Robotics: Science and Systems 2011
Los Angeles, CA, USA, June 27-30, 2011

An art gallery approach to ensuring that landmarks are distinguishable

Lawrence H. Erickson and Steven M. LaValle
Department of Computer Science
University of Illinois at Urbana-Champaign
{lericks4, lavalle}@uiuc.edu

***Abstract*—How many different classes of partially distinguishable landmarks are needed to ensure that a robot can always see a landmark without simultaneously seeing two of the same class? To study this, we introduce the *chromatic art gallery problem*. A guard set $S \subset P$ is a set of points in a polygon P such that for all $p \in P$, there exists an $s \in S$ such that s and p are mutually visible. Suppose that two members of a finite guard set $S \subset P$ must be given different colors if their visible regions overlap. What is the minimum number of colors required to color any guard set (not necessarily a minimal guard set) of a polygon P? We call this number, $\chi_G(P)$, the *chromatic guard number* of P. We believe this problem has never been examined before, and it has potential applications to robotics, surveillance, sensor networks, and other areas. We show that for any spiral polygon P_{spi}, $\chi_G(P_{spi}) \leq 2$, and for any staircase polygon (strictly monotone orthogonal polygon) P_{sta}, $\chi_G(P_{sta}) \leq 3$. For lower bounds, we construct a polygon with $4k$ vertices that requires k colors. We also show that for any positive integer k, there exists a monotone polygon M_k with $3k^2$ vertices such that $\chi_G(M_k) \geq k$, and for any odd integer k, there exists an orthogonal polygon R_k with $4k^2 + 10k + 10$ vertices such that $\chi_G(R_k) \geq k$.**

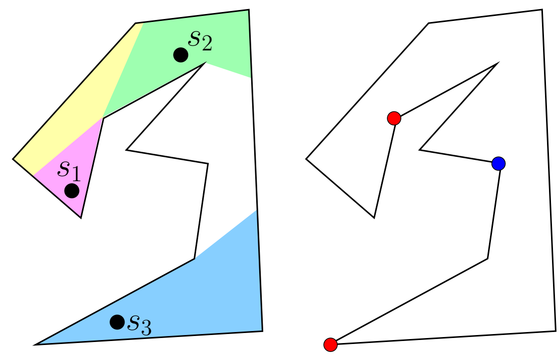

Fig. 1. [left] Three guards in a polygon P. The yellow region denotes the points that are visible from both s_1 and s_2, so s_1 and s_2 must be given different colors. The purple region is the area visible only from s_1, and the green region is the area visible only from s_2. The blue region is visible from s_3. Since this region does not intersect with the points visible from s_1 or s_2, the guard s_3 may be colored the same as s_1 or s_2. [right] A guard placement and coloring that uses only two colors. This is the minimum number of colors required for this polygon.

I. Introduction

Suppose a robot is navigating a region populated with colored landmarks. The robot is equipped with the following primitives: drive toward the landmark, drive away from the landmark, and drive in circles around the landmark. If this robot were in an area where two landmarks with the same color are visible, then its motion primitives may become unpredictable. If it can see two different green landmarks, then what is it to do when told "drive toward the green landmark"? This raises a natural question: How many classes of partially distinguishable guards are required to guard a given area (see Figure 1)? Equivalently, how many classes of landmarks are required so that the robot can always see a landmark (so that it can always navigate), but never two landmarks of the same class (so that it does not get confused)? In this paper, we try to answer this question for bounded simply connected polygonal areas. We assume that a robot cannot see a given landmark if the polygon boundary is in the way.

There are many reasons why one would want to minimize the number of landmark classes. Adding more classes of landmarks means that a more sophisticated sensing system is required. An eight color camera is easier to construct than a 32-bit color camera. Even if a camera can see thousands or millions of colors, differences in light or shade could still make classification difficult. This was demonstrated in [12], in which more powerful cameras (in terms of number of colors) were found to be worse for human iris identification than weaker cameras, as the more powerful cameras would see false differences in different pictures of the same eye. Minimizing the number of landmark classes could also make it so that only the "most different" classes are used, increasing the separation between sensor data points and decreasing classification errors. The problem of discovering the most distinctive visual landmarks for mobile robot navigation was addressed in [11]. Other research in which landmarks are specifically selected to reduce classification error include [20] and [24].

This is closely related to the original art gallery problem. It is impossible to list all of the significant results about art galleries in general polygons, but some of the most important works include results on tight bounds [2], [4] and exterior visibility [5]. Orthogonal art galleries are one of the most commonly studied variants, with notable results including tight bounds on the number of required guards [8], [15], [21] and bounds on the number of guards required for exterior visibility problems [7]. Results specific to monotone polygons include bounds on edge guards [1] and approximation algorithms with bounds independent of the number of polygon vertices [18]. Most of the important results from before 1987 are discussed in [22]. We prove lower bounds on the chromatic art gallery number for general, monotone, and orthogonal polygons.

We also prove upper bounds on the chromatic art gallery number for spiral polygons and staircase polygons (also known as strictly monotone orthogonal polygons). Spiral polygons are a heavily studied area in visibility. Special results for this class of polygons are available for the watchman route problem [19], the weakly cooperative guard problem [14], the visibility graph recognition problem [3], point visibility isomorphisms [16], and triangulation [25]. However, we are most interested in spiral polygons because of their use as building blocks. An algorithm for decomposing general polygons into a minimum number of spiral polygons was described in [10]. We choose to focus on spiral polygons because we think they could be a useful component in solving the chromatic guard number problem for general polygons, and staircase polygons for their similar potential as pieces of orthogonal polygons.

Section II contains the formal definition of the problem. Section III contains proofs for lower bounds on the chromatic guard number for general polygons, monotone polygons, and orthogonal polygons. Section IV contains upper bounds on the chromatic guard number for spiral polygons and staircase polygons. Section V discusses directions of future research.

II. Problem definition

Let a *polygon* P be a closed, simply connected, polygonal subset of $\mathbb{R}^2$ with boundary ∂P. A point $p \in P$ is *visible* from point $q \in P$ if the closed segment $\overline{pq}$ is a subset of P. The *visibility polygon* $V(p)$ of a point $p \in P$ is defined as $V(p) = \{q \in P \mid q \text{ is visible from } p\}$. Let a *guard set* S be a finite set of points in P such that $\bigcup_{s \in S} V(s) = P$. The members of a guard set are referred to as *guards*. A pair of guards $s, t \in S$ is called *conflicting* if $V(s) \cap V(t) \neq \emptyset$. Let $C(S)$ be the minimum number of colors required to color a guard set S such that no two conflicting guards are assigned the same color. Let $T(P)$ be the set of all guard sets of P. Let $\chi_G(P) = \min_{S \in T(P)} C(S)$. We call this value $\chi_G(P)$ the *chromatic guard number* of the polygon P. Note that the number of guards used can be as high or low as is convenient. We want to minimize the number of colors used, not the number of guards.

The notion of conflict can be phrased in terms of *link distance*. The link distance between two points $p, q \in P$ (denoted $LD(p,q)$) is the minimum number of line segments required to connect p and q via a polygonal path. Each line segment must be a subset of P.

Theorem 1. *Two guards $s_1, s_2 \in P$ conflict if and only if $LD(s_1, s_2) \leq 2$.*

Proof: If $LD(s_1, s_2) = 1$, then s_1 and s_2 are mutually visible, and obviously conflict.

If $LD(s_1, s_2) = 2$, then there exists a point $r \in P$, such that $\overline{s_1 r}, \overline{r s_2} \subseteq P$. Since $\overline{s_1 r} \subseteq P$, $r \in V(s_1)$. Since $\overline{r s_2} \subseteq P$, $r \in V(s_2)$. Because r is in $V(s_1)$ and $V(s_2)$, the intersection of $V(s_1)$ and $V(s_2)$ is non-empty; therefore s_1 and s_2 conflict.

If s_1 and s_2 conflict, then let r be a point in the intersection of $V(s_1)$ and $V(s_2)$. Since $r \in V(s_1)$, $\overline{s_1 r} \subseteq P$. Since $r \in V(s_2)$, $\overline{r s_2} \subseteq P$. Because $\overline{s_1 r}, \overline{r s_2} \subseteq P$, $LD(s_1, s_2) \leq 2$. ■

III. Lower bounds on the chromatic guard number

A finite set of lines in the plane is a *simple arrangement* if each pair of lines intersects and no three lines intersect at the same point. A simple arrangement of lines can be used to construct a polygon that requires a linear number of colors relative to the number of vertices in the polygon.

Theorem 2. *For every integer $k \geq 3$, there exists a polygon P_k with $4k$ vertices such that $\chi_G(P_k) \geq k$.*

Proof: The polygon P_k will be constructed from k gadgets, each consisting of four line segments. Each gadget consists of a nearly triangular well and a line that connects to the next gadget. The goal is to arrange k of these gadgets so that every pair of guards conflict, and each guard can guard no more than two convex vertices.

Let T be a simple arrangement of k lines. Now, make a closed convex k-gon bounding box B that contains each intersection among the lines of T in its interior, and has a boundary vertex on each line of T. Place the well of a very thin gadget at each of the boundary vertices (see Figure 2). Let p_1 and p_2 be two convex vertices in the same well associated with line T_1. Note that, as the opening of the well is made smaller, and the width of the segment joining p_1 and p_2 is made narrower, the distance between a point $q \in V(p_1) \cap B$ and the closest point to q in $T_1 \cap B$ becomes arbitrarily small. Note also that any guard placed in the well must lie on a line segment $\ell \subset V(p_1) \cup V(p_2)$ that extends from $\overline{p_1 p_2}$ to a polygon edge connecting two reflex vertices on the other side of P_k.

Since each guard in a well has an ℓ segment that is arbitrarily close to its line from the arrangement, and all the lines in the arrangement intersect, the ℓ segments from two guards in different wells must intersect (assuming that the wells are thin and the well openings are narrow enough), so two guards in different wells must conflict. A guard s located in B must conflict with every guard, as every ℓ segment intersects B, and $B \subset V(s)$. Therefore, all guards placed in P_k will pairwise conflict. Since P_k has $2k$ convex vertices, and each guard can see at most two convex vertices, k guards are required; hence $\chi_G(P_k) \geq k$. Since P_k is made from k gadgets, each of which has four edges, P_k has $4k$ vertices. ■

A polygon P is *monotone* if there exists a line L such that the intersection of P and any line perpendicular to L has at most one connected component. A polygon P is *strictly monotone* if there exists a line L such that any line perpendicular to L intersects ∂P at two or fewer points.

Theorem 3. *For every integer $k \geq 3$, there exists a strictly monotone polygon M_k with $3k^2$ vertices such that $\chi_G(M_k) \geq k$.*

Proof: The polygon M_k is a variant of the standard "comb" used to show the occasional necessity of $\lfloor n/3 \rfloor$ guards in the standard art gallery problem [2]. The vertex list of M_k is $[(1, 2k-2), (2, 2k-3), (4, 2k-3), (5, 2k-2), (6, 2k-$

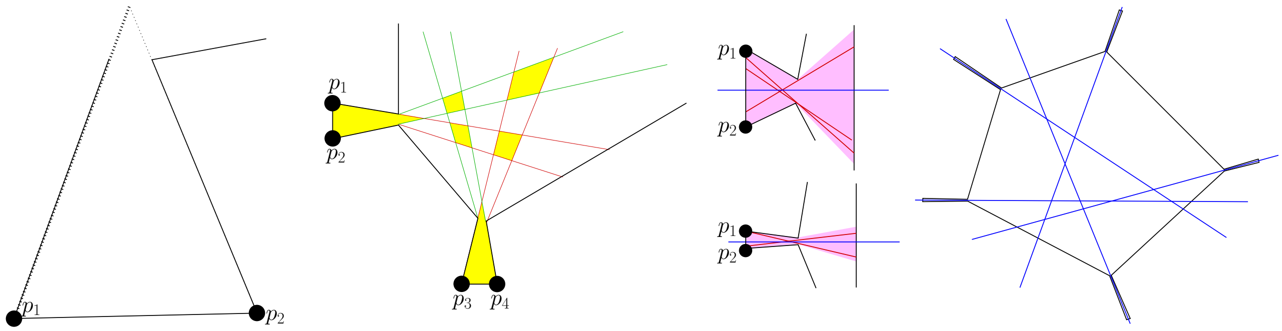

Fig. 2. [left] A gadget. The points p_1 and p_2 are the convex vertices. For a guard to see p_1 and p_2 simultaneously, it would have to be placed in the triangular region (bounded on top by the dotted lines) that does not extend far out of the well. [middle left] Two gadgets. The cones show the region outside of the well where a convex vertex is visible. The yellow regions are where a single guard can see two convex vertices. There is no place where a guard can see three convex vertices. [middle right] As the well opening is made smaller and the well is made more narrow, $V(p_1) \cup V(p_2)$ (purple region) becomes more narrow and any ℓ line segments (colored in red) from a guard in the well must get closer to arrangement line T_1 (colored in blue). [right] A polygon P_k for $k = 5$. The blue lines represent a simple arrangement T of $k = 5$ lines. Each line in the arrangement is associated with the well of a gadget.

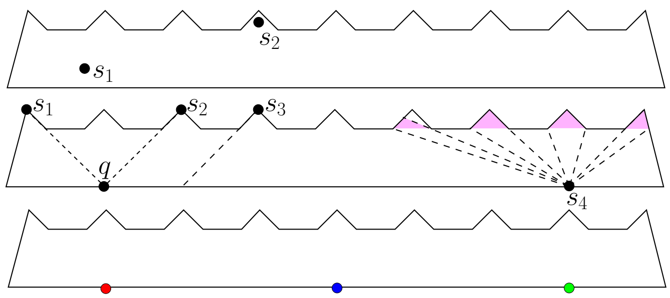

Fig. 3. [top] The polygon M_k for $k = 3$. The guard s_1 is a body guard, and the guard s_2 is an notch guard. [middle] Notch guards must be placed at least $k = 3$ notches away from each other to avoid conflicts. Guards s_1 and s_2 conflict as $V(s_1) \cap V(s_2) = q$, but s_3, which is $k = 3$ notches away from s_1 does not conflict with s_1. A body guard s_4 can only guard $k = 3$ notches by itself. Portions of the rightmost four notches visible from s_4 are highlighted in purple. [bottom] A guard placement that requires three colors.

$3) \ldots (4k^2-4, 2k-3), (4k^2-3, 2k-2), (4k^2-2, 0), (0, 0)]$. This polygon has $3k^2$ vertices, and it consists of a trapezoidal region (the *body region*) that has k^2 *notches* attached to the shorter edge. Call the vertices with a y coordinate of $2k-2$ *apex points*. Note that each notch has a unique apex point. A guard with coordinates (x, y) will be referred to as a *notch guard* if $y > 2k-3$ and will be referred to as a *body guard* if $y \leq 2k-3$ (see Figure 3).

Each body guard can guard up to k distinct notches. However, since the visibility polygon of a body guard includes the entire body region, and every guard's visibility polygon intersects the body region, a body guard will conflict with every other guard in the polygon. Let m_{body} be the number of body guards used in a guard set of M_k.

Each notch guard can guard only one notch. However, two notch guards will not conflict if they are placed far enough away from each other. Since the bottom edge of M_k has a y coordinate of 0, two notch guards are forced to conflict only if the distance between the apex points of their corresponding notches is $4k-4$ or less. Let a set of k notches be *consecutive* if the maximum distance between the apex points of any two notches in the set is $4k-4$. Let m_{notch} be the maximum number of notch guards in any consecutive set of k notches in M_k.

Suppose the polygon M_k has a guard set S assigned to it that requires only $\chi_G(M_k)$ colors. Consider k consecutive notches in M_k that contain m_{notch} notch guards in total. All of these notch guards will conflict with each other, and all of these notch guards will conflict with all of the body guards. Therefore, $\chi_G(M_k) \geq m_{notch} + m_{body}$. Now, note that each body guard can guard at most k notches. Since there are k^2 notches, by the pigeonhole principle, notch guards can guard at most km_{notch} notches (see Figure 3). Since each notch must be guarded, $km_{notch} + km_{body} \geq k^2$, so $m_{notch} + m_{body} \geq k$. Therefore $\chi_G(M_k) \geq m_{notch} + m_{body} \geq k$.

■

A polygon P is *orthogonal* (sometimes called *rectilinear* in other publications) if all of its angles are right angles.

Theorem 4. *For every odd integer $k \geq 3$, there exists a monotone orthogonal polygon R_k with $4k^2+10k+10$ vertices such that $\chi_G(R_k) \geq k$.*

Proof: We begin by introducing a family of orthogonal polygons with two parameters, $m, i \in \mathbb{Z}^+$. The vertex list for polygon $R_{m,i}$ is $[(0,0), (0, i+1), (1, i+1), (1, i), (2, i), (2, i+1), (3, i+1), (3, i), \ldots, (2m-2, i), (2m-2, i+1), (2m-1, i+1), (2m-1, 0)]$. This takes the form of a $(2m-1) \times i$ rectangle with m 1×1-sized notches along the top edge (see Figure 4). Any guard in $R_{m,i}$ with a y-coordinate greater than i will be called a *notch guard*. All other guards will be called *body guards*.

There are m notches. Each notch has a ceiling of length 1. These ceilings are a subset of the polygon, so they must be covered. A body guard can cover the most ceiling if it is placed on the bottom of the polygon. Let $C(s)$ be the total length of ceiling that a body guard s can see. Suppose a body guard s is placed on the bottom of the polygon underneath the left edge of a notch (thus maximizing the amount of ceiling it can see to its right). This guard can see all of the notch that it is underneath. It can see a length of $(i-2)/i$ of the next notch to the right, $(i-4)/i$ of the notch after that, and so on (see Figure 4). Therefore, s can see $\sum_{j=0}^{i/2} 2j/i$ ceiling to its right. We double this term to account for the ceiling it might be able to see on its left to get

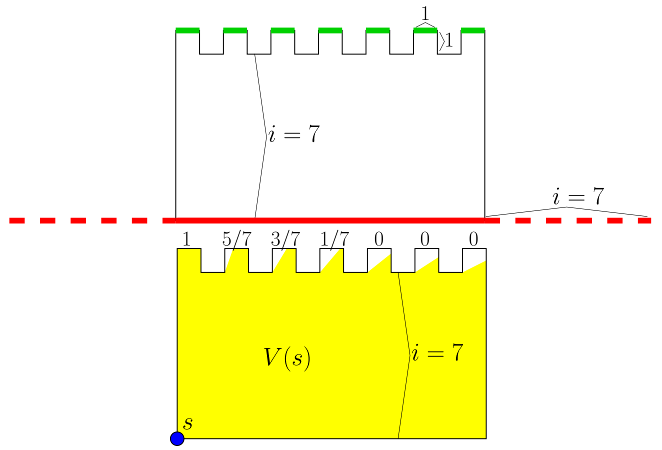

Fig. 4. [top] The polygon $R_{m,i}$ for $i = 7$ and $m = 7$. Each notch has a height and width of 1. The bottom edge is highlighted in red, and the ceiling edges are highlighted in green. The dotted red line represents the extra length i that we can assume exists on either side of the bottom edge for the purposes of placing nonconflicting notch guards. [bottom] A guard s is placed on the bottom edge of the polygon is a position where the total length of ceiling edge in $V(s)$ to the right of s is maximized. The visibility polygon $V(s)$ is highlighted in yellow. The number above each notch shows how much ceiling edge length in that notch is in $V(s)$.

$$C(s) \leq \sum_{j=0}^{i/2} \frac{4j}{i} = \frac{4}{i}\left(\frac{\left(\frac{i}{2}\right)\left(\frac{i}{2}+1\right)}{2}\right) = \frac{i}{2}+1. \quad (1)$$

Suppose that a certain color is used instead for notch guards. Each notch guard can guard a ceiling of length 1. However, while each body guard must have its own unique color, a single color can be assigned to multiple notch guards. So, given the dimensions of the polygon, how many notch guards can share one color? Note that the visibility polygon of a notch guard must include a portion of the bottom edge of the polygon. Since two notch guards that use the same color have visibility polygons that do not intersect, this space along the bottom edge of the polygon is a resource that can only support a finite number of notch guards of the same color. The bottom edge of the polygon has length $2m-1$. However, to account for the fact that the bottom of the visibility polygons of notches close to the edge could have an additional length of up to i if the convex portion of the polygon were wider, we can treat the bottom edge as though it has length $2m+2i-1$ (see Figure 5). It is clear that placing a notch guard s along the ceiling of a notch minimizes the length of the bottom edge inside $V(s)$. It is also clear that for any point p on the bottom edge of the polygon, there exists a point on the ceiling of a notch that is visible from p. Suppose a guard s is placed on the ceiling of a notch at a length $0 \leq t \leq 1$ from the left vertex of the ceiling. Since the height of the notch is 1, the leftmost point of $V(s)$ on the bottom edge will extend a distance ti past the x-coordinate of the leftmost point in the notch. Similarly, the rightmost point of $V(s)$ will extend a distance of $(1-t)i$ past the x-coordinate of the rightmost point in the notch (see Figure 5). Therefore, the length of the bottom edge inside $V(s)$ is $i+1$ (we have to include the length of 1 directly underneath

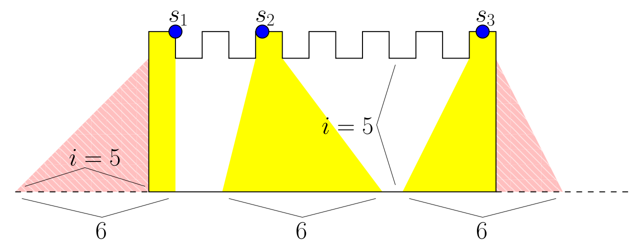

Fig. 5. The polygon $R_{m,i}$ for $i = 5$ and $m = 7$. Three guards have been placed on ceiling edges and their visibility polygons are highlighted in yellow. The striped pink regions are portions of the visibility polygons that have been cut off by the left or right side of $R_{m,i}$. Note that if the length of the bottom edge of $R_{m,i}$ extended an extra i in both directions, then the length of the bottom edge of each visibility polygon would be $i+1 = 6$, regardless of the guard's location on its notch's ceiling.

the notch). This means that the amount of the bottom edge seen by a single notch guard placed on a ceiling is not related to its exact location within that ceiling. Since no two notch guards with the same color can have any of their visibility polygons overlap, a single color can be used to guard at most $(2m+2i-1)/(i+1)$ notches.

Choose any guard set for $R_{m,i}$. Let x_{notch} be the number of colors used in the notch guards, and let x_{body} be the number of colors used in the body guards. Since each guard must be a notch or a body guard, we get

$$x_{notch} + x_{body} = \chi_G(R_{m,i}). \quad (2)$$

Since each color used for a body guard can guard at most $i/2+1$ length of ceiling, and each color used for notch guards can guard at most $(2i+2m-1)/(i+1)$ length of ceiling, and there is m total length of ceiling, we get

$$\left(\frac{2i+2m-1}{i+1}\right) x_{notch} + \left(\frac{i}{2}+1\right) x_{body} \geq m. \quad (3)$$

Let $k = (i-3)/2$ and let polygon R_k be the polygon where $m = (i^2-i)/4+1$ with $i \in \{x \in \mathbb{Z}^+ | x \equiv 1 \mod 4\}$. By the quadratic formula (and keeping in mind that i must be positive), this implies that $i = 1/2 + \sqrt{4m-(15/4)}$. This turns Equation 3 into

$$\left(\frac{i}{2}+1\right)(x_{notch}+x_{body}) \geq \frac{i^2-i}{4}+1. \quad (4)$$

The term $(i^2-i)/4+1$ is equal to $((i^2+2i)-(3i+6)+10)/4$; hence Equation 4 can be rewritten as

$$\chi_G(R_k) = x_{notch} + x_{body} \geq \frac{i}{2} - \frac{3}{2} + \frac{10}{2i+4} \geq \frac{i-3}{2} = k. \quad (5)$$

The polygon therefore requires at least $(i-3)/2 = \sqrt{m-(15/16)} - (5/4)$ colors. The polygon R_k has $4m$ vertices and $\chi_G(R_k) \geq \sqrt{m-(15/16)} - (5/4)$. Since $k = (i-3)/2 = \sqrt{m-(15/16)} - (5/4)$, R_k has $4k^2+10k+10$ vertices and requires k colors. The integer k must be odd to ensure that the number of vertices is divisible by 4. ■

While these constructions do not work when the desired number of required colors is 1 or 2, it is trivially easy to construct such polygons, as $\chi_G(P) \geq 1$ for all polygons, and $\chi_G(P) \geq 2$ for all non-star-shaped polygons.

IV. Upper bounds on the chromatic guard number

One could just give every guard its own color. Any polygon P with n vertices can be guarded by $\lfloor n/3 \rfloor$ guards (the art gallery theorem [2]), so $\chi_G(P) \leq \lfloor n/3 \rfloor$. However, this bound is unsatisfying, because colors can often be reused. There exist polygons with an arbitrarily high number of vertices that require only two colors. We prove bounds better than $\lfloor n/3 \rfloor$ for two categories of polygons.

A. Spiral polygons

A *chain* is a series of points $[p_1, p_2, \ldots, p_n]$ along with line segments connecting consecutive points. A *subchain* is a chain that forms part of the boundary of a polygon. The points p_1 and p_n are called *endpoints*, and all other points are *internal vertices*. A *convex subchain* is a subchain where all the internal vertices have an internal angle of less than π radians. A *reflex subchain* is a subchain where all the internal vertices have an internal angle of greater than π radians. Note that convex and reflex subchains can trivially consist of a single line segment (if there are no internal vertices). A spiral polygon is a polygon with exactly one maximal reflex subchain (all reflex subchains of the spiral polygon must be contained within the maximal reflex subchain).

Theorem 5. *For any spiral polygon P, $\chi_G(P) \leq 2$.*

Proof: The spiral polygon consists of two subchains, a reflex subchain, and a convex subchain. Let v_s and v_t be the endpoints of the reflex subchain. Without loss of generality, assume that the path along the convex subchain from v_s to v_t runs clockwise. The guards will all be placed along the edges of the convex subchain.

Call the nth guard placed s_n. Place s_1 at v_s. Let p_n be the point most clockwise along the convex subchain that is visible from s_n. Let b_n be the most counterclockwise vertex along the reflex subchain visible from s_n. Let g_n be the vertex immediately clockwise from b_n. Let r_n be the point on the convex subchain colinear with g_n and b_n and visible from both. Note that p_n and r_n define the endpoints of an interval along the convex subchain. Place s_{n+1} at a point on this interval that is not one of the endpoints. Note that this means $s_{n+1} \notin V(s_n)$. Terminate when a guard can see v_t (see Figure 6).

We can show that this is a guard set for the polygon by triangulating the polygon using the polygon vertices, the members of S, and the points p_i and showing that each triangle has a member of S as one of its vertices. Suppose that the polygon bounded by the edges starting from p_n counterclockwise along the boundary of P until b_n and the edge between p_n and b_n has already been triangulated such that each triangle contains a vertex in the set $\{s_i | i \leq n\}$. We must show that s_{n+1} can guard the subpolygon bordered by the edges counterclockwise from p_{n+1} to p_n, the edge between p_n and b_n, the vertices counterclockwise from b_n to b_{n+1}, and the edge between b_{n+1} and p_{n+1} (call this subpolygon P_{n+1}). If each of these vertices in the subpolygon is visible from s_{n+1}, then the subpolygon can be triangulated by connecting each vertex to s_{n+1}, meaning that s_{n+1} guards the entire subpolygon (see Figure 6).

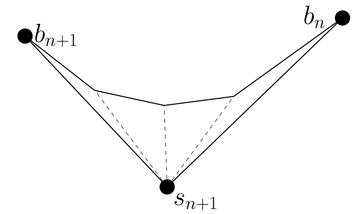

Fig. 7. A polygon consisting of the edges on the reflex subchain between b_n and b_{n+1} and the edges $\overline{s_{n+1}b_n}$ and $\overline{s_{n+1}b_{n+1}}$. Since all the vertices on the reflex subchain are reflex, this polygon has only one triangulation, where all triangles have s_{n+1} as an endpoint.

Since s_{n+1} is placed on the interval in between p_n and r_n, it must be able to see the entire edge between g_n and b_n, meaning that b_n is visible from s_{n+1}. By definition, the vertex b_{n+1} is visible from s_{n+1}. Examine the polygon consisting of the edges along the reflex subchain between b_n and b_{n+1}, $\overline{s_{n+1}b_n}$, and $\overline{s_{n+1}b_{n+1}}$. Since all the vertices along the reflex subchain are reflex, they cannot have edges between each other in a triangulation, so in any triangulation, they must all be connected to s_{n+1} (see Figure 7). By definition, the point p_{n+1} is visible from s_{n+1}. The point p_n is visible to s_{n+1} because s_{n+1} is on the convex subchain interval between p_n and r_n. If two points on the convex subchain interval between p_n and r_n are not mutually visible, then there must be a reflex vertex between b_n and g_n on the reflex subchain, but by definition, there are no such vertices. Because the vertices in between p_n and p_{n+1} lie on a convex subchain, if s_{n+1} can see both p_n and p_{n+1}, then s_{n+1} can see all the vertices in between. This means that P_{n+1} can be triangulated with every triangle having s_{n+1} as an endpoint, so s_{n+1} guards P_{n+1} (the triangle with endpoints p_{n+1}, b_{n+1}, and s_{n+1} is degenerate, as those three points are colinear, but this is not a problem). This technique still works if s_{n+1} can see v_t (in this case, $p_{n+1} = b_{n+1} = v_t$). This implies inductively that S is a guard set for P.

Because all the guards are along the convex subchain, if two guards conflict, their visibility polygons must intersect somewhere along the convex subchain. Also, since $s_n \notin V(s_{n+1})$ and $s_n \notin V(s_{n-1})$, s_{n+1} cannot conflict with s_{n-1}, or there would be no room along the convex subchain to place s_n. Therefore, all evenly indexed guards can be colored red, and all oddly indexed guards can be colored blue, so $\chi_G(P) \leq 2$. ■

B. Staircase polygons

An *alternating subchain* is a subchain with at least one internal vertex, with the first and last internal vertices being convex, and with consecutive internal vertices alternating between convex and reflex. A *staircase polygon* is an orthogonal polygon consisting of two convex vertices, v_w and v_z, connected by two alternating subchains. For simplicity, we will

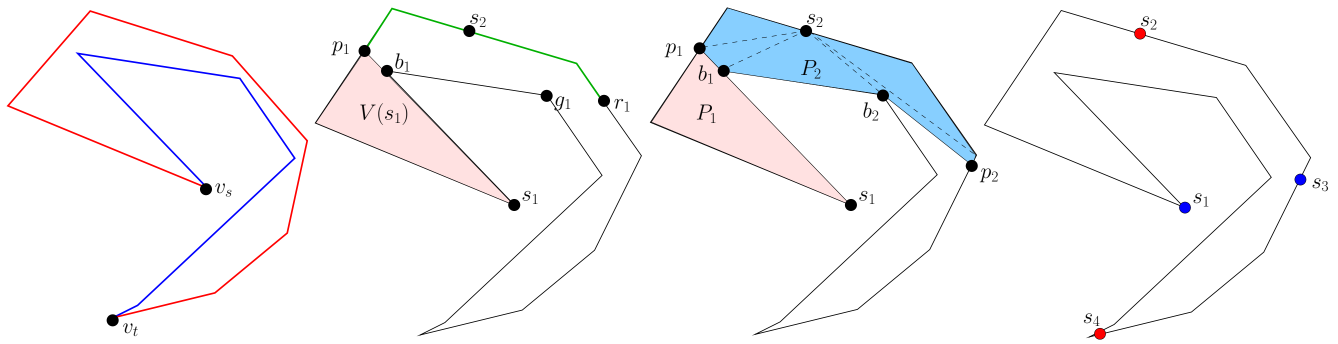

Fig. 6. [top left] A spiral polygon P. The convex subchain is highlighted in red, and the reflex subchain is highlighted in blue. [top right] The first guard s_1 is placed on vertex v_s. The points p_1, b_1, g_1, and r_1 are marked and the interval in which s_2 can be placed is highlighted in green. [bottom left] Recursively showing that placed guards form a guard set. The subpolygon P_1 is assumed to be guarded by s_1. The region that s_2 is responsible for is P_2, bounded by the reflex subchain between b_1 and b_2, the edge between p_2 and b_2, the convex subchain between p_2 and p_1, and the edge between b_1 and p_1. The subpolygon P_2 has been triangulated, indicating that s_2 can guard the whole subpolygon. The triangle with endpoints p_2, b_2, and s_2 is degenerate, as those three points are colinear. [bottom right] A guard placement and 2-coloring.

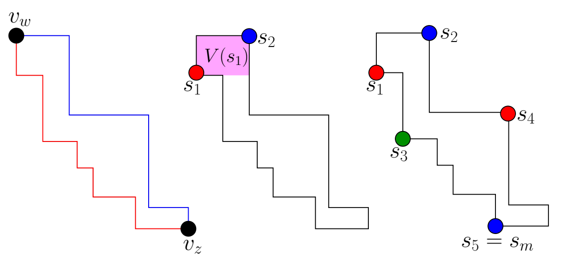

Fig. 8. [left] A staircase polygon P with vertices v_w and v_z identified. The lower subchain is highlighted in red, and the upper subchain is highlighted in blue. [middle] The guard s_1 is placed on the neighbor of v_w on the lower subchain. The guard s_2 is placed on the rightmost convex vertex in $V(s_1)$. [right] A guard placement and coloring for P that uses only three colors.

assume without loss of generality that orthogonal polygons are always oriented such that each edge is either vertical or horizontal, and that v_w is the top left vertex, and that v_z is the bottom right vertex. Put the polygon on a coordinate plane with v_w at the $(0, 0)$ coordinate, let right be the positive x direction, and let up be the positive y direction. The term “staircase polygon” is a synonym for strictly monotone orthogonal polygon (mentioned in [7], which solved the prison yard problem for this class of polygons). Note that the bound from Theorem 4 is for monotone orthogonal polygons, not *strictly* monotone orthogonal polygons.

Theorem 6. *For any staircase polygon P, $\chi_G(P) \leq 3$.*

Proof: Due to our assumptions about the orientation of the polygon P, one of the alternating subchains is going to be above the other one. Call the higher subchain the *upper subchain* and call the other subchain the *lower subchain*. Place a guard s_1 on the neighbor of v_w along the lower subchain. If guard s_i has been placed on the lower subchain, then place guard s_{i+1} on the right-most convex vertex on the upper subchain that is contained in $V(s_i)$. If guard s_i has been placed on the upper subchain, then place guard s_{i+1} on the right-most convex vertex on the lower subchain that is contained in $V(s_i)$. Stop placing guards when a guard can see v_z, and let m be the number of guards placed (see Figure 8).

First, it must be shown that s_i and s_{i+2} are not placed on the same vertex. Suppose without loss of generality that s_i is on the lower subchain. Note that the rightmost convex vertex on the lower subchain in $V(s_{i+1})$ must also be the lowest convex vertex on the lower subchain in $V(s_{i+1})$. Note also that a ray extended downward from s_{i+1} must intersect the horizontal edge incident to s_{i+2} (otherwise s_{i+2} would not be the rightmost convex vertex on the lower subchain). If this is the same horizontal edge that is incident to s_i, then the point where the ray intersects the horizontal edge incident to s_i must be a convex vertex (call it v_f). Since the convex vertex v_f neighbors the convex vertex v_i along a horizontal edge, and since v_f is to the right of v_i, v_f must be v_z. Therefore, s_{i+2} would only be placed on the same vertex as s_i if v_z is visible from s_{i+1}. Since we stop placing guards once a guard can see v_z, two guards will never be placed on the same vertex.

Next, it must be shown that this is a guard set for the staircase polygon. Suppose without loss of generality that guard s_i is placed on the lower subchain. Assume that the set $[s_1, s_2 \dots s_i]$ forms a guard set for the subpolygon that lies above the guard s_i (call this subpolygon P_i). We must show that the set $[s_1, s_2 \dots s_{i+1}]$ forms a guard set for the subpolygon that lies to the left of guard s_{i+1} (call this subpolygon P_{i+1}). Let p_{i+1} be the point where a ray extended downward from s_{i+1} intersects the lower subchain. Note that each vertex on the lower subchain between s_i and p_{i+1} is visible from s_{i+1}. We have to show that s_{i+1} guards $P_{i+1} \backslash P_i$. Let v_i^r be the reflex vertex to the right of s_i on the lower subchain. Let Q_{i+1} be the subpolygon below s_{i+1} and to the left of s_{i+1} (see Figure 9). Clearly, $Q_{i+1} \supseteq P_{i+1} \backslash P_i$ (as s_{i+1} cannot be lower than s_i). Note that every vertex of Q_{i+1} that is not connected to s_{i+1} by an edge of Q_{i+1} is on the lower subchain. For any given vertex v in Q_{i+1} that is not connected to s_{i+1} by an edge of Q_{i+1}, all edges of Q_{i+1} not incident to s_{i+1} that lie above v must also lie to the left of v, and all edges of Q_{i+1} not incident to s_{i+1} that lie to the right of v must also lie below v. Since s_{i+1} is never lower than v, and never to the right of v, every vertex v of Q_{i+1} must be visible from s_{i+1}. This means that one could triangulate Q_{i+1} such that each triangle has s_{i+1} as one of its corners. Therefore,

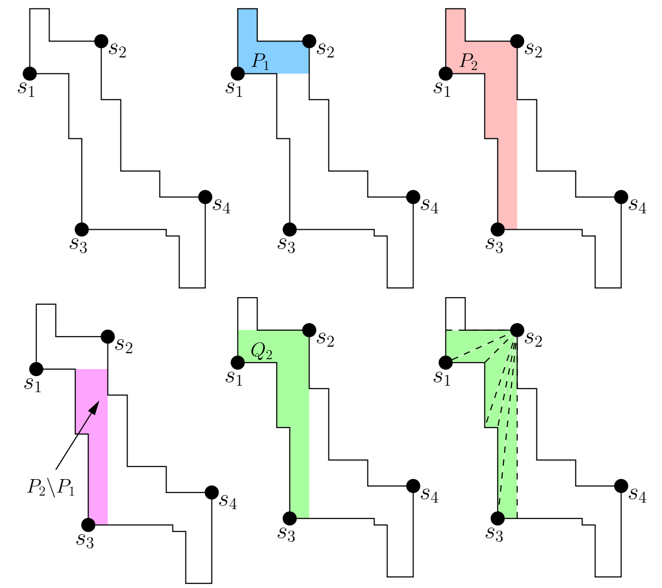

Fig. 9. [top left] A polygon P with a guard placement. [top middle] The region P_1 that s_1 is responsible for guarding. [top right] The region P_2 that s_1 and s_2 are reponsible for guarding. [bottom left] The region $P_2 \backslash P_1$ that s_2 is responsible for guarding. [bottom middle] The region Q_2, which consists of the portion of P below and to the left of s_2. This region is a superset of $P_2 \backslash P_1$. [bottom right] A triangulation of Q_2 where all triangles have a vertex at the location of s_2, showing that s_2 guards Q_2.

the guard s_{i+1} can guard Q_{i+1} by itself. Therefore, the set $[s_1, s_2 \ldots s_m]$ forms a guard set for P.

Finally, it must be shown that the guard set $[s_1, s_2 \ldots s_m]$ can be colored with three colors. Suppose guard s_i is placed on the lower chain. Let y_i be the y-coordinate of the lowest point visible from s_i. Note that, because s_i is on a convex right-angle vertex on the lower subchain, $V(s_i)$ is bordered on the bottom by a horizontal line at the same height as the horizontal edge incident to s_i; therefore y_i is just the y-coordinate of s_i. Let y_{i+3} be the y coordinate of the highest point in $V(s_{i+3})$. Because s_{i+3} is on a convex right-angle vertex on the upper subchain, $V(s_{i+3})$ is bordered on top by a horizontal line at the same height as the horizontal edge incident to s_{i+3}; therefore y_{i+3} is just the y-coordinate of s_{i+3}. Now, we must show that $y_i > y_{i+3}$. In the portion of the proof that showed that each guard is placed on a unique vertex, we demonstrated that the y-coordinate of s_{i+1} (call it y_{i+1}) has to be higher than the y-coordinate of s_{i+3}. If $y_i \leq y_{i+3}$, then $y_i \leq y_{i+3} < y_{i+1}$. However, this is impossible, because s_{i+1} was placed on the rightmost (and thus, lowest) vertex on the upper chain that was in $V(s_i)$. Therefore, $y_i > y_{i+3}$. Since the highest point in $V(s_{i+3})$ is lower than the lowest point in $V(s_i)$, s_i and s_{i+3} cannot conflict (see Figure 10).

Since s_i and s_{i+3} do not conflict, we can color all guards with an index of $0 \mod 3$ with green, all guards with an index of $1 \mod 3$ with red, and all guards with an index of $2 \mod 3$ with blue. Therefore $\chi_G(P) \leq 3$.

We have assumed throughout this proof that guard s_i was placed on the lower subchain. However, the arguments made above still apply if s_i was placed on the upper subchain (reflect the polygon over the $y = -x$ line).

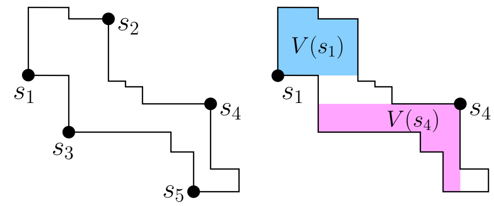

Fig. 10. [left] A staircase polygon P with a guard placement. [right] The regions $V(s_1)$ and $V(s_4)$ are shown. Note that the lowest point in $V(s_1)$ is higher than the highest point in $V(s_4)$, as the horizontal line incident to s_1's vertex is higher than the horizontal line incident to s_4's vertex.

■

V. Conclusion

We have introduced the chromatic art gallery problem, which asks for the minimum number of landmark colors required to ensure that a robot travelling in a given polygon can always see at least one landmark, but never simultaneously sees two of the same color. We have constructed a polygon with n vertices that requires $\Omega(n)$ colors, and we have constructed monotone and orthogonal polygons that require $\Omega(\sqrt{n})$ colors. We have also found constant upper bounds on the chromatic guard number for the spiral and staircase polygons. These two families of polygons may be useful as building blocks for polygons in more general families.

The results from Section III seem to indicate that the environments that have the highest chromatic guard number have a large central convex region with several smaller niches attached to it. Therefore, if one were designing an environment where robots were to navigate via visual landmarks, it may be advantageous to design the environment without such a region, as that region would require more landmark classes and would potentially be more susceptible to classification errors.

Some directions of future research would be finding bounds for other families of polygons, and finding tight bounds for the general, monotone, and orthogonal polygons. Visibility in curvilinear bounded regions has also been researched [9]. Allowing polygons with holes is another possibility, as is placing further restrictions on the placement of guards, perhaps forcing the guards to be strongly cooperative [23] or weakly cooperative [17].

The problem could also be attacked from a visibility graph context. The structure of standard visibility graphs for general polygons is still not completely understood, but [6] gives four necessary conditions for visibility graphs. It is likely that analogues of these four conditions could be made for "2-link" visibility.

There are also algorithmic questions. While finding the minimum number of art gallery guards for a given polygon is NP-complete [13], it is not necessary to find the minimum number of art gallery guards to find the minimum number of colors required for a polygon (see Figure 11). There is also the possiblity that the graphs representing the conflict relationships between guards (each graph vertex is a guard, and two vertices are connected by an edge if the corresponding guards conflict) is an easy family of graphs to color. However, these graphs

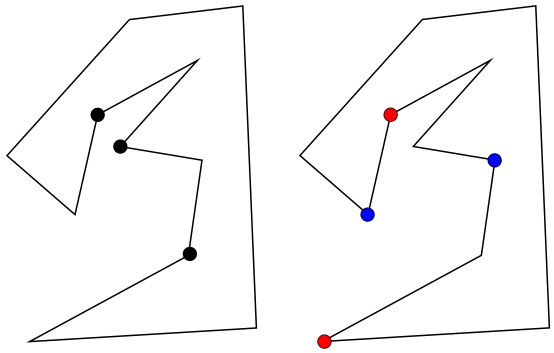

Fig. 11. [left] A polygon P with a minimal guard set of three guards. [right] A guard placement and coloring for P that uses the minimum two colors, but four guards. This demonstrates that a guard placement that uses the minimum number of colors does not need to use the minimum number of guards.

are not generally perfect graphs, and there are relatively few non-perfect families of graphs that are easy to color.

Finally, for practical robotics purposes, it would be useful to make a more realistic model of when guards conflict. For example, it may be interesting to research the case where the robot has limited vision, so that two guards sufficiently far from each other will not conflict even if there is no obstacle between them.

Acknowledgments

This work is supported in part by and NSF grant 0904501 (IIS Robotics), NSF grant 1035345 (CNS Cyberphysical Systems), DARPA SToMP grant HR0011-05-1-0008, and MURI/ONR grant N00014-09-1-1052.

References

[1] I. Bjorling-Sachs and D. L. Souvaine. A tight bound for edge guards in monotone polygons. In *Canadian Conference on Computational Geometry*, 1992.

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

[3] H. Everett and D. G. Corneil. Recognizing visibility graphs of spiral polygons. *Journal of Algorithms*, 11(1): 1–26, 1990.

[4] S. Fisk. A short proof of Chvatal's watchman theorem. *Journal of Combinatorial Theory Series B*, 24(3):374, 1978.

[5] Z. Furedi and D. J. Kleitman. The prison yard problem. *Combinatorica*, 14(3):287–300, 1994.

[6] S. K. Ghosh. On recognizing and characterizing visibility graphs of simple polygons. *Discrete and Computational Geometry*, 17(2):143–162, 1997.

[7] F. Hoffman and K. Kriegel. A graph-coloring result and its consequences for polygon-guarding problems. *SIAM Journal on Discrete Mathematics*, 9(2):210–224, 1996.

[8] J. Kahn, M. Klawe, and D. Kleitman. Traditional galleries require fewer watchmen. *SIAM Journal on Algebraic and Discrete Methods*, 4(2):194–206, 1983.

[9] M. I. Karavelas, C. D. Toth, and E. P. Tsigaridas. Guarding curvilinear art galleries with vertex or point guards. *Computational Geometry: Theory and Applications*, 42 (6–7):522–535, 2009.

[10] J. M. Keil. Decomposing a polygon into simpler components. *SIAM Journal on Computing*, 15(4):799–817, 1985.

[11] M. Knapek, R.S. Oropeza, and D. J. Kriegman. Selecting promising landmarks. In *Proc. IEEE International Conference on Robotics and Automation*, pages 3771–3777, 2000.

[12] Emine Krichen, Mohamed Chenafa, Sonia Garcia-Salicetti, and Bernadette Dorizzi. Color-based iris verification. In Seong-Whan Lee and Stan Li, editors, *Advances in Biometrics*, volume 4642 of *Lecture Notes in Computer Science*, pages 997–1005. Springer Berlin / Heidelberg, 2007.

[13] D. T. Lee and A. K. Lin. Computational complexity of art gallery problems. *IEEE Transactions on Information Theory*, 32(2):276–282, 1986.

[14] B. C. Liaw and R.C.T. Lee. An optimal algorithm to solve the minimum weakly cooperative guards problem for 1-spiral polygons. *Information Processing Letters*, 52 (2):69–75, 1994.

[15] A. Lubiw. Decomposing polygonal regions into convex quadrilaterals. In *Proc. Symposium on Computational Geometry*, pages 97–106, 1985.

[16] G. MacDonald and T. Shemer. Isomorphism of spiral polygons. *Discrete and Computational Geometry*, 16: 277–304, 1996.

[17] T. S. Michael and V. Pinciu. Art gallery theorems for guarded guards. *Computational Geometry: Theory and Applications*, 26(3):247–258, 2003.

[18] B. Nilsson. *Approximate guarding of monotone and rectilinear polygons*, pages 1362–1373. Lecture Notes in Computer Science. Springer Berlin/Heidelberg, 2005.

[19] B. J. Nilsson and D. Wood. Optimum watchmen routes in spiral polygons. In *Canadian Conference on Computational Geometry*, 1990.

[20] C. F. Olson. Selecting landmarks for localization in natural terrain. *Autonomous Robotics*, 12(1):201–210, 2002.

[21] J. O'Rourke. An alternate proof of the rectilinear art gallery theorem. *Journal of Geometry*, 21(1):118–130, 1983.

[22] J. O'Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, Cambridge, UK, 1987. ISBN 0-19-503965-3.

[23] V. Pinciu. *A coloring algorithm for finding connected guards in art galleries*, pages 257–263. Lecture Notes in Computer Science. Springer Berlin/Heidelberg, 2003.

[24] S. Thrun. Finding landmarks for mobile robot navigation. In *Proc. IEEE International Conference on Robotics and Automation*, pages 958–963, 1998.

[25] G. Toussaint. Efficient triangulation of simple polygons. *The Visual Computer*, 7(5–6):280–295, 1991.