Conflict-free Chromatic Art Gallery Coverage

Author(s):
Bärtschi, Andreas; Suri, Subhash

Publication date:
2012

Permanent link:
https://doi.org/10.3929/ethz-b-000061697

Originally published in:
Leibniz International Proceedings in Informatics (LIPIcs) 14, https://doi.org/10.4230/LIPIcs.STACS.2012.160

Conflict-free Chromatic Art Gallery Coverage*

Andreas Bärtschi[1] and Subhash Suri[2]

1 **Institute of Theoretical Computer Science, ETH Zürich**
8092 Zürich, Switzerland (`andrbaer@student.ethz.ch`)
2 **Department of Computer Science, University of California**
Santa Barbara, CA 93106, USA (`suri@cs.ucsb.edu`)

Abstract

We consider a *chromatic* variant of the art gallery problem, where each guard is assigned one of k distinct colors. A placement of such colored guards is *conflict-free* if each point of the polygon is seen by some guard whose color appears exactly once among the guards visible to that point. What is the smallest number $k(n)$ of colors that ensure a conflict-free covering of all n-vertex polygons? We call this the *conflict-free chromatic art gallery problem.* The problem is motivated by applications in distributed robotics and wireless sensor networks where colors indicate the wireless frequencies assigned to a set of covering "landmarks" in the environment so that a mobile robot can always communicate with at least one landmark in its line-of-sight range without interference. Our main result shows that $k(n)$ is $O(\log n)$ for orthogonal and for monotone polygons, and $O(\log^2 n)$ for arbitrary simple polygons. By contrast, if *all* guards visible from each point must have distinct colors, then $k(n)$ is $\Omega(n)$ for arbitrary simple polygons and $\Omega(\sqrt{n})$ for orthogonal polygons, as shown by Erickson and LaValle [3].

1998 ACM Subject Classification I.3.5 Computational Geometry and Object Modeling

Keywords and phrases art gallery problem, conflict-free coloring, visibility

Digital Object Identifier 10.4230/LIPIcs.STACS.2012.160

1 Introduction

The Art Gallery Theorem is a classical result in computational geometry, first posed by Klee and proved by Chvátal [2], which says that $\lfloor n/3 \rfloor$ (point) guards are always sufficient, and sometimes necessary, to cover a simply-connected n-vertex polygon. In the last 30 years, many extensions, variations, and generalizations involving different types of guards, polygons, and visibility constraints have been investigated. (See [6] and [8], for instance.)

Besides their mathematical elegance and appeal, the interest in art gallery problems is also spurred by applications in distributed surveillance, monitoring, and robotics. In many of these applications, the "guards" are "landmarks" deployed in an environment to help provide navigation and localization service to mobile robots. The mobile device communicates with these landmarks through wireless, or other "line-of-sight" signaling mechanisms. In order for the signaling mechanism to work correctly, the different landmarks visible to the robot at any position must operate on different frequency—the robot is unable to receive the signal if multiple landmarks in its range are transmitting at the same frequency. This motivates a "chromatic" version of the art gallery theorem, where the goal is not to optimize the *number of guards*, but rather the *number of distinct colors* needed to distinguish the guards.

* Andreas Bärtschi's research was partially supported by a scholarship from the Student Exchange Office of ETH Zürich. Subhash Suri's research was supported in part by the NSF grant IIS-0904501.

29th Symposium on Theoretical Aspects of Computer Science (STACS'12).
Editors: Christoph Dürr, Thomas Wilke; pp. 160–171
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany
SYMPOSIUM ON THEORETICAL ASPECTS OF COMPUTER SCIENCE

Problem Motivation and the Results

Radio transceivers are cheap but tuning them to many different frequencies requires costly hardware. If the polygons can be covered by guards of very few *distinct colors* (frequencies), then it would enable inexpensive robot localization and navigation. This was the motivation behind the work of Erickson and LaValle [4] who sought to guard the polygon so that each point of the polygon is seen by guards of distinct colors only—that is, the robot located anywhere in the polygon is able to communicate without interference with *any* of the guards in its line-of-sight. Surprisingly, Erickson and LaValle discovered that this *strong chromatic* condition does not lead to much savings in the number of colors: there are simple polygons that require $\Omega(n)$ colors, and even monotone orthogonal polygons require $\Omega(\sqrt{n})$ colors [3].

Motivated by this negative result, we consider a weaker chromatic condition, which is sufficient for the original robotics application of interference-free communication with a guard at all locations. Specifically, we call a placement of colored guards *conflict-free* if each point of the polygon is seen by some guard whose color appears exactly once among the guards visible to that point. Thus, for any placement of the robot in the polygon, there is at least one guard that can communicate with the robot without interference. We want to determine the smallest number $k(n)$ of colors that ensure a conflict-free coloring of some guard set in all n-vertex polygons. We call this the *conflict-free chromatic art gallery problem.*

The main result of our paper is to prove that $k(n)$ is $O(\log n)$ for orthogonal and for monotone polygons, and $k(n) = O(\log^2 n)$ for arbitrary simple polygons. Thus, not only does the conflict-free coloring yield significantly smaller bounds for distinct colors, it also fulfills the hopeful vision of robotics application that a few colors suffice.

Related Work and Hypergraph Coloring

The chromatic art gallery problem is related to hypergraph coloring, where one must assign colors to the vertices of a hypergraph $H = (V, \mathcal{E})$, so that its edges, which are subsets of vertices, are appropriately colored. In the most basic form, called the proper coloring, every edge e with at least two vertices must be non-monochromatic; that is, there must be two vertices $x, y \in e$ whose colors are distinct. In the conflict-free coloring of H, every edge e must have a vertex that is uniquely colored among the vertices in e. Smorodinsky [9, 11] considers several simple geometric hypergraphs, such as those induced by disks or rectangles. For instance, the rectangle hypergraph has a finite set of axis-aligned rectangles, and each maximal subset of rectangles with a common intersection forms an hyperedge. For these hypergraphs, it is known that the conflict-free chromatic number is $\Omega(\log n)$ and $O(\log n)$ [7, 10].

To see the connection between chromatic art gallery and the hypergraph coloring, consider a guard set S, and let $\mathcal{R}$ be the set of the guards' visibility regions in the polygon. Then we have a hypergraph $H = (V, \mathcal{E})$, whose vertices correspond to S and in which a subset $S_e \subseteq S$ corresponds to an edge if there is a point p_e in the polygon contained exactly in the visibility regions of the guards in S_e and no others. A *conflict-free hypergraph coloring* of H is easily seen to be also a conflict-free coloring of the guard set S. Of course, in the chromatic art gallery, we need to simultaneously choose the guard set and color it, so it does not quite reduce to the hypergraph coloring. Even if we were to consider a fixed guard set, the visibility regions are not as well-behaved as disks or rectangles, and no non-trivial bound is known for their conflict-free chromatic number.

The previous result that is most directly relevant to our work is the mentioned version of the chromatic art gallery, with a stronger chromatic condition on the guard's coloring. This original version relates to a *strong hypergraph coloring* of the corresponding hypergraph H.

Organization

Section 2 introduces some basic definitions and concepts. In Section 3, we prove the $O(\log n)$ bound for the conflict-free coloring of orthogonal polygons, and the general proof strategy that is used later for simple polygons as well. In Section 4, we prove the $O(\log n)$ bound for monotone polygons, which is the key to establishing the $O(\log^2 n)$ upper bound for general polygons in Section 5.

2 The conflict-free chromatic art gallery problem

Let P be a simple polygon, whose boundary we denote as $\partial P \subset P$. We say that two points $p, q \in P$ are *visible* to each other if the line segment $\overline{pq}$ is a subset of P. The *visibility region* of a point p is defined as $V(p) := \{q \in P \mid q \text{ is visible from } p\}$. A finite point set $S \subset P$ is called a *guard set* if $\bigcup_{p \in S} V(p) = P$ and we call the points in S *guards*. A coloring $c\colon S \to \{1, \ldots, k\}$ of the guards with k colors is called *conflict-free* if each point $p \in P$ is seen by a guard whose color appears exactly once among all guards that see p. Let $k_{cf}(S)$ be the minimum number of colors required to color a guard set S conflict-free and let $\mathcal{S}(P)$ be the set of all guard sets of P. Then the *conflict-free chromatic guard number* of a polygon P is defined as $\chi(P) := \min_{S \in \mathcal{S}(P)} k_{cf}(S)$. We want to determine the smallest number $k(n)$ such that for all n-vertex polygons P_n we have $\chi(P_n) \leq k(n)$.

The classical art gallery theorem says that $\lfloor n/3 \rfloor$ guards are both necessary and sufficient for covering a n-vertex polygon, but the number of colors needed to ensure conflict-free covering may be significantly smaller. For instance, the construction that forces $\lfloor n/3 \rfloor$ guards (Fig. 1) only requires *two colors*. A polygon is called *orthogonal*, if its edges meet at right angles. A polygon P is called *monotone with respect to a line* ℓ if every line orthogonal to ℓ intersects the boundary of P at most twice. P is called *x-monotone* (y-monotone) if P is monotone with respect to the x-axis (respectively the y-axis).

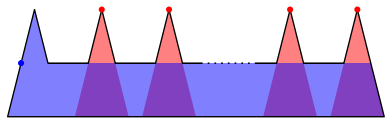

Figure 1 This polygon requires $\lfloor n/3 \rfloor$ guards but its conflict-free chromatic guard number is just 2.

The following concept of *independence* is central to our proofs, and forms a basis for coloring by partitioning into independent subpolygons.

▶ **Definition 1** (Independence). Let P be a polygon. We call two subpolygons P_1 and P_2 of P *independent* if there are no points $p_1 \in P_1$ and $p_2 \in P_2$ that are mutually visible.

▶ **Lemma 2.** *Let $\{A_1, \ldots, A_m\}$ be a partition of the polygon P into m families of pairwise independent subpolygons. That is, each $A_i = \{P_{i1}, \ldots, P_{ik_i}\}$ is a collection of subpolygons that are pairwise independent and all the subpolygons in the m families form a partition of P. Then we have* $\chi(P) \leq \sum_{i=1}^{m} \max_{P_{ij} \in A_i} \{\chi(P_{ij})\}$.

Proof. Let $\{C_1, \ldots, C_m\}$ be m disjoint color sets, where $|C_i| = \max_{P_{ij} \in A_i} \{\chi(P_{ij})\}$. Then we can guard every subpolygon $P_{ij} \in A_i$ conflict-free in itself with guards that get colors from C_i, giving a total number of $|C_1| + \ldots + |C_m|$ colors. We claim that this coloring ensures that every point $p \in P$ sees a guard of unique color among all guards that see p. To prove this claim, without loss of generality, suppose that p is contained in a subpolygon P_{ij_1} of A_i and s_1 is its guard of unique color in P_{ij_1}. Any other guard s_2 in P that has the same color as s_1 must lie in a subpolygon $P_{ij_2} \neq P_{ij_1}$, which is contained in A_i and hence independent of P_{ij_1}. Thus s_2 does not see p, and s_1 is not only a guard of unique color among all guards in P_{ij_1}, but among all guards in P. Thus, we have found a conflict-free covering with $|C_1| + \ldots + |C_m|$ colors, which completes the proof. ◀

Lemma 2 naturally suggests a divide-and-conquer strategy: we partition the polygon into four sets of subpolygons and then conquer each set by recursively splitting the regions into sets of independent regions and applying Lemma 2.

▶ Remark. We only require the interiors of subpolygons P_1 and P_2 to be independent, and allow mutual visibility among their boundary points as long as these points *also belong* to the boundary of another subpolygon that is responsible for their conflict-free covering. In particular, for a line segment e contained in two boundaries ∂P_1 and ∂P_2, we will explicitly mention whether P_1 or P_2 is "responsible" for guarding e.

3 Orthogonal Polygons

Our basic strategy is to partition the orthogonal polygon P into four types of monotone orthogonal subpolygons. These subpolygons have a boundary consisting of a single base edge and another subchain that is either x-monotone or y-monotone. The chain can be either above the base edge or below in the former case, and to the left or to the right in the latter case. We use mnemonic identifiers U (up), D (down), L (left) and R (right) to refer to these four types. When we show all or parts of the partition, we display these types with the colors red, green, black and blue, always using the following consistent mapping U $\to$ red, D $\to$ green, L $\to$ black and R $\to$ blue.

The partitioning process

Given a polygon P we construct a partition by iteratively adding monotone subpolygons. In each *odd-numbered* step we add subpolygons of Type U and D, and in each *even-numbered* step we add subpolygons of Type L and R. Figures 2 and 3 illustrate the construction.

Step 1 Let e be the lowest horizontal edge of P's boundary. Let Q be the set of all points $q \in P$ which are vertically visible from e and lie on or above e. Q is the first subpolygon in our partitioning, and it is of type U. Because P is a simply-connected region, with no holes, it is easy to see that Q splits it in parts that lie entirely to its left or entirely to its right, and each part R shares exactly one edge with Q, which is a vertical line segment.

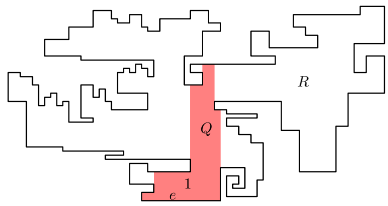

Figure 2 The first step of the partitioning process.

Step 2 The line segments on the boundary of Q become the base edges for new subpolygons of Type L and R, which are defined analogously as the first subpolygon, with vertical visibility replaced by horizontal visibility. We note that the remaining regions lie entirely above or below a subpolygon of type L or R and share exactly one horizontal line segment with these subpolygons, but not with the first subpolygon Q.

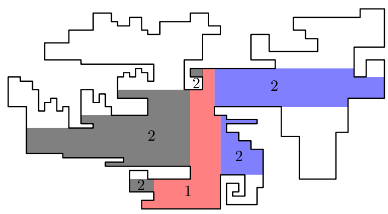

Figure 3 The second step of the partitioning process.

Step 3 The horizontal line segments from Step 2 in turn generate subpolygons of Type U and D.

We repeat steps 2 and 3 until we have a complete partition. In each odd-numbered step we construct red (U) and green (D) polygons and in each even-numbered step black (L) and blue (R) subpolygons.

▶ **Lemma 3.** *The partitioning process terminates within $n + 2$ steps.*

Proof. In each step at least one subpolygon is added to the partition. Such a subpolygon touches at least one edge $e = \{u, v\}$ previously not touched. In at most two additional steps, both the endpoints of e, u and v, become completely surrounded by subpolygons of the partition. The polygon is completely covered if all vertices are surrounded, hence the partitioning process ends after at most $n + 2$ steps. ◀

The schematic tree

The recursive partitioning generates four *families* of polygons: up-polygons A_U, down-polygons A_D, left-polygons A_L, and right-polygons A_R. Ideally, we would like to invoke Lemma 2 on this partition partitioned $\{A_U, A_D, A_L, A_R\}$. Unfortunately the subpolygons in each family are *not independent*, see Fig. 4 for an example. We, therefore, introduce a condition that allows us to subdivide the group A_U into sets of independent subpolygons. In the following, we focus exclusively on the red (up) polygon group; the other three groups are handled in the same way.

We first introduce a schematic tree that is a convenient graphical representation of the polygon partition we have. This graph is a 4-colored directed graph, where each vertex represents a subpolygon of the partition of the same color. There exists a directed edge from a subpolygon P_i to a subpolygon P_j if and only if P_j has been constructed over a line segment e that is part of P_i's boundary. As mentioned earlier, we consider e to be part of P_i but not of P_j. Since P has no holes, T contains no cycle and is a tree. The first constructed

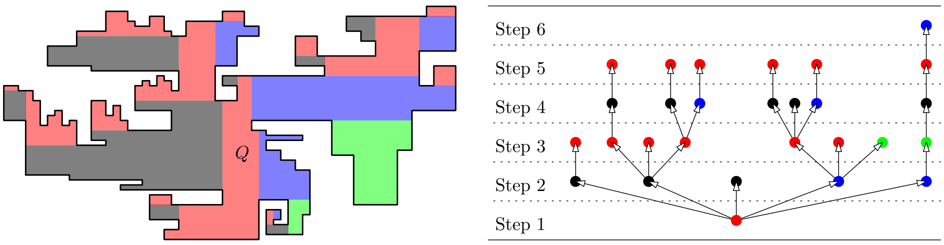

Figure 4 The complete partition and the corresponding schematic tree.

subpolygon Q has no incoming edge, and it represents the root of our tree. (The base edge of Q is considered to be a part of this subpolygon.) Since all other vertices have indegree 1, T is a rooted directed tree and any subpolygon constructed in Step k has depth $k - 1$ in T. Hence all red and green vertices have even height and all vertices of color blue or black have odd height. Therefore every directed path in T alternates between vertices of red or green color and vertices of blue or black color.

▶ Remark. Let $p_i \in P_i$ and $p_j \in P_j$ be two points of two subpolygons of the partition. Then the shortest path between p_i and p_j in P goes through a subpolygon P_k if and only if P_k lies on the shortest path between P_i and P_j in T.

▶ **Lemma 4.** *Let P be a polygon with the given partition and the schematic tree T. Let P_i and P_j be two arbitrary subpolygons of type U. Then, either (i) P_i and P_j are independent, or (ii) there exists a red-black-alternating (or a red-blue-alternating) directed path in T between P_i and P_j.*

Proof. Suppose P_i and P_j are not independent, then there exist points $p_i \in P_i$ and $p_j \in P_j$ that are mutually visible. The shortest path in P between p_i and p_j, therefore, must be a line segment. The way we included the base edges to be part of just one subpolygon excludes the possibility of the line segment being horizontal or vertical. Without loss of generality, let us assume that the line segment is directed up and to the left, with p_i at the bottom-right, and p_j at the top-left. Since P_i is a U polygon, the visibility ray $\overrightarrow{p_i p_j}$ can only leave it through its left boundary, and therefore it must enter a type L subpolygon. Next, by the upward direction of $\overrightarrow{p_i p_j}$, it can leave this L subpolygon only through a top boundary edge, which forces it to enter a U subpolygon. This process repeats until we reach P_j, showing that the sequence of polygons traversed by the shortest path from p_i to p_j is an alternating U-L sequence, which corresponds to a red-black-alternating path in T. ◀

Conquering red-black-alternating trees: Staircase and recursion

Deriving a bound on the conflict-free chromatic guard number for family A_U directly seems difficult, because of inter-dependence of the subpolygons within the family. Instead, we use the property of Lemma 4 to look at that portion of A_U that forms a *red-black-alternating tree*. That is, consider the union of the subpolygons that corresponds to a red-black alternating tree in T. Suppose P_n is such an n-vertex orthogonal polygon, namely, whose partition is a red-black alternating tree. We will cover a part of P_n with a *staircase polygon* in such a way that all other relevant parts (containing red subpolygons) are independent and proceed recursively for all of them.

Recall that a *staircase (orthogonal) polygon* is an orthogonal polygon whose boundary can be split into two subchains with alternating convex and reflex interior vertices, with the two endpoints being convex. A staircase polygon in which one of the subchains has only one interior vertex is called a *convex fan*. Convex fans are star-shaped and can clearly be guarded with one guard (and one color).

▶ **Lemma 5.** *The conflict-free chromatic number for a staircase polygon P is at most* 3.

Proof. Consider the following placement of colored guards in a staircase polygon: Starting from the top, we place a guard s_1 on the first convex vertex of the lower subchain. Then we iteratively place a guard s_{i+1} on the lowest convex vertex visible from s_i, alternating between the two subchains until the staircase polygon is covered. To each guard s_i we assign the color in $\{1, 2, 3\}$ with the same residue class as i modulo 3. One can check that the coloring is conflict-free, and a complete proof can be found in [4]. ◀

Let $f(n)$ denote the smallest number of colors that ensure a conflict-free covering of all type U subpolygons in any orthogonal P_n corresponding to a red-black-alternating tree. In other words, for every P_n there is a guard set $S \subset P_n$ that can be colored with $f(n)$ colors such that each point of a type U subpolygon is seen by some guard whose color appears exactly once among the guards visible to that point. In the following we give a placement of colored guards, which shows that $f(n)$ is $O(\log n)$.

Since P_n consists of type U and type L subpolygons, it "grows to the left". Therefore we will cover P_n with staircases ascending to the left in a natural way: Let e be a horizontal edge with two reflex vertices. We call the horizontal line through e a *decision line*, see Figure 5. A decision line splits P_n in a lower part and two or more independent upper parts, of which at most one upper part contains more than $\lfloor n/2 \rfloor$ vertices. Starting from the lowest and rightmost vertex of P_n we construct a staircase ascending to the left, which at every decision line follows the upper part with the most vertices. We guard this staircase with colors $\{1, 2, 3\}$. Furthermore at every intersection of the staircase's lower subchain with a base edge of a type U subpolygon, we insert a convex fan that is oriented to the left and to

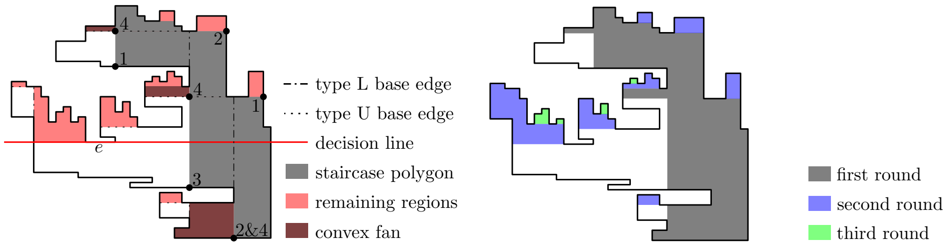

Figure 5 The first staircase subpolygon and covering with staircases and convex fans.

the top. These convex fans are bounded from the right by the staircase polygon and hence independent. We guard every convex fan with a guard of color 4 placed on the intersection. By iteratively adding staircases together with convex fans we can prove an upper bound on $f(n)$:

▶ **Lemma 6.** *Suppose P_n is an orthogonal polygon with a partition that has a red-black-alternating schematic tree. Then a conflict-free coloring of all the red regions of P_n needs at most $4 \log n$ colors. The same bound also holds for a red-blue-alternating schematic tree.*

Proof. We cover a part of the type U subpolygons with a staircase and convex fans as described. The remaining regions of the type U subpolygons are parts of smaller red-black-alternating trees. These smaller trees are all bounded from below by a decision line and from above and from the side by P_n's boundary, hence they are independent. Furthermore all of the smaller trees contain at most $\lfloor n/2 \rfloor$ of P_n's vertices because during the construction we choose at every decision line the upper part with the most remaining vertices.

Thus, the chromatic number follows the recurrence $f(n) \leq f(n/2) + 4$, which yields $f(n) \leq 4 \log n$. The same holds also for the red-blue-alternating trees by symmetry. ◀

A logarithmic upper bound for orthogonal polygons

We will show how one can cover *all* type U subpolygons in an arbitrary orthogonal polygon P_n with $O(\log n)$ colors. Let T be the schematic tree of the partition and let A and B be two disjoint color sets of size $f(n)$. We use the A and B to iteratively cover red-black-alternating and red-blue-alternating subtrees of T. In each step we must ensure that the subtrees of the same type are independent so that we can use the same colors for all of the subtrees:

Step 1 Take a not yet covered subpolygon P_s corresponding to a vertex v_s of minimal depth in T. Let T_s denote the inclusion-maximum red-black-alternating subtree rooted at v_s. By Lemma 6 we can guard all type U subpolygons corresponding to red vertices in the tree T_s with A.

Step 2 For every type U subpolygon in T_s (which now are all guarded) check whether it has red grandchildren in T that are not yet guarded (and thus must be connected through a blue vertex). These grandchildren are pairwise independent by Lemma 4, hence for each grandchild v it is possible to cover the inclusion-maximum red-blue-alternating subtree rooted at v with guards colored with colors in B conflict-free by Lemma 6. We have no conflicts with the type U subpolygons covered before since A and B are disjoint.

Step 3 As in Step 2, cover the independent inclusion-maximum red-black-alternating subtrees rooted at not yet covered red grandchildren of type U subpolygons in one of the red-blue-alternating subtrees. We use the color set A, which gives no conflicts with the guards in the red-blue-alternating subtrees, since they have colors from B. Furthermore we have also no conflicts with the guards in a previous red-black-alternating subtree by

Lemma 4, since the shortest path must go through the root of a red-blue-alternating subtree and hence through both a type L and a type R subpolygon.

Step 4 Repeat Step 2 and Step 3 as long as there are grandchildren. Otherwise we either have covered all type U subpolygons, or there remain type U subpolygons connected through a green vertex, which are thus independent by Lemma 4. In that case we start over with Step 1.

In this way, we get a conflict-free covering of all type U subpolygons in P_n with at most $|A| + |B| = 2f(n) = O(\log n)$ colors. We can apply the same procedure to type D, L and R subpolygons in alternating trees, since these cases are axis symmetric or rotationally symmetric. For each type we use two new color sets of size $f(n)$, which yields a conflict-free coloring of all subpolygons of the partition of an orthogonal polygon, where we use at most $8f(n) = O(\log n)$ colors in total. We have established the main result of this section.

▶ **Theorem 7.** *The conflict-free chromatic guard number for orthogonal polygons on n vertices is $k(n) = O(\log n)$.*

4 Monotone Polygons: a Step Towards Simple Polygons

The recursive partitioning technique of the previous section will form the basis for our proof of the general (non-orthogonal) polygons as well. However, the more complex visibility structure of non-orthogonal polygons forces us first to establish an intermediate result for *monotone* polygons. Specifically, our proof structure works by partitioning the polygon into families of simpler *staircase-shaped* subpolygons. In the orthogonal case, staircase polygons are easily covered using 3 colors (Lemma 5), but non-orthogonal staircases appear to be more complicated. The main result of this section is to show that this basic building block has conflict-free chromatic guard number $O(\log n)$. We then use this result to show that arbitrary simple polygons have conflict-free chromatic guard number $O(\log^2 n)$. A second (albeit minor) is that a naive recursive partitioning using x-aligned and y-aligned visibility may not even terminate in general polygons, and so we appropriately modify the partitioning to ensure finite termination. In the following, we assume without loss of generality that our polygon is x-monotone.

Monotone polygons

The monotone polygons are easily reduced to a collection of *independent* monotone polygons with a specialized structure, where one of the chains is either a line segment or a *concave chain*. Specifically, given an x-monotone polygon, consider the shortest path between the leftmost and the rightmost vertices. This path splits the polygon into a family of x-monotone pieces, with pieces of the shortest path forming one of their chains. In addition, all the subpolygons lying *below* the shortest paths are mutually independent, as are those lying above the path. Due to lack of space, the proof of the following lemma is omitted from this extended abstract.

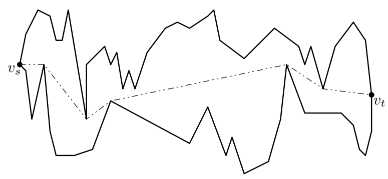

Figure 6 Partitioning a monotone polygon into independent monotone subpolygons over concave subchains.

▶ **Lemma 8.** *The conflict-free chromatic guard number for monotone polygons is at most twice the conflict-free chromatic guard number for monotone polygons over a concave chain.*

In the following, we show that monotone polygons over a concave chain have conflict-free chromatic guard number $O(\log n)$. The basic units of interest, however, turn out to be monotone polygons over an edge or over a *convex chain*. The following subsection handles their coloring, which in turn forms the basis for coloring of monotone polygons over concave chains.

Monotone polygons over a convex subchain

Let P_n be a monotone polygon over a single horizontal edge. Let $g(n)$ denote the smallest number of colors that ensure a conflict-free covering for any such P_n. Similar to our method for constructing staircases in orthogonal polygons, we consider decision lines through either horizontal edges with adjacent reflex vertices or through a vertex for which both of its neighbors have a higher y-coordinate. A decision line splits P_n in a lower part and two or more independent upper parts, of which at most one part contains more than $\lfloor n/2 \rfloor$ vertices. Then we construct a subpolygon that contains the base edge of P_n and at each decision line follows the part with the most remaining vertices, see the left picture in Figure 7. This subpolygon is *star-shaped* and can thus be guarded with a single guard. The remaining regions are mutually independent, x-monotone over a horizontal edge and contain at most $\lfloor n/2 \rfloor$ of P_n's vertices. We get the recurrence $g(n) \leq g(n/2) + 1$, which yields $g(n) \leq \log n$.

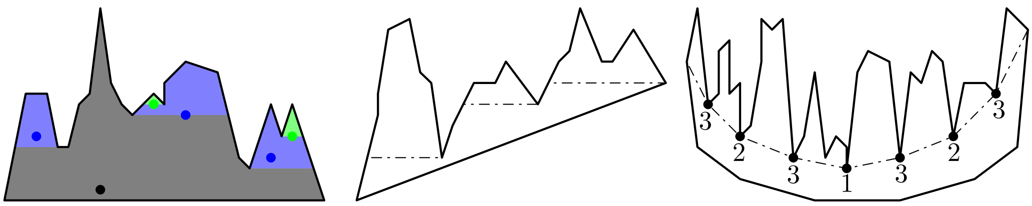

Figure 7 x-monotone polygons over a single horizontal edge, a sloped edge and a convex chain.

Now let's look at a monotone polygon over a single non-horizontal edge, without loss of generality ascending to the right. We show in the middle picture of Figure 7 that P_n can be partitioned into a set of independent monotone polygons over a horizontal edge and a tilted monotone polygon over a horizontal edge. Hence by Lemma 2 for any such polygon we have $\chi(P_n) \leq 2g(n) \leq 2 \log n$.

Monotone polygons P_n over a convex subchain are also easily covered with $O(\log n)$ colors. The shortest path in P_n from the leftmost vertex to the rightmost vertex cuts off independent monotone polygons over a single edge. The remaining subpolygon is bounded by a concave chain on top and the convex chain at the bottom. We can cover such a polygon using $\log n$ colors, by the following recursive process: place a guard of color $i = 1$ at the middle vertex of the concave chain; increment the color to $i = 2$, place guards of color 2 at the middle vertex of the two subchains, and so on. Clearly this requires $\log n$ colors, so it remains to show that the polygon is covered and the coloring is conflict-free. Let p be a point in the remaining subpolygon and let $l(p)$ be the list of all guard colors p can see. Between any two guards on the concave subchain that have the same color there must lie a guard of lower color between them. Hence the minimal color in $l(p)$ is a unique color among all guards that contain p in their visibility region. Therefore by Lemma 2, for any monotone polygon P_n over a convex chain we have $\chi(P_n) \leq 2g(n) + \log n \leq 3 \log n$.

Monotone polygons over a concave subchain

For monotone polygons P_n over a concave chain, we cut off independent monotone subpolygons over a horizontal edge as we did before in the case of a non-horizontal base edge. This results in two additional independent subpolygons whose boundary consists of a lower subchain which is concave and strictly increasing (respectively strictly decreasing) and an upper subchain which is monotonically increasing (respectively monotonically decreasing). In both of these subpolygons we place colored guards on the concave subchains as we did in the case of monotone subpolygons over a convex subchain. We show the partition and the guard placement and coloring in Figure 8. Let P be the subpolygon over the strictly increasing concave subchain. If a point p in P is guarded by a guard on the concave subchain, it has a guard of unique color among all other guards on the concave subchain that see p. However, there may be regions in P not guarded by the guards on the concave subchain. For these regions we have the following technical lemma, whose proof is omitted due to lack of space.

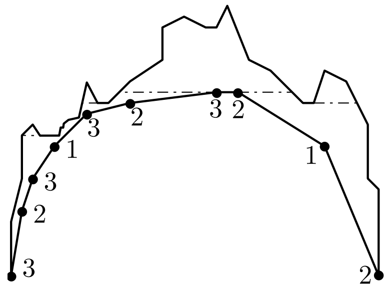

Figure 8 Guard placement for monotone polygon over a concave subchain.

▶ **Lemma 9.** *If a point $p \in P$ is not visible from any of the guards on the concave subchain, then p lies in a not yet guarded simply connected region, which has the shape of a monotone subpolygon over a convex chain. Furthermore all such regions are independent.*

Thus, we have a partition into monotone polygons over a single horizontal edge (where we need at most $\log n$ colors), monotone polygons over a convex chain (at most $3 \log n$ colors) and the two independent subpolygons guarded by the guards on the concave chain (at most $\log n$ colors). By Lemma 2 we have that for any monotone polygon P_n over a concave chain, $\chi(P_n) \leq 5 \log n$. In view of Lemma 8, we now have the main result of this section.

▶ **Theorem 10.** *The conflict-free chromatic guard number for monotone polygons on n vertices is $k(n) = O(\log n)$.*

5 Arbitrary Simple Polygons

Our proof structure for orthogonal polygons has the following form. We first partitioned the polygon into four different types of subpolygons and showed that the process terminates after a finite number of steps (Lemma 3). We then derived a necessary condition for two subpolygons of the same type not to be independent (Lemma 4). We then found a conflict-free covering using three colors for the basic building blocks, the staircase polygons (Lemma 5). We used this to get an upper bound of $4 \log n$ for polygons corresponding to red-black-alternating subtrees (Lemma 6). Finally we put all subtrees together to achieve an $O(\log n)$ upper bound on the chromatic guard number $k(n)$ for orthogonal polygons.

Our proof for non-orthogonal simple polygons follows the same outline, with appropriate differences spelled out. Specifically, given a n-vertex polygon P_n, we construct a partition $\{A_U, A_D, A_L, A_R\}$, where A_U, A_D, A_L, A_R, respectively, is the collection of up-polygons (depicted in red), down-polygons (in green), left-polygons (in black) and right-polygons (in blue). We rotate P_n in such a way that we can start with a horizontal line segment which gives rise to a first subpolygon of type U. Since the polygon's edges are no longer axis parallel, the partitioning process can be trapped between to edges e and f that ascend to the same direction. This gives rise to a long and possibly infinite alternating path, see the left picture in Figure 9.

In order to deal with this difficulty, we do the following. When an edge e of P gets touched during the partitioning process for the *second time* by a subpolygon of the same type, *without any vertex being touched in the meantime*, we extend this subpolygon until it touches a vertex of e or f, see the right picture in Figure 9. This ensures that in at least

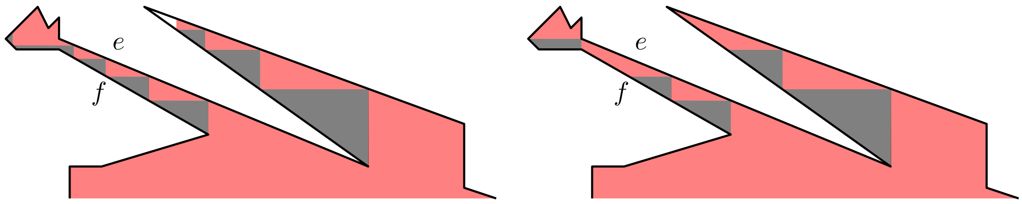

Figure 9 Replacing a red-black-alternating path with an augmented U subpolygon.

every third step of the partitioning process a vertex gets touched. Now any vertex v can be touched at most five times, since after the first time v gets touched, in each following step at least an additional $90°$ of v's interior angle are covered by a subpolygon of the partition. Along the lines of Lemma 3, this modification allows us to prove the following result, whose proof is omitted from this extended abstract due to lack of space.

▶ **Lemma 11.** *The revised partitioning process gives a complete partition after a finite number of (at most* $15n$*) steps.*

This replacement of alternating paths with a single polygon slightly changes the definition of the subpolygon types in the partitioning, but its does not change the relations between subpolygons of the same type when it comes to visibility—we simply replaced an alternating path with a *shorter* alternating path. This means that the schematic tree of the revised partitioning process has the same properties as the original partitioning process in orthogonal polygons, in particular we get as a corollary from Lemma 4:

▶ **Lemma 12.** *Let* P *be a polygon with the given revised partition and the schematic tree* T*. Let* P_i *and* P_j *be two arbitrary subpolygons of type U. Then, either (i)* P_i *and* P_j *are independent, or (ii) there exists a red-black-alternating (or a red-blue-alternating) directed path in* T *between* P_i *and* P_j*.*

This allows us to invoke the same coloring strategy as used in orthogonal polygons. We first focus on polygon regions corresponding to red-black-alternating trees. A polygon P_n corresponding to a red-black-alternating tree consists of type U and type L subpolygons; it "grows to the left". In place of Lemma 5, which states a constant conflict-free chromatic guard number for staircase polygons, we have Theorem 10, which gives an $O(\log n)$ for monotone polygons. We cover a part of P_n with a polygon that is both x- and y-monotone: Starting from the lowest and rightmost vertex of P_n, at every decision line we follow the upper part with the most vertices. We need $O(\log n)$ colors to do this plus an additional color to cover the convex fans to its left as before. We are left with independent subtrees, all of size $\leq \lfloor n/2 \rfloor$. We recursive each of them and cover all type U subpolygons of P_n in at most $\log n$ rounds. This leads to the following result.

▶ **Lemma 13.** *Suppose* P_n *is a simple polygon with a partition that has a red-black-alternating schematic tree. Then a conflict-free coloring of all the red regions of* P_n *needs at most* $O(\log^2 n)$ *colors. The same bound also holds for a red-blue-alternating schematic tree.*

The composition of red-black-alternating trees and red-blue-alternating trees that we described earlier depended only on the condition of Lemma 4, which we preserved in the revised partition of arbitrary polygons, see Lemma 12. Considering this, we can put subtrees together as we did in the case of orthogonal polygons. Thus we finally get an upper bound for simple polygons.

▶ **Theorem 14.** *The conflict-free chromatic guard number for simple non-orthogonal polygons on n vertices is $k(n) = O(\log^2 n)$.*

6 Conclusions

The art gallery problems provide a conceptually clean and mathematically elegant framework to study many applied questions related to surveilling, monitoring and covering of a physical environment. In this paper, we studied a *chromatic* variant of the art gallery, where the primary concern is to minimize the number of distinct *colors* assigned to guards. Our two main results are that (i) every n-vertex simple polygon has a conflict-free chromatic art gallery coverage with $O(\log^2 n)$ colors, and (ii) if the polygon is orthogonal, then the number of colors is only $O(\log n)$. A stronger form of coloring, which requires all guards visible to a point to be distinct in colors, needs $\Omega(n)$ colors for simple polygons and $\Omega(\sqrt{n})$ for orthogonal polygons [3], showing that the weaker conflict-free condition gives a significant improvement in the number of colors.

Our work suggests several directions for future research. Perhaps the most natural question is to investigate the lower bounds on the number of colors needed. Currently, we have none. What is the tight bound for the simple non-orthogonal polygons? Finally, the line-of-sight visibility model is a crude model for wireless communication. Recently, Fabila-Monroy et al. [5] have investigated the art gallery problems that allows the signal to penetrate k walls. One could consider our chromatic art gallery in a similar setting.

Acknowledgment We thank Luca Foschini for some insightful discussions during this research.

References

1 A. Bärtschi. Coloring Variations of the Art Gallery Problem. Master's thesis, Department of Mathematics, ETH Zürich, August 2011.

2 V. Chvátal. A combinatorial theorem in plane geometry. *Journal of Combinatorial Theory, Series B*, 18(1):39 – 41, 1975.

3 L. H. Erickson and S. M. LaValle. An art gallery approach to ensuring that landmarks are distinguishable. In *Proc.of Robotics: Science and Systems*, June 2011.

4 L. H. Erickson and S. M. LaValle. A chromatic art gallery problem. Technical report, Departement of Computer Science, University of Illinois at Urbana-Champaign, 2011.

5 R. Fabila-Monroy, A. Ruiz-Vargas, and J. Urrutia. On Modem Illumination Problems. In *Proc. XIII Encuentros de Geometria Computacional*, Zaragoza, Spain, June 2009.

6 J. O'Rourke. *Art Gallery Theorems and Algorithms.* Oxford University Press, 1987.

7 J. Pach and G. Tardos. Coloring axis-parallel rectangles. *Journal of Combinatorial Theory, Series A*, 117(6):776–782, 2010.

8 T.C. Shermer. Recent results in art galleries. *Proc. of the IEEE*, 1992.

9 S. Smorodinsky. *Combinatorial Problems in Computational Geometry.* PhD thesis, School of Computer Science, Tel-Aviv University.

10 S. Smorodinsky. On the chromatic number of some geometric hypergraphs. In *Proc. of 17th ACM-SIAM Symposium on Discrete Algorithms*, pages 316–323, 2006.

11 S. Smorodinsky. Conflict-Free Coloring and its Applications, 2010. `http://arxiv.org/abs/1005.3616v2`.