcal{F}$ is at most $n\beta(n)$. Then we can maintain a CF-coloring on a set $\mathcal{S}$ of objects from $\mathcal{F}$ under insertions, such that the number of used colors is $O(\beta(n) \log^3 n)$ and the number of recolorings per insertion is $O(\log n)$, where n is the current number of objects in $\mathcal{S}$.*

3.2. A fully-dynamic solution

We now generalize the semi-dynamic solution presented above so that it can also handle deletions. As before we assume we have a family $\mathcal{F}$ of objects such that any set of n objects from $\mathcal{F}$ can be unimax-colored with $\gamma_{\text{um}}(n)$ colors. We further assume that such a coloring admits *weak deletions*: once we have colored a given set $\mathcal{S}$ of n_0 objects using $\gamma_{\text{um}}(n_0)$ colors, we can delete objects from it using $r(n_0)$ recolorings per deletion such that the number of colors never exceeds $\gamma_{\text{um}}(n_0)$. The functions $\gamma_{\text{um}}(n)$ and $r(n)$ are assumed to be non-decreasing.

Fig. 2. Two sets in migration. The set $S_i^{(j)}$ is the set S_j before the migration happened (minus the objects we already placed in S_i^*). The set S_i^* is the set of objects of S_i that already received their final color. When all objects are in S_i^* , the migration is done and S_i^* becomes the set S_i which is no longer in migration. Note that for the downwards migration, instead of $S_\ell^{(\ell+1)}$, we can have $S_\ell^{(\ell-2)}$.

Let S be the current set of objects. We again employ ideas from the logarithmic method, but we need to relax the conditions on the set sizes. More precisely, we maintain an integer ℓ and a partition of S into $\ell + 1$ sets $S_0, \dots, S_\ell$, such that the following *size invariant* is maintained.

(Inv-S) For all $0 \leq i < \ell$ we have $|S_i| \leq 2^i$, and we have $2^{\ell-2} \leq |S_\ell| \leq 2^\ell$.

Note that the second part of (Inv-S) implies that $\ell = \Theta(\log n)$, where $n := |S|$. As before, for each i with $0 \leq i \leq \ell$ we have color sets $\mathcal{C}(i, t)$ available, each of size $\gamma_{\text{um}}(2^i)$. This time the number of colors sets $\mathcal{C}(i, t)$ is $\ell + 2$ for each i , instead of $\ell - i$. Additionally, we allow the use of colors sets $\mathcal{C}(\ell + 1, t)$. Hence, for each $i = 0, \dots, \ell + 1$, we have $\ell + 2$ color sets of size $\gamma_{\text{um}}(2^i)$. A set S_i can now be in four states: *empty*, *non-empty*, *in upwards migration*, and *in downwards migration*. It is worth pointing out that only S_ℓ can be in downwards migration.

Our coloring of the sets S_i satisfies several *color invariants*, which depend on the state of S_i . The invariant for sets S_i whose state is *non-empty* is relatively straightforward.

(Inv-C-NonEmp) The objects in a set S_i whose state is *non-empty* are unimax-colored using a color set $\mathcal{C}(i, t)$ not used elsewhere.

Before we can state the invariants for sets S_i in migration we need to introduce some notation.

A set S_i that is in upwards migration is the disjoint union of subsets $S_i^{(0)}, \dots, S_i^{(i-1)}$, some of which may be empty, plus at most one other object (that was just inserted). Set $S_i^{(j)}$ is the set of objects that used to be S_j before the migration and is now migrating to S_i (and, hence, needs to be recolored). When set S_ℓ is in downwards migration, it is the disjoint union of either subsets $S_\ell^{(\ell-2)}, S_\ell^{(\ell-1)}$, and $S_\ell^{(\ell)}$, or of subsets $S_\ell^{(\ell-1)}, S_\ell^{(\ell)}$, and $S_\ell^{(\ell+1)}$. In the former case, we define $\ell' := \ell$, in the latter case $\ell' = \ell + 1$. Defining ℓ' in this way simplifies the descriptions. See Fig. 2 for an illustration of both migrations.

The idea is now that we have a “global” coloring for the set S_i —this is the new coloring we are migrating to—and separate “local” colorings for each of the sets $S_i^{(m)}$. The global color of an object is called its *final color*, and the local color is called its *temporary color*. To know which of these two colors is the actual color of an object, we maintain a (possibly empty) subset $S_i^* \subseteq S_i$ such that the actual color of an object in S_i^* is its final color, and the actual color of an object in $S_i \setminus S_i^*$ is its temporary color. Thus, when $S_i^* = S_i$ then all objects received their new color and the migration is finished.

We first give the color invariants that hold for all sets in migration, and then give additional invariants that depend on whether the set is in upwards or downwards migration.

(Inv-C-Mig-1) If S_i is in migration, then we have a unimax coloring on S_i using a color set $\mathcal{C}(i, t)$ not used elsewhere. As already mentioned, the color an object receives in this coloring is called its *final color*; which can be different from its actual color.

(Inv-C-Mig-2) If S_i is in migration, then there is an integer z such that S_i^* contains all objects from S_i whose final color is greater than z , some objects of color z , and at most one other object. (Recall that S_i^* is the set of objects that are colored with their final color.)

For sets in upwards migration we also have the following invariant.

(Inv-C-Up) If S_i is in upwards migration, then for each $S_i^{(m)}$ we have a unimax coloring using a color set $\mathcal{C}(m, t)$ not used elsewhere. The color an object receives in this coloring is its *temporary color*, which is the color of the object before the migration. (During the migration process, this object is recolored with its final color.)

Finally, when S_ℓ is in downwards migration—recall that S_i can only be in downwards migration when $i = \ell$ —we have the following invariant.

(Inv-C-Down) If S_ℓ is in downwards migration, then the following holds.

- Each set $S_\ell^{(m)}$ with $m < \ell' - 2$ is empty.
 - The set $S_\ell^{(\ell'-2)}$ is unimax colored using at most $\ell' - 1$ color sets not used elsewhere. More precisely, for each $m = 0, \dots, \ell' - 2$ at most one color set $\mathcal{C}(m, t)$ is used in the coloring of $S_\ell^{(\ell'-2)}$.
 - The set $S_\ell^{(\ell'-1)}$ is unimax colored using at most ℓ' color sets not used elsewhere. More precisely, for each $m = 0, \dots, \ell' - 1$ at most one color set $\mathcal{C}(m, t)$ is used in the coloring of $S_\ell^{(\ell'-1)}$.
 - The set $S_\ell^{(\ell')}$ is unimax colored using at most one color set $\mathcal{C}(\ell', t)$ not used elsewhere.
- The color an object receives in these colorings is its *temporary color*, which, again, is the color of the object before the migration and that is going to be changed for the final color during the migration process, this object is recolored with its final color.)

Note that, when S_ℓ is in downwards migration, the sets $S_\ell^{(\ell'-1)}$ and $S_\ell^{(\ell'-2)}$ are colored using several color sets. In fact, it is the case that the algorithm distinguishes several subsubsets of $S_\ell^{(\ell'-1)}$ (and, similarly, of $S_\ell^{(\ell'-2)}$), each of which is unimax-colored using a different color set. It is easily checked that the coloring obtained in this manner for $S_\ell^{(\ell'-1)}$ is a unimax-coloring, which allows weak deletions by doing a weak deletion on the relevant subsubset.

The next lemma implies that it is sufficient to maintain a coloring satisfying all invariants; its proof is similar to the proof of Theorem 9 that the coloring used by our insertion-only method is conflict-free, with sets marked non-empty taking the role of sets marked full.

Lemma 11. *A coloring satisfying all color invariants is conflict-free.*

Next we describe how to insert or delete an object s . We start with insertions.

1. Let i be the smallest index such that S_i is empty. Find the smallest j such that $S_1 \cup \dots \cup S_{i-1}$ fit into S_j . Note that $j \leq i$. If $j = \ell + 1$, set $\ell := \ell + 1$.
2. Set $S_j^* := \{s\}$, set $S_j^{(m)} := S_m$ for all $0 \leq m \leq i - 1$, and set $S_j := S_j^* \cup S_j^{(0)} \cup \dots \cup S_j^{(i-1)}$. Mark $S_0, \dots, S_{i-1}$ as *empty*, and then mark S_j as *in upwards migration*. Mark all colors of the objects in $S_j^{(0)} \cup \dots \cup S_j^{(i-1)}$ as *temporary*.
3. Take an unused color set $\mathcal{C}(j, t)$ and compute a unimax coloring of S_j using colors from $\mathcal{C}(j, t)$. The color that an object in S_j receives in this coloring is its *final color*. Except for the newly inserted objects s , we do not recolor any objects to their final color in this step; they all keep their temporary colors.
4. For each set S_k in migration—this includes the set we just created in Step 3—we proceed as follows.
 - If $k < \ell$, we pick one object whose final color is different from its actual color and that has the highest final colors among such objects. We recolor the object so that its actual color becomes its final color, and we add it to S_k^* .
 - If $k = \ell$, we pick two objects whose final color is different from its actual color and that have the highest final colors among such objects. We recolor them to their final color, and add them to S_k^* .

In both cases we make an arbitrary choice in case of ties, and in the second case when only one object still needs to be recolored we just recolor that one. If all objects in S_k now have their final color—thus $S_k^* = S_k$ —we mark S_k as (that is, we change its status to) *non-empty*.

We need the following lemma.

Lemma 12. *Suppose that when we insert an object s into S , the first empty set is S_i . Then the sets $S_0, \dots, S_{i-1}$ are marked non-empty.*

Proof. Suppose for a contradiction that S_k , for some $0 \leq k < i$ is in migration. If S_k is in upwards migration then the proof is similar to the proof of Lemma 7. Suppose now S_k is in downwards migration (in which case, $k = \ell$). At the moment S_ℓ is marked as being in downwards migration, we know that $S_{\ell-1}$ is empty—see the description of a deletion further down. Hence, at least $2^{\ell-1}$ objects need to be inserted before an insertion can reach S_ℓ . Since we do two recolorings per insertion, this implies that S_ℓ is no longer in migration when s is inserted, yielding a contradiction. $\square$

We also need the analog of Lemma 8.

Lemma 13. *In Step 3 of the insertion procedure, at least one of the color sets $C(j, t)$ with $0 \leq t \leq \ell + 1$ is currently unused.*

Proof. A color set $C(j, t)$ can only be already in use for sets S_m with $m \neq j$ that are in migration. By Lemma 12 we have $m > j$ for such sets. When $m < \ell$ then S_m uses at most one color set $C(j, t)$, namely for $S_m^{(j)}$; see Invariant (Inv-C-Up). Hence, there are at most $\ell - 2$ color sets $C(j, t)$ already in use by sets S_m with $m \neq \ell$. (Note that S_0 is never in migration, hence we have $\ell - 2$ instead of $\ell - 1$.) It remains to consider S_ℓ . In fact, it suffices to consider the case where S_ℓ is in downwards migration as the other cases only use at most one color set $C(j, t)$ for S_ℓ . If S_ℓ is in downwards migration, then only $S_\ell^{(\ell'-2)}$ and $S_\ell^{(\ell'-1)}$ can use a color set $C(j, t)$; see (Inv-C-Down). Then, the total number of color sets $C(j, t)$ used is at most $\ell - 2 + 2 = \ell$, leaving at least one unused color set. $\square$

We can now prove the correctness of the insertion procedure.

Lemma 14. *The insertion procedure maintains all size and color invariants.*

Proof. It is easy to check that the size invariant is maintained. Color invariant (Inv-C-NonEmp) remains true because we only create sets marked *non-empty* in Step 4 and when we do they are unimax-colored. Invariant (Inv-C-Mig-1) still holds as well, because we only create a new set in migration in Step 2 and then in Step 3 we generate a unimax coloring for it using a set not used elsewhere. Invariant (Inv-C-Mig-2) holds because of the way Step 4 works. We only need to check Invariant (Inv-C-Up) for the set S_j that is marked as being in upwards migration in Step 2, and there it holds because (Inv-C-NonEmp) holds before the insertion. Finally, (Inv-C-Down) cannot be violated because it holds before the insertion and our insertion algorithm does not mark a set as being in downwards migration. $\square$

Next we describe the deletion of an object s .

1. Let i be such that $s \in S_i$.
2. (a) If $i \neq \ell$ or ($i = \ell$ and $|S_\ell| > 2^{\ell-2}$), and in addition S_i is not in migration, then do a weak deletion of s in S_i . Mark S_i as empty if applicable.
- (b) If $i \neq \ell$ or ($i = \ell$ and $|S_\ell| > 2^{\ell-2}$), and in addition S_i is in migration, then do the following.
 - Do a weak deletion of s in $S_i^{(m)}$, where $S_i^{(m)}$ is the subset of S_i containing s . Note that this involves changing the actual color of at most $r(2^m)$ of the objects in $S_i^{(m)} \setminus S_i^*$.
 - Do a weak deletion of s on S_i , thus changing the final color of at most $r(2^i)$ objects. Note that this may break invariant (Inv-C-Mig-2). Repair (Inv-C-Mig-2) by removing at most $r(2^i)$ objects from S_i^* and putting at most $r(2^i)$ other objects from S_i into S_i^* instead. We do this such that the size of S_i^* does not change. Observe that objects added to S_i^* are recolored to their final color, while objects removed from S_i^* are recolored to their temporary color.
 Mark S_i as empty if applicable.
- (c) Otherwise we have $i = \ell$ and $|S_\ell| = 2^{\ell-2}$, so the deletion of s breaks the condition on the size of S_ℓ . In this case we merge the last three sets $S_{\ell-2}, S_{\ell-1}, S_\ell$, as follows. Set $\ell' := \ell$. If the three sets $S_{\ell-2}, S_{\ell-1}, S_\ell$ together fit into $S_{\ell-1}$ then set $\ell := \ell - 1$, otherwise keep ℓ as it is.
 - If $S_{\ell'-2}$ is empty or non-empty, set $S_\ell^{(\ell'-2)} := S_{\ell'-2}$; otherwise, set $S_\ell^{(\ell'-2)} := \bigcup_{i=0}^{\ell'-2} S_{\ell'-2}^{(i)}$. Note that in the later case, $S_\ell^{(\ell'-2)}$ is now using at most $\ell' - 2$ color sets, namely at most one color set $C(m, t)$ for each $m \leq \ell' - 2$.
 - If $S_{\ell'-1}$ is empty or non-empty, set $S_\ell^{(\ell'-1)} := S_{\ell'-1}$; otherwise, set $S_\ell^{(\ell'-1)} := \bigcup_{i=0}^{\ell'-1} S_{\ell'-1}^{(i)}$. Note that in the later case, $S_\ell^{(\ell'-1)}$ is now using at most $\ell' - 1$ color sets, namely at most one color set $C(m, t)$ for each $m \leq \ell' - 2$.
 - Set $S_\ell^{(\ell')} := S_{\ell'}$ and $S_\ell^* := \emptyset$.
 Now compute the final coloring on S_ℓ using an unused color set $C(\ell, t)$ and mark S_ℓ as being in downwards migration.
3. If $i = \ell$, do two recolorings in S_ℓ starting from the highest final colors and add the two recolored objects to S_ℓ^* . If only one such object remains, do the recoloring only on that one. If $S_\ell^* = S_\ell$, mark S_ℓ non-empty and free the unused color sets $C(\ell, t)$.

Lemma 15. Suppose that upon deletion of s from S_ℓ , Step 2c applies and thus we mark S_ℓ as in downwards migration. Then S_ℓ was not in migration just before s is deleted.

Proof. Suppose for a contradiction that S_ℓ is in migration. Let s' be the last object that caused a migration on S_ℓ . If s' was being inserted then $|S_\ell| \geq 2^{\ell-1}$, because of the choice of j in Step 1 of the insertion algorithm. If s' was being deleted then we also have $|S_\ell| \geq 2^{\ell-1}$, because of (Inv-S) and because we decrement ℓ in Step 2c of the deletion procedure when $S_{\ell-2}, S_{\ell-1}, S_\ell$ fit into $S_{\ell-1}$. Hence, before s is deleted, at least $|S_\ell|/2 - 1$ other objects have been deleted (recall that $|S_\ell| \geq 2^{\ell-1}$ when s' is inserted, and another migration occurs when $|S_\ell| < 2^{\ell-2}$), generating $|S_\ell| - 2$ recolorings for S_ℓ . This with the first two recolorings from when the migration happens is enough to recolor all objects in S_ℓ , contradicting that S_ℓ is still in migration. $\square$

Lemma 16. When we compute the final coloring of S_ℓ in Step 2c of the deletion procedure, at least one of the color sets $C(\ell, t)$ with $0 \leq t \leq \ell + 1$ is currently unused.

Proof. Before the deletion that caused the move, thanks to Lemma 15 and the fact that only the last set can be in downwards migration, no set is in downwards migration. Therefore, before the deletion, each set $S_0, \dots, S_\ell$ uses at most one color set $C(\ell, t)$. Hence at most $\ell + 1$ colors sets are being used, leaving at least one color set available. $\square$

Lemma 17. The deletion procedure maintains all size and color invariants.

Proof. Invariant (Inv-S) is maintained since when the set S_ℓ becomes too small, Step 2c redistributes the last three sets such that the invariant holds again. Invariant (Inv-C-NonEmp) is maintained by construction in Step 2a. Invariants (Inv-C-Mig-1), (Inv-C-Mig-2) and (Inv-C-Up) are maintained by construction in Step 2b. We now argue that (Inv-C-Down) is maintained. It is obvious that the first part of (Inv-C-Down) is maintained. The second and third part are maintained because (Inv-C-NonEmp) and (Inv-C-up) hold before the deletion; (Inv-C-NonEmp) is needed when $S_{\ell-2}$ resp. $S_{\ell-1}$ are empty or non-empty, otherwise we need (Inv-C-up). The last part is maintained because before the deletion, due to Lemma 15 and the fact that Invariant (Inv-C-NonEmp) holds before the deletion. $\square$

We obtain the following result.

Theorem 18. Let $\mathcal{F}$ be a family of objects such that any subset of n objects from $\mathcal{F}$ admits a unimax coloring with $\gamma_{\text{um}}(n)$ colors and that allows weak deletions at the cost of $r(n)$ recolorings, where $\gamma_{\text{um}}(n)$ and $r(n)$ are non-decreasing. Then we can maintain a CF-coloring on a set S of objects from $\mathcal{F}$ under insertions and deletions, such that the number of used colors is $O(\sum_{i=0}^k \gamma_{\text{um}}(2^i) \log n)$, where $k = \Theta(\log n)$. The number of recolorings per insertion is $O(\log n)$, and the number of recolorings per deletion is $O(r(8n) + 1)$, where n is the current number of objects in S .

Proof. The correctness of our insertion and deletion procedures follows from Lemma 11 together with Lemmas 14 and 17. The number of colors used is at most $\sum_{i=0}^{\ell+1} (\ell + 2) \gamma_{\text{um}}(2^i)$, where $\ell = \Theta(\log n)$. Since $\gamma_{\text{um}}(n)$ is non-decreasing, this is bounded by $O(\sum_{i=0}^{\ell} \gamma_{\text{um}}(2^i) \log n)$. The number of recolorings per insertion is at most 2 per set S_i , so $O(\log n)$ in total and at most $2r(2^{\ell'})$ per deletion for at most three sets, so $O(r(8n))$ in total. $\square$

Application: points with respect to rectangles. We now make use of Theorem 18 to maintain a CF-coloring of points with respect to rectangles. But first we present a simple technique to color points with respect to intervals in $\mathbb{R}^1$, which we use as a subroutine.

Lemma 19. We can maintain a unimax coloring of n points in $\mathbb{R}^1$ with respect to intervals under deletions, using $\lceil \log n_0 \rceil$ colors and at the cost of one recoloring per deletion. Here n_0 is the initial number of points.

Proof. We start with a static unimax coloring of points with respect to intervals using $\lceil \log n_0 \rceil$ colors [17]. Recoloring after deleting a point p with color i is done as follows. If both neighbors of p have a higher color then we do nothing, otherwise we pick a neighbor with color smaller than i and recolor it to i . To prove the coloring stays unimax we only need to consider intervals I containing a neighbor of p ; for other intervals nothing changed. Now consider $I \cup \{p\}$. If before the deletion the maximum color was larger than i then that color is still present and unique. Otherwise i was the unique maximum color. Now either I contains a neighbor of p that was recolored to i , or no point in I was recolored; in both cases the maximum color in I is unique. $\square$

Remark. We can also get a fully dynamic solution using $O(\log n)$ colors and $O(\log n)$ recolorings per insertion or deletion, by storing the points in a red-black tree and coloring them with their height in the tree.

We now explain how to color points in the plane with respect to rectangles. Let S be a set of points and $\mathcal{F}$ be the family of all rectangles in the plane. The following lemma shows how to perform weak deletions.

Lemma 20. *There is a conflict free coloring of n points with respect to rectangles using $O(\sqrt{n} \log n)$ colors that allows weak deletions at the cost of one recoloring per deletion.*

Proof. We first partition the point set into at most $\sqrt{n}$ subsets such that each set is monotone using Dilworth's theorem [12]. Then, each point set behaves exactly as points with respect to intervals in one dimension. Indeed, if a rectangle contains two points, since the sequence of points is monotone, it also contains all the points in between. We can then apply Lemma 19 to finish the proof. $\square$

We can directly conclude the following corollary.

Corollary 21. *Let S be a set of points in the plane and $\mathcal{F}$ be a family of rectangles. Then we can maintain a CF-coloring on S under insertions and deletions such that the number of used colors is $O(\sqrt{n} \log^2 n)$ and the number of recolorings per insertion is $O(\log n)$ and the number of recolorings per deletion $O(1)$, where n is the current number of objects in S .*

4. Concluding remarks

We studied the maintenance of a CF-coloring under insertions and deletions of objects, presenting the first fully-dynamic solution for objects in $\mathbb{R}^2$. We showed how to maintain a CF-coloring for unit squares and for bounded-size rectangles, with $O(\log n)$ resp. $O(\log^2 n)$ colors and $O(\log n)$ recolorings per update. The method extends to arbitrary rectangles with coordinates from a fixed universe of size N , yielding $O(\log^2 N \log^2 n)$ colors and $O(\log n)$ recolorings per update. We also presented general techniques for the semi-dynamic (insertion-only) and the fully-dynamic case (insertions and deletions). Our insertions-only technique can be applied to objects with near-linear union complexity, giving for instance a CF-coloring of $O(\log^3 n)$ colors for pseudodisks using $O(\log n)$ recolorings per update. This is the first results on semi-dynamic CF-colorings for this general class of objects. Our fully dynamic solution applies to any class of object on which weak deletions are possible, giving for instance a CF-coloring of $O(\sqrt{n} \log^2 n)$ colors for points with respect to rectangles at the cost of $O(\log n)$ recolorings per insertion and $O(1)$ recolorings per deletion. This constitutes the first fully-dynamic CF-coloring for objects in the plane.

References

- [1] Deepak Ajwani, Khaled Elbassioni, Sathish Govindarajan, Saurabh Ray, Conflict-free coloring for rectangle ranges using $o(n^{0.382})$ colors, *Discrete Comput. Geom.* 48 (1) (2012) 39–52.
- [2] Boris Aronov, Mark de Berg, Esther Ezra, Micha Sharir, Improved bounds for the union of locally fat objects in the plane, *SIAM J. Comput.* 43 (2014) 534–572.
- [3] Amotz Bar-Noy, Panagiotis Cheilaris, Svetlana Olonetsky, Shakhar Smorodinsky, Online conflict-free colouring for hypergraphs, *Comb. Probab. Comput.* 19 (4) (2010) 493–516.
- [4] Jon Louis Bentley, James B. Saxe, Decomposable searching problems I: static-to-dynamic transformation, *J. Algorithms* 1 (4) (1980) 301–358.
- [5] Panagiotis Cheilaris, Luisa Gargano, Adele A. Rescigno, Shakhar Smorodinsky, Strong conflict-free coloring for intervals, *Algorithmica* 70 (4) (2014) 732–749.
- [6] Chen Ke, How to play a coloring game against a color-blind adversary, in: *Proceedings of the 22nd ACM Symposium on Computational Geometry*, 2006, pp. 44–51.
- [7] Ke Chen, Amos Fiat, Haim Kaplan, Meital Levy, Jiri Matousek, Elchanan Mossel, János Pach, Micha Sharir, Shakhar Smorodinsky, Uli Wagner, Emo Welzl, Online conflict-free coloring for intervals, *SIAM J. Comput.* 36 (5) (2007) 1342–1359.
- [8] Ke Chen, Haim Kaplan, Micha Sharir, Online conflict-free coloring for halfplanes, congruent disks, and axis-parallel rectangles, *ACM Trans. Algorithms* 5 (2) (2009) 16, 1–24.
- [9] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, Clifford Stein, *Introduction to Algorithms*, 3rd edition, The MIT Press, 2009.
- [10] Mark de Berg, Otfried Cheong, Marc van Kreveld, Mark Overmars, *Computational Geometry: Algorithms and Applications*, 3rd edition, Springer-Verlag, 2008.
- [11] Mark de Berg, Tim Leijssen, Aleksandar Markovic, André van Renssen, Marcel Roeloffzen, Gerhard Woeginger, Dynamic and kinetic conflict-free coloring of intervals with respect to points, in: *Proceedings of the 28th International Symposium on Algorithms and Computation*, 2017.
- [12] Robert P. Dilworth, A decomposition theorem for partially ordered sets, *Ann. Math.* 51 (1) (1950) 161–166.
- [13] Guy Even, Zvi Lotker, Dana Ron, Shakhar Smorodinsky, Conflict-free colorings of simple geometric regions with applications to frequency assignment in cellular networks, *SIAM J. Comput.* 33 (1) (2003) 94–136.
- [14] Sarel Har-Peled, Shakhar Smorodinsky, Conflict-free coloring of points and simple regions in the plane, *Discrete Comput. Geom.* 34 (1) (2005) 47–70.
- [15] Elad Horev, Roi Krakovski, Shakhar Smorodinsky, Conflict-free coloring made stronger, in: *Proceedings of the 12th Scandinavian Symposium and Workshops on Algorithm Theory*, 2010, pp. 105–117.
- [16] Shakhar Smorodinsky, *Combinatorial Problems in Computational Geometry*, PhD thesis, Tel-Aviv University, 2003.
- [17] Shakhar Smorodinsky, Conflict-free coloring and its applications, in: I. Barany, K.J. Boroczky, G. Fejes Toth, J. Pach (Eds.), *Geometry-Intuitive, Discrete, and Convex*, 2013, pp. 331–390.