Navigation among visually connected sets of partially distinguishable landmarks

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

***Abstract*— A robot navigates in a polygonal region populated by a set of partially distinguishable landmarks. The robot's motion primitives consist of actions of the form "drive toward a landmark of class x". To effectively navigate, the robot must always be able to see a landmark. Also, if the robot sees two landmarks of the same class, its motion primitives become ambiguous. Finally, if the robot wishes to navigate from landmark s_0 to landmark s_{goal} with a simple graph search algorithm, then there must be a sequence of landmarks $[s_0, s_1, s_2, \ldots, s_k = s_{goal}]$, in which landmark s_i is visible from s_{i-1}. Given these three conditions, how many landmark classes are required for navigation in a given polygon P? We call this minimum number of landmark classes the *connected landmark class number*, denoted $\chi_{CL}(P)$. We study this problem for the monotone polygons, an important family of polygons that are frequently generated as intermediate steps in other decomposition algorithms. We demonstrate that for all odd k, there exists a monotone polygon M_k with $\frac{3}{4}(k^2 + 2k + 1)$ vertices such that $\chi_{CL}(P) \geq k$. We also demonstrate that for any n-vertex monotone polygon P, $\chi_{CL}(P) \leq n/3 + 12$.**

I. Introduction

Suppose a point robot is searching for a treasure in a simply connected polygonal environment populated by a set of partially distinguishable landmarks that the robot can detect visually. The robot's motion primitives are "drive toward the red landmark", "drive toward the green landmark", "drive toward the blue landmark", etc. If the robot can see the treasure, then the robot "wins". What properties can a set of partially distinguishable landmarks have that would allow such a robot to find an arbitrarily placed treasure?

Suppose a landmark set satisfied the following three conditions:

1) There is no point where the robot can see two landmarks of the same class at the same time (the robot's motion primitives are never ambiguous).
2) For each point in the environment, there exists a landmark visible from that point (the robot is never without bearings on which to base its primitives).
3) For any two landmarks s_1 and s_{goal}, there is a sequence of landmarks $[s_1, s_2, \ldots, s_{k-1}, s_k = s_{goal}]$ such that s_i is visible from s_{i-1} for all i from 2 to k.

An environment containing such a landmark set could be easily searched for treasure. The treasure must be visible from one of the landmarks, and the robot could use a graph searching algorithm to travel to each landmark until the treasure is found. The robot would not even require a priori knowledge about the number of landmarks or their arrangement. Due to the first condition, the landmarks are always locally distinguishable to the robot. This provides the robot sufficient information to use the graph searching algorithm described in [6].

Problems about landmark sets that satisfy the second condition are typically called *art gallery problems*, and there is a great deal of prior work on the subject. Determining the minimum number of landmarks required is NP-complete [13]. Tight bounds on the number of landmarks required are demonstrated in [2] and [5]. An approximation algorithm with a $\log(n)$ approximation factor is described in [8], and an exact algorithm with empirically fast performance is described in [3]. Results are also available for rectilinear polygons [10], [14], [21], monotone polygons [19], and regions with curved boundaries [11].

The problem of determining the minimum number of landmark classes required for a landmark set that satisifies the first and second conditions above is called the *chromatic art gallery problem*. This problem is discussed in [4], where upper and lower bounds on the number of landmark classes required to satisfy these constraints are determined for several families of polygons.

The third condition means that the *visibility graph* of the guard set must be connected. Landmark sets that satisfy the second and third conditions are called *cooperative guard sets*. Tight bounds on the number of landmarks required for cooperative guard set were determined in [9], and with a different method in [22]. Tight bounds on a weaker version of this problem, where each landmark must see at least one other landmark, were described in [17].

What kinds of robots are capable of exploiting a landmark set that meets the three conditions? Light gradient ascent is one navigation method that could produce a robot capable only of the "drive toward a landmark of class x" primitives and nothing else (each landmark is a light source of a certain color). Navigation via light gradients was described in [18]. Since these are sufficient conditions for a robot with the aforementioned primitives, landmark sets with these three restrictions could be effectively used for navigation by any robotic system capable of performing these primitives, even if their actuation and sensor suites allow them to perform other primitives as well. Many robotic systems that navigate by visual landmarks fall into this category, including the ones used in [15], [23] and the RGBD-sensing robots of [1].

The problem of adding some set of artificial, indistinguishable landmarks to an environment for the purposes of robotic navigation was studied in [16]. Indistinguishable landmarks are used so that no landmark coding or identification system is needed. The results in this paper could be applied in a similar manner by adding some partially distinguishable landmarks to the environment for the purposes of simplifying navigation. Instead of totally eliminating a coding or identification system, our aim is to make such a system as simple as possible by minimizing the number of landmark classes.

There are two primary reasons why one would wish to minimize the number of landmark classes that the robot uses to navigate. First, if the robot is required to distinguish from among fewer landmark classes, then it may be possible to construct the robot with a less powerful (and thus, less costly, less complex, etc.) sensor suite. Second, the use of a small number of landmark classes may allow only the most "different" types of landmarks to be used for navigation. For example, a robot that navigates via red and green landmarks is probably more robust than one that navigates via red, green, and teal landmarks, as green and teal may be easily confused. This is closely related to the data association problem that frequently arises in landmark-based navigation, discussed in [12], [20], and [24].

The contributions of this paper are non-trivial lower and upper bounds on the number of landmark classes required for monotone polygons. Although robotic environments are typically not monotone, these bounds are still useful because there exist many algorithms for decomposing polygons into monotone pieces. We believe that if care is taken with the monotone decomposition, the results in this paper can be used to create a cooperative guard set for a general polygon that requires few landmark classes.

Section II formally describes the problem. Section III describes a pathological family of monotone polygons that require $\Omega(\sqrt{n})$ colors in an n-vertex polygon. Section IV describes a way of placing guards in an n-vertex monotone polygon that requires only $n/3 + 12$ colors. Section V discusses the implications of this research and directions of further research.

II. Problem Statement

Let P be a closed bounded simply connected polygonal subset of $\mathbb{R}^2$. Let ∂P be the boundary of P. Since P is a closed region, $\partial P \subset P$. For a point $s \in P$, the *visibility polygon* $Vis(s)$ is the set $\{p \in P \mid \overline{sp} \in P\}$. A finite set of points $S \subset P$ is a *guard set* of P if $\bigcup_{s \in S} Vis(s) = P$.

Let $G(S,P)$ be the *conflict graph* of a finite set of points $S \subset P$, in which the vertex set of $G(S,P)$ is S, and $s_i, s_j \in S$ share an edge if $Vis(s_i) \cap Vis(s_j) \neq \emptyset$. We will say that two members of S *conflict* if they share an edge in $G(S,P)$.

Let $VG(S,P)$ be the *visibility graph* of a finite set of points $S \subset P$, in which the vertex set of $VG(S,P)$ is S, and $s_i, s_j \in S$ share an edge if $s_i \in Vis(s_j)$. Note that this is a different structure than the conflict graph $G(S,P)$. In fact, $VG(S,P)$ is a subgraph of $G(S,P)$ (they have the same vertex set, but $G(S,P)$ generally has more edges). A guard set S is *cooperative* if $VG(S)$ is connected.

Let S_{all} be the set of all cooperative guard sets of P. Let $\chi(G)$ be the chromatic number of the graph G, the minimum number of colors required in a proper coloring. We are interested in the *connected landmark class number* of P, denoted as $\chi_{CL}(P)$, and defined to be $\min_{S \in S_{all}} \chi(G(S,P))$. Due to the relationship of the the connected landmark class number to cooperative guard sets and graph coloring, "landmark classes" will usually be referred to as "guard colors" from this point on.

A polygon P is *monotone* if there exists a line H such that the intersection of P and any line perpendicular to H is connected. Without loss of generality, we can assume that H is a horizontal line (by rotating the polygon). For a polygon P so rotated, P is monotone if there exist two distinct vertices of ∂P, v_{front} and v_{back}, such that v_{front} is the vertex of ∂P with the smallest x-coordinate, v_{back} is the vertex of ∂P with the largest x-coordinate, and the remainder of ∂P forms two paths from v_{front} to v_{back} that each have monotonically increasing x-coordinates. The *upper subchain* is the path with the higher y-coordinates, and the *lower subchain* is the path with the lower y-coordinates. The vertices v_{front} and v_{back} are considered to be part of both the upper and lower subchains.

III. Lower Bounds

From [4], it is known that there is a monotone polygon with $3k^2$ vertices that requires k colors in any guard set. This provides a trivial lower bound, because the same polygon must require at least k colors in any strongly connected guard set. The following theorem provides a better bound.

Theorem 1: For any odd $k \in \mathbb{N}$, there exists a monotone polygon M_k with $\frac{3}{4}(k^2 + 2k + 1)$ vertices such that $\chi_{CL}(M_k) \geq k$.

Proof: This lower bound is provided by a variant of the "comb" polygon (with $(k+1)^2/4$ notches) used to demonstrate the lower bound for the art gallery problem in [2] (see Figure 1). Each of the triangular notches is very thin; therefore it is impossible to use the same guard to guard two notches simultaneously. We define two types of guards. *Notch guards* are guards placed inside the triangular notches. *Body guards* are guards placed in the lower trapizoidal region. Let x_{notch} and x_{body} denote the number of colors assigned to each type of guard. Note that each body guard requires its own unique color, all body guards are mutually visible, and two notch guards in different notches require a path through $VG(S,P)$ for which the non-endpoint vertices are body guards.

Due to the thinness of the notches, a body guard can only guard a single notch tip. A single color assigned to notch guards could be used to guard each notch tip, but then each notch guard would require its own body guard (with a unique color) to ensure that $VG(S,P)$ is connected, which would require the number of colors to be linear in the number of

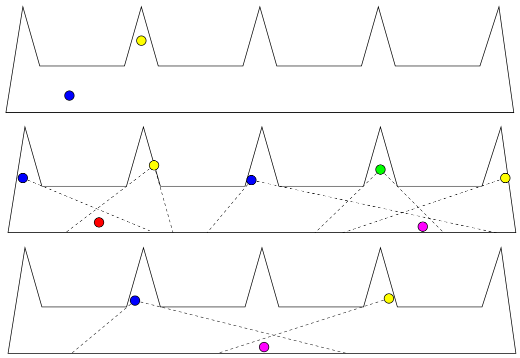

Fig. 1. [top] A "comb" polygon with five notches. The yellow guard is a notch guard, and the blue guard is a body guard. [middle] Notch guards require body guards to connect them to the rest of the visibility graph. No notch guards that share a connecting body guard can have the same color. [bottom] A body guard can guard one notch by itself.

notches. A set T of notch guards can share one "connecting" body guard, but only if $\bigcap_{t \in T} Vis(t) \neq \emptyset$, which means that all members of T must use different colors (this body guard could guard a notch by itself). Suppose the notches are grouped into sets of size m, guarded by $m-1$ notch guards and 1 body guard, in which the same $m-1$ colors are used in the notch guards of each set. There is no reason to make the sets different sizes (if one set uses y colors in its notch guards, and another uses $z < y$ colors, then the second set could include more notches without raising the total number of colors used). If the notches are grouped in this way, then $x_{notch} \geq m-1$ and $x_{body} \geq (k+1)^2/4m$. Since $\chi_{CL}(M_k) = x_{notch} + x_{body}$, we obtain

$$\chi_{CL}(M_k) \geq \min\{m - 1 + \frac{(k+1)^2}{4m} \mid m \in \mathbb{N}\}. \quad (1)$$

The value of m that minimizes the right side is $m = \frac{k+1}{2}$, which after substitution reduces it to $\chi_{CL}(M_k) \geq k$. Note that we do not claim that a grouping of notches in which every notch is a part of a size $\frac{k+1}{2}$ set that all share a common body guard actually exists, merely that any realizable grouping would use at least as many colors.

Since each notch requires 3 vertices and the polygon has $(k+1)^2/4$ notches, M_k has $\frac{3}{4}(k^2+2k+1)$ vertices. ■

IV. Upper Bounds

An upper bound on the number of landmarks required for an n-vertex monotone polygon is $\lfloor (n-2)/2 \rfloor$, provided by [9] and [22]. Coincidentally, this bound is actually tight even for monotone polygons, because the pathological family of polygons provided by Pinciu [22] is monotone (see Figure 2). Since one could simply place a set of landmarks according to Pinciu's algorithm and give them all different colors, we have a trivial bound of $\chi_{CL}(P) \leq \lfloor (n-2)/2 \rfloor$. However, one can generally do better than this by reusing colors. In this section, we demonstrate that $\chi_{CL}(P) \leq n/3 + 12$.

We begin by defining a decomposition of the monotone polygon P into a number of smaller *hemimonotone* polygons. A polygon P is hemimonotone if there exist two vertices of the boundary of P, v_{left} and v_{right} such that P can be aligned so that the boundary of P consists of

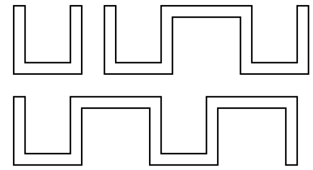

Fig. 2. Polygons from [22] that demonstrate that a monotone polygon may require $\lfloor n-2 \rfloor/2$ guards in a cooperative guard set.

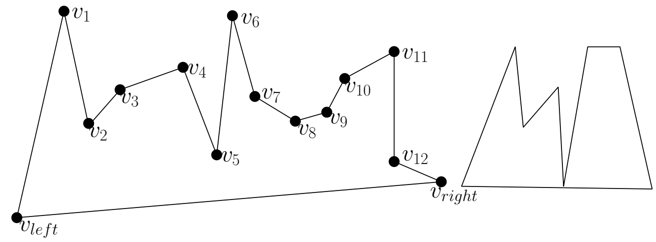

Fig. 3. [left] A hemimonotone polygon that is monotone relative to a horizontal line. [right] A degenerate hemimonotone polygon where the base is intersected by a vertex on the non-base path.

two paths whose x-coordinates are monotonically increasing and one of the paths is a straight line segment $\overline{v_{left}v_{right}}$ which will be referred to as the *base* (see Figure 3). Note that v_{left} and v_{right} cannot be reflex vertices. Let $L = [v_{left}, v_1, v_2, v_3, \ldots, v_{right}]$ be the sequence of vertices of ∂P, ordered as per their appearance on the path between v_{left} and v_{right} (the path that is not the base). We allow the hemimonotone polygons to be *degenerate*, meaning that it is possible for v_i to intersect $\overline{v_{left}v_{right}}$ (see Figure 3).

If a monotone polygon P is structured such that $v_{front} \in Vis(v_{back})$, then P can be easily decomposed into two hemimonotone polygons by using the line segment $\overline{v_{front}v_{back}}$. However, if those two points cannot see each other, a more complicated approach is required.

There are three types of line segments and pairs of lines segments that can be exploited to guarantee that a pair of hemimonotone polygons are separated enough to use identical color sets. Suppose there exists a line segment with one endpoint e_1 (the endpoint of the segment with the lower x-coordinate) on the lower subchain that first intersects a reflex vertex r_1 of the upper subchain, then intersects a reflex vertex r_2 of the lower subchain, and then has an endpoint e_2 on the upper subchain (or alternately, a line segment with the same behavior with the upper and lower subchains reversed). Call such a line segment a *valid bitangent*. A valid bitangent divides the polygon into four regions (see Figure 4).

- P_1 - A monotone polygon containing v_{front}.
- P_2 - A hemimonotone polygon with $\overline{e_1r_2}$ as a base.
- P_3 - A hemimonotone polygon with $\overline{r_1e_2}$ as a base.
- P_4 - A monotone polygon containing v_{back}.

Note that for a point $p_1 \in P_1 \setminus \overline{e_1r_1}$, $Vis(p_1) \cap (P_3 \cup P_4) = \emptyset$, because r_1 blocks p_1's view of P_3, and r_2 blocks p_1's view of P_4. Similarly, for a point $p_4 \in P_4 \setminus \overline{r_2e_2}$, $Vis(p_4) \cap (P_1 \cup P_2) = \emptyset$, because r_2 blocks p_4's view of P_2, and r_1 blocks p_4's view of P_1. Therefore, p_1 and p_4 do not conflict, which means that, other than the portion of $\overline{e1e2}$ that each polygon contains, guards placed in P_4 can reuse the colors of the guards placed in P_1.

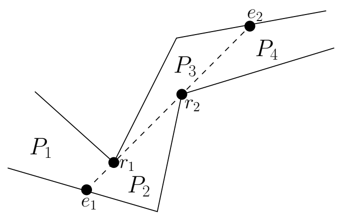

Fig. 4. A valid bitangent segment and the four regions created by such a segment.

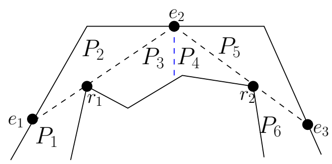

Fig. 5. An L-gadget and the six regions formed by such a structure. Note that a vertical segment is added below e_2 to ensure that P_2 and P_3 are hemimonotone.

Certain pairs of line segments can also be exploited to ensure separation. If there exists a pair of line segments, $\overline{e_1e_2}$ and $\overline{e_2e_3}$, such that e_1, e_2, and e_3 are on the upper subchain, e_2 has a higher x-coordinate than e_1, e_3 has a higher x-coordinate than e_2, $\overline{e_1e_2}$ intersects the lower subchain at a reflex vertex r_1, and $\overline{e_2e_3}$ intersects the lower subchain at a reflex vertex r_2 (again, one could swap the use of upper and lower in the preceeding description). These two line segments (which will be referred to collectively as an *L-gadget*), along with a line segment perpendicular to H that intersects e_2, divide the polygon into six regions (see Figure 5).

- P_1 - A monotone polygon containing v_{front}.
- P_2 - A hemimonotone polygon with $\overline{e_1e_2}$ as a base.
- P_3 - A hemimonotone polygon with $\overline{r_1e_2}$ as a base.
- P_4 - A hemimonotone polygon with $\overline{e_2r_2}$ as a base.
- P_5 - A hemimonotone polygon with $\overline{e_2e_3}$ as a base.
- P_6 - A monotone polygon containing v_{back}.

A point $p_1 \in P_1 \setminus \overline{e_1r_1}$ cannot conflict with a point $p_6 \in P_6 \setminus \overline{r_2e_3}$, because $Vis(p_1)$ cannot extend far enough right to reach e_2, and $Vis(p_6)$ cannot extend far enough left to reach e_2. Therefore, other than the aforementioned line segments, P_1 and P_6 can be guarded by guards that use the same sets of colors.

Another exploitable pair of line segments is a more complicated type of bitangent that will be referred to as a *hybrid bitangent*. This consists of a line segment $\overline{e_1e_2}$ that intersects a reflex vertex r_1 of the upper subchain. There is also a bitangent line segment $\overline{ce_3}$, with endpoint $c \in \overline{r_1e_2}$ that first intersects a reflex vertex of the lower subchain r_2, then a reflex vertex of the upper subchain r_3. The other endpoint e_3 is on the lower subchain. This divides P into six regions (see Figure 6).

- P_1 - A monotone polygon containing v_{front}.
- P_2 - A hemimonotone polygon with $\overline{e_1e_2}$ as a base.
- P_3 - A hemimonotone polygon with $\overline{r_1c}$ as a base.
- P_4 - A hemimonotone polygon with $\overline{ce_2}$ as a base.
- P_5 - A degenerate hemimonotone polygon with $\overline{ce_3}$ as a base.

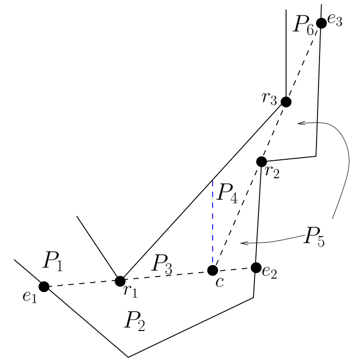

Fig. 6. A hybrid bitangent and the six regions formed by such a structure. Note that a vertical segment is added above c to ensure that P_3 and P_4 are hemimonotone.

- P_6 - A monotone polygon containing v_{back}.

Note that for a point $p_1 \in P_1 \setminus \overline{e_1r_1}$, $Vis(p_1) \subset P_1 \cup P_2$, because r_1 prevents the visibility polygon from extending into the other four regions. Similarly, for a point $p_6 \in P_6 \setminus \overline{r_3e_3}$, $Vis(p_6) \subset P_6 \cup P_5$, because r_3 prevents the visibility polygon from spreading into the other four regions.

Lemma 2: Let P be a monotone polygon that has been partially decomposed through the addition of a set M of line segments. Let $P' \subset P$ be the subpolygon with edges in $M \cup \partial P$ that contains v_{back}. A line segment m_i may be added to M that decomposes P' such that m_i has the following properties:

- The first line segment placed, m_1, contains v_{front} as an endpoint and intersects a reflex vertex.
- The last line segment placed, m_k, contains v_{back} as an endpoint and intersects a reflex vertex.
- If m_i does not have an endpoint at v_{front} or v_{back}, then it is a valid bitangent, the latter half of an L-gadget (the $\overline{e_2e_3}$ edge), or the latter half of a hybrid bitangent (the $\overline{ce_3}$ edge).

Proof: Note that we can assume that $v_{front} \notin Vis(v_{back})$ (otherwise the polygon could just be decomposed with $\overline{v_{front}v_{back}}$). Since $v_{front} \notin Vis(v_{back})$, the Euclidean shortest path from v_{front} to v_{back} must contain a segment $\overline{v_{front}r}$, where r is a reflex vertex. The segment $\overrightarrow{v_{front}r} \cap P$ can be used as m_1.

Next is the case where the m_i does not contain v_{front} or v_{back}. Assume that all prior edges have been placed according to the guidelines. Assume without loss of generality that the endpoint of m_{i-1} that is in P' is on the lower subchain. Call that point p. The other endpoint of $m_{i-1} \cap P'$ must then be on the upper subchain. Call that point q. We assume that no point in m_{i-1} can see v_{back}.

Let y be a ray that extends upwards from p. Let $end(y)$ be the point on ∂P that intersects y that is furthest from y's source point. Since P' is monotone relative to a horizontal line, $end(y)$ must initially be some point in the upper subchain. Rotate the direction of y clockwise until any further rotation would cause $end(y)$ to be on the lower subchain (if there is no such point, then rotate to the first point that would cause $end(y)$ to be on the lower subchain). One of two things must have happened to make further rotation impossible. Either the upper subchain has abruptly ended, in which case y intersects some reflex vertex r_{upper} of the upper subchain, or some reflex vertex r_{lower} of the lower subchain is getting in the way. In the first case, y forms the second half of an L-gadget (where $r_1 = q$, $e_2 = p$, $r_2 = r_{upper}$, $e_3 = end(y)$, and e_1 is a point to the left of q on the m_{i-1} segment), so use $y \cap P'$ as m_i. In the second case, move the source of y towards q along $\overline{pq}$, but ensure that y continues to intersect r_{lower}. Continue moving the source of y toward q until further movement would cause $end(y)$ to move to the lower subchain (again, if no such point exists, move the source until the first point where $end(y)$ is on the lower subchain). If the source of y reaches q, and $end(y)$ is still on the upper subchain, then y forms a valid bitangent ($q = r_1$, $r_{lower} = r_2$, $end(y) = e_2$, and e_1 is a point on the lower subchain below m_{i-1}), so use $y \cap P'$ as m_i. If the movement of y causes the ray to intersect a reflex vertex on the upper chain in the segment $\overline{yr_{lower}}$, then y forms a valid bitangent ($r_1 = r_{upper}, r_2 = r_{lower}, e_2 = end(y)$, and e_1 is on the other side of m_{i-1}, but the valid bitangent need not actually be extended that far, because that region is already decomposed). A vertical line must be extended upward from y to ensure that all subpolygons are hemimonotone. If the movement of the source of y causes $end(y)$ to move to the lower subchain, then one of two things happened. Either the upper subchain ended abruptly, and y intersects some reflex vertex r'_{upper} of the upper subchain, or a reflex vertex r'_{lower} of the lower subchain got in the way. In the first case, y forms the second half of a hybrid bitangent (with $q = r_1$, $p = e_2$, $r_{lower} = r_2$, $r'_{upper} = r_3$, $end(y) = e_3$, and the source of y as c), so use y as m_i. In the second case, continue moving the source of y closer to q, but ensure that y intersects r'_{lower} instead of r_{lower} (and replace the reflex vertex on the lower chain intersected by y repeatedly if the need arises). This process must eventually terminate, because there are only a finite number of reflex vertices on the lower subchain.

For the last edge m_k, note that some point on m_{k-1} must be able to see v_{back}. If the q point of m_{k-1} can see v_{back}, then $\overline{qv_{back}}$ can be used as m_k, because q is a reflex vertex. If $q \notin Vis(v_{back})$, then let x be the point in $\overline{qp} \cap Vis(v_{back})$ closest to q. There must be some reflex vertex r that prevents a point closer to q from seeing v_{back}. Therefore, $\overline{xv_{back}}$ can be used as m_k, because it intersects reflex vertex r. ∎

Guard sets based on the decomposition defined in Lemma 2 that meet certain conditions have connected visibility graphs, as described in the following lemma. For the purposes of the following two lemmas, index the members of set $M = \{m_1, m_2, \ldots, m_{|M|}\}$ in the order that they were added. Also, let $T = \{t_1, t_2, t_3, \ldots, t_{|M|-1}\}$ be a set of points such that t_i is the intersection of m_i and m_{i+1}.

Lemma 3: In a hemimonotone decomposition of a monotone polygon P by the process described in Lemma 2, and a guard set S of P where $T \subset S$, and for all $s \in S$ there exists an $\ell \in M$ such that $s \in \ell$, the graph $VG(S, P)$ is connected.

Proof: If $|T| = 0$, then $v_{front} \in Vis(v_{back})$ and $\overline{v_{front}v_{back}}$ is the only member of M, and a guard set that is placed entirely a single line segment within the polygon must have a connected visibility graph.

If $|T| \neq 0$, then note that, for all $1 < i \leq |T|$, $t_i \in Vis(t_{i-1})$. Therefore, the vertices of T form a connected component in $VG(S, P)$. Since guards are only placed on members of M, and each member of M contains a vertex in T, every vertex of $S \setminus T$ is adjacent to a vertex of T in $VG(S, P)$. Therefore, $VG(S, P)$ is connected. ∎

Additionally, two hemimonotone polygons generated by this decomposition can be covered by guards using the same color set if their bases are sufficiently separated.

Lemma 4: For points $p_i \in m_i \setminus \{t_i\}$, and $p_{i+3} \in m_{i+3} \setminus \{t_{i+2}\}$, $Vis(p_i) \cap Vis(p_{i+3}) = \emptyset$.

Proof: Since m_{i+1} and m_{i+2} are consecutive, either they form an L-gadget, they form a hybrid bitangent, or one of them is a valid bitangent. Therefore, m_{i+1} and m_{i+2} separate P into several subpolygons $[P_1, P_2, P_3, \ldots P_k]$, where $v_{front} \in P_1$, $v_{back} \in P_k$, and for all $2 \leq i \leq k-1$, P_i is hemimonotone. Let ℓ_1 be the boundary between P_1 and P_2, and let ℓ_k be the boundary between P_{k-1} and P_k. Note that no point in $P_1 \setminus \ell_1$ can conflict with any point in $P_k \setminus \ell_k$ (see Figures 4, 5, and 6). Since $m_i \cap \ell_1 = \{t_i\}$ and $m_{i+3} \cap \ell_k = \{t_{i+2}\}$, $Vis(p_i) \cap Vis(p_{i+3}) = \emptyset$. ∎

Now that the decomposition is defined, some results about the number of guards required for the hemimonotone polygons are necessary.

Lemma 5: A hemimonotone polygon P with n vertices can be guarded by a set of $\lfloor n/3 \rfloor$ guards, where all guards are placed on $\overline{v_{left}v_{right}}$.

Proof: The polygon P is monotone with respect to a line H, which we can assume is monotone.

This will proceed via induction on n, the number of vertices in P. The base cases are when $n = 3$, $n = 4$, and $n = 5$. If P is convex, a guard at v_{left} will suffice to guard the whole polygon. Therefore, assume that P is not convex (which means that $n = 4$ or $n = 5$). If $n = 4$ and P is not convex, then either v_1 or v_2 is reflex (assume without loss of generality that v_1 is the reflex vertex). This means that P can be triangulated by adding an edge between v_1 and v_{right}. The two triangles making up P share a vertex at v_{right}, so a single guard on v_{right} can guard all of P.

If $n = 5$, then each triangulation of P contains a vertex that meets all the triangles. If this vertex is v_{left} or v_{right}, then P can be guarded with a single guard at that vertex. If this vertex is v_2, then v_1 and v_3 must be convex (or else they would block v_2's view of v_{left} and v_{right}, respectively). In this case, P can be decomposed into two convex quadrilaterals through the addition of a line segment ℓ that is perpendicular to H and intersects v_2 and $\overline{v_{left}v_{right}}$, and let p be the intersection of ℓ and $\overline{v_{left}v_{right}}$. Since P is monotone with respect to H, the two vertices resulting from the splitting of v_2 must both be convex (or else a line formed by extending ℓ would intersect ∂P at three points). Also note that both vertices at point p must be convex (the sum of their angles in radians must be π, because the vertices are formed by splitting up a line segment). Since v_{left}, v_{right}, v_1, v_3, both vertices at v_2, and both vertices at p are convex, and

$p \in \overline{v_{left}, v_{right}}$, P can be guarded with a single guard at p.

If the vertex that meets all the triangles in a triangulation of P is v_1 or v_3 (assume without loss of generality that it is v_1, then v_2 must be convex, so as to not block v_1's view of v_3). Let ℓ be a line segment that intersects v_2 and $\overline{v_{left}v_{right}}$, and let p be the intersection of ℓ and $\overline{v_{left}v_{right}}$. This divides P into two quadrilaterals. Since v_2, v_{right}, and v_{left} are all convex, v_1 and v_3 are visible from p. Add $\overline{v_1p}$ and $\overline{v_3p}$ to make a triangulation of P with a Steiner point at p where all triangles share p as a vertex. Therefore, a guard placed at p can guard all of P.

Now, for the inductive step. In the inductive step, the goal is to decompose the n-vertex hemimonotone polygon P into two polygons Q and R, where Q is a hemimonotone polygon with at most $n-3$ vertices, and R can be guarded by one guard placed on $\overline{v_{left}, v_{right}}$ (there is a possibility that the guard that guards R is placed in Q). Since the base cases apply when $n \leq 5$, we will assume that P has at least 6 vertices (note that this means P must contain a v_4 vertex). There are eight subcases, based on whether v_1, v_2, and v_3 are convex or reflex. Figure 7 provides illustrations of these subcases.

1) *v_1, v_2, and v_3 are convex* - Let R be the pentagon with vertices v_1, v_2, v_3, v_4, and v_{left}. Since v_1, v_2, and v_3 are convex, $v_4 \in Vis(v_{left})$. This pentagon can be triangulated with v_{left} present in each triangle, so a guard placed at v_{left} guards R. The polygon Q lacks v_1, v_2, and v_3, so it has $n-3$ vertices.
2) *v_1 and v_2 are convex, v_3 is reflex* - If $v_4 \in Vis(v_{left})$, then R has vertices v_1, v_2, v_3, v_4, and v_{left}, and can be triangulated and guarded in the same way as the previous case. If $v_4 \notin Vis(v_{left})$, then let point p be the intersection of $\overline{v_{left}v_{right}}$ and the line formed by extending $\overline{v_3v_4}$. The polygon R has vertices v_{left}, v_1, v_2, v_3, and p. The polygon R is convex and can be guarded with a single guard at v_{left}. The polygon Q lacks v_1, v_2, v_3, and v_{left}, but does contain p as a vertex, so Q has $n-3$ vertices.
3) *v_1 and v_3 are convex, v_2 is reflex* - Let p be the point at the intersection of $\overline{v_{left}v_{right}}$ and a line perpendicular to H that intersects v_2. Since v_1 and v_3 are convex, $p \in Vis(v_4) \cap Vis(v_{left})$. The polygon R has vertices $v_{left}, v_1, v_2, v_3, v_4$, and p. Since P is monotone with respect to H and $\overline{pv_2}$ is perpendicular to H, $p \in Vis(v_1) \cap Vis(v_3)$. Therefore, R can be triangulated with all triangles sharing p as a vertex. Therefore, a guard at p guards all of R. The polygon Q lacks v_1, v_2, v_3, and v_{left}, but does contain p as a vertex, so Q has $n-3$ vertices.
4) *v_1 is convex, v_2 and v_3 are reflex* - Note that if $v_4 \in Vis(v_{left})$, then $v_3 \in Vis(v_{left})$ because otherwise, v_2 would have to block v_{left}'s view of v_3, which means that it would also block v_{left}'s view of v_4, as v_3 is reflex. Therefore, if $v_4 \in Vis(v_{left})$, then R is a pentagon with vertices v_{left}, v_1, v_2, v_3, and v_4, and R can be triangulated as per case 1. If $v_4 \notin Vis(v_{left})$, then let p be the intersection of $\overline{v_{left}v_{right}}$ and the line formed by extending $\overline{v_3v_4}$. Let q be the intersection of $\overline{v_{left}v_{right}}$ and the line perpendicular to H that intersects v_2. The polygon R has vertices v_{left}, v_1, v_2, v_3, and p. The polygon Q lacks v_1, v_2, v_3, and v_{left}, but does contain p as a vertex, so Q has $n-3$ vertices. The polygon R can be guarded by a single guard at point q (note that it is possible for q to not be in R). Since $\overline{pv_2}$ is perpendicular to H, $q \in Vis(v_1) \cap Vis(v_3)$. Therefore, if $q \in R$, then R can be triangulated with a Steiner point at q where all triangles share a vertex at q, so a guard at q guards R. If $q \notin R$, then the polygon R' with vertices $v_{left}, v_1, v_2, v_3, q$ can be triangulated with all triangles sharing a vertex at q. Since $R \subset R'$, and a guard at q guards R', a guard at q guards R.
5) *v_1 is reflex, v_2 and v_3 are convex* - Let point p be the intersection of $\overline{v_{left}v_{right}}$ and the line perpendicular to H that intersects v_1. Let R be the polygon with vertices $v_{left}, v_1, v_2, v_3, v_4$, and p. Since v_2 and v_3 are convex, and $\overline{pv_1}$ is perpendicular to H, $v_2, v_3, v_4 \in Vis(p)$. Therefore, R can be triangulated with all triangles sharing p as a vertex, so a guard at p guards R. The polygon Q lacks v_1, v_2, v_3, and v_{left}, but does contain p as a vertex, so Q has $n-3$ vertices.
6) *v_1 and v_3 are reflex, v_2 is convex* - Let point p be the intersection of $\overline{v_{left}v_{right}}$ and the line perpendicular to H that intersects v_2. Let point q be the intersection of $\overline{v_{left}v_{right}}$ and the line that intersects v_3 and v_4 (such an intersection may not exist if $\overline{v_{left}v_{right}}$ is not long enough). If q exists and $d(v_{right}, q) < d(v_{right}, p)$, then let R be the polygon with vertices v_{left}, v_1, v_2, v_3, and q. In this case, Q lacks v_{left}, v_1, v_2, and v_3, but contains q as a vertex, so Q has $n-3$ vertices. A guard at p is sufficient to guard R, as v_1 is the only reflex vertex of R, and v_1 obviously cannot block p's view of v_{left} or v_2, and the other vertices are on the other side of $\overline{pv_2}$. Therefore, R can be triangulated with a Steiner point at p with all triangles sharing a vertex at p, so a guard at p guards R. If q does not exist or $d(v_{right}, q) \geq d(v_{right}, p)$, then let R be the polygon with vertices $v_{left}, v_1, v_2, v_3, v_4$, and p. In this case, Q lacks v_{left}, v_1, v_2, and v_3, but contains p as a vertex, so Q has $n-3$ vertices. The line segment $\overline{pv_2}$ divides R into two quadrilaterals with one reflex vertex a piece. Since p is not adjacent to either reflex vertex in either quadrilateral, a guard at p can guard both quadrilaterals, and thus all of R.
7) *v_1 and v_2 are reflex, v_3 is convex* - Let point p be the intersection of $\overline{v_{left}, v_{right}}$ and the line perpendicular to H that intersects v_3. Let R be the polygon with vertices $v_{left}, v_1, v_2, v_3, v_4$, and p. The line segment $\overline{pv_3}$ divides R into a pentagon and a triangle. A guard at p obviously guards the entire triangle, and the pentagon has only three convex vertices, where p is the middle convex vertex, so a guard at p guards the entire pentagon as well. The polygon Q lacks vertices v_{left}, v_1, v_2, v_3, but contains p as a vertex, so Q has

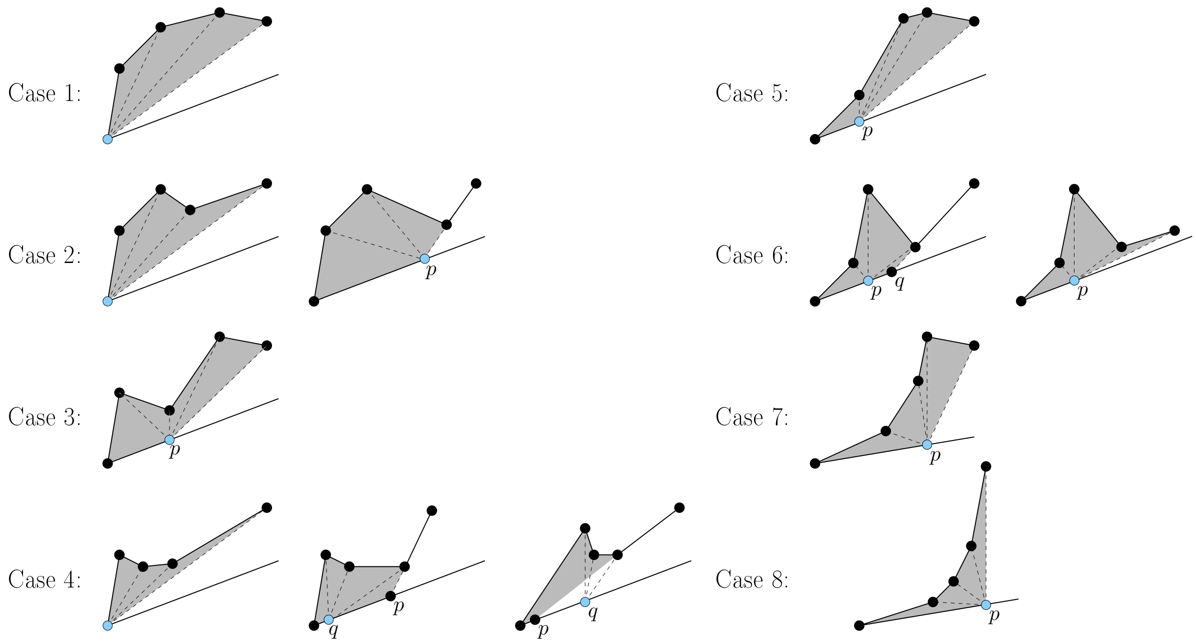

Fig. 7. Illustrations of the eight inductive step cases for Lemma 5. The leftmost vertex is v_{left}. The shaded area represents the subpolygon R. The guard is placed on the blue vertex. The black vertices represent other vertices of P or R. The points p and q are labelled if applicable. These hemimonotone polygons are monotone relative to a horizontal line.

$n-3$ vertices.

8) v_1, v_2, *and* v_3 *are reflex* - Let point p be the intersection of $\overline{v_{left}, v_{right}}$ and the line perpendicular to H that intersects v_4. Let R be the polygon with vertices $v_{left}, v_1, v_2, v_3, v_4$, and p. The polygon Q lacks vertices v_{left}, v_1, v_2, v_3, but contains p as a vertex, so Q has $n-3$ vertices. Since R has only three convex vertices and p is the middle convex vertex, a guard at p guards all of R.

■

Combining these lemmas leads to the desired result.

Theorem 6: For any n-vertex monotone polygon P, $\chi_{CL}(P) \leq n/3 + 12$.

Proof: Decompose P using Lemma 2. Let the guard set S initially consist of the points in T. Due to Lemma 4, the members of T require only 4 colors (because $t_i \in m_{i+1} \setminus \{t_{i+1}\}$, and $t_{i+4} \in m_{i+4} \setminus \{t_{i+3}\}$, so $Vis(t_i) \cap Vis(t_{i+4}) = \emptyset$). Use the methods described in Lemma 5 to make a guard set for each hemimonotone polygon in the decomposition, and add those guard sets to S. Lemma 3 ensures that $VG(S, P)$ is connected.

At most two hemimonotone polygons share m_i as a base. In the decomposition, at least $k-4$ vertices of a k-vertex hemimonotone polygon are vertices of the original, undecomposed monotone polygon that are not shared between two or more hemimonotone polygons (the possible exceptions are v_{left}, v_{right}, v_1, and v_k). Let j be the integer such that the hemimonotone polygons with m_j, m_{j+1}, and m_{j+2} as bases require a maximal sum total of guards. There are 6 hemimonotone polygons total. The total number of vertices in these hemimonotone polygons can be at most $n+24$ (the n vertices of P, plus 4 vertices in each hemimonotone polygon that could be added by the decomposition). By Lemma 5,

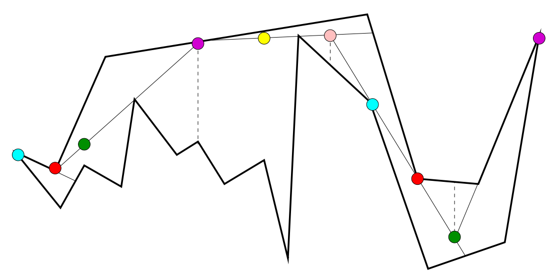

Fig. 8. A monotone polygon decomposed, guarded, and colored. The thin solid lines are the hemimonotone bases, and the dotted lines extend vertically from the intersection of bases if needed to ensure that the subpolygons are hemimonotone.

these polygons can be guarded by $n/3 + 8$ guards, all of which may require unique colors. Since this is the set of hemimonotone polygons with three consecutive bases that require a maximal number of guards, these $\lfloor n/3 \rfloor + 8$ colors are sufficient to color all the guards generated by Lemma 5. With the additional 4 colors required to color the members of T, up to $n/3 + 12$ total colors are used.

Therefore, $\chi_{CL}(P) \leq n/3 + 12$. ■

An example polygon that has been decomposed into hemimonotone polygons and given a guard placement and coloring is shown in Figure 8.

V. Conclusion

We have presented a method of placing a set of landmarks in a monotone polygon so that the visibility graph of the landmarks is connected, each point in the polygon is visible from a landmark, and the landmarks can be divided into a relatively small number of classes such that no two landmarks of the same class are visible from a common point. Such a set can be exploited for navigation by a robot equipped with motion primitives of the form “drive toward

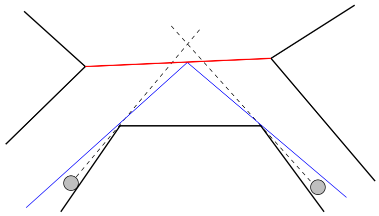

Fig. 9. The black lines are the polygon boundary, and the red line was a boundary added as part of a monotone decomposition. The two guards do not conflict within the monotone polygon under the red line, because they are separated by the blue L-gadget. However, they do conflict in the whole polygon, because they can both see the point at the intersection of the dotted lines.

a landmark of class x".

Of course, robot environments are usually not monotone. However, monotone polygons are frequently generated as intermediate steps in other decomposition algorithms. For example, since monotone polygons can be easily triangulated in linear time, they are often generated during triangulation algorithms [7], [25]. In fact, the technique described in [25] directly creates hemimonotone polygons (referred to as "one-sided" polygons in that work). This invites the possibility of taking a polygon, performing a monotone decomposition on it, and then using the techniques in this paper to color each of the monotone pieces. If some sort of separation guarantee could be made between each of the monotone pieces, then it could be possible to show that for any n-vertex polygon P (not just monotone), $\chi_{CL}(P) \leq n/3 + C$, where C is a constant ([4] shows that this bound must be at least $\lfloor n/4 \rfloor$). Care must also be taken so that the decomposition does not violate the separation conditions within a single monotone polygon (see Figure 9).

There are numerous variations of this problem that would likely produce different bounds. There are many art gallery results specific to rectilinear polygons including [10], [14], [21], and it is likely that rectilinear polygons would require fewer landmark classes than general polygons. One could also discard the assumption that the environment is simply connected and determine bounds on the number of landmark classes in terms of both the number of vertices in the polygon and the number of holes. This paper has assumed that the robot can see the landmark as long as there is no obstacle in the way, regardless of the distance between them. It is unknown how the bounds would change if the robot could only detect landmarks present within some finite distance. Bounds based on the number of inflections could even be determined for curvilinear regions, as bounds were determined for the corresponding art gallery problem in [11].

Finally, while the three conditions listed in the introduction are sufficient for a robot with the "drive toward a landmark of class x" primitives to search an environment for treasure, it is possible that weaker conditions may suffice as well. A categorization of the types of landmark sets that would permit searching by robots of this type would be helpful for describing the power of this very general robot model.

Acknowledgements

This work is supported in part by 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] A. Bachrach, A. S. Huang, D. Maturana, P. Henry, M. Krainin, D. Fox, and N. Roy, "Visual navigation for micro air vehicles," 2011, (workshop manuscript).

[2] V. Chvatal, "A combinatorial theorem in plane geometry," *Journal of Combinatorial Theory Series B*, vol. 18, no. 1, pp. 39–41, 1975.

[3] M. C. Couto, P. J. de Rezende, and C. C. de Souza, "An exact algorithm for minimizing vertex guards on art galleries," *International Transactions in Operational Research*, vol. 18, no. 4, pp. 425–448, 2011.

[4] L. H. Erickson and S. M. LaValle, "An art gallery approach to ensuring that landmarks are distinguishable," in *Robotics: Science and Systems*, Los Angeles, CA, USA, June 2011.

[5] S. Fisk, "A short proof of Chvatal's watchman theorem," *Journal of Combinatorial Theory Series B*, vol. 24, no. 3, p. 374, 1978.

[6] P. Fraigniaud, D. Ilcinkas, G. Peer, A. Pelc, and D. Peleg, "Graph exploration by a finite automaton," *Theoretical Computer Science*, vol. 345, no. 2–3, pp. 331–344, 2005.

[7] M. R. Garey, D. S. Johnson, F. P. Preparata, and R. E. Tarjan, "Triangulating a simple polygon," *Information Processing Letters*, vol. 7, pp. 175–179, 1978.

[8] S. K. Ghosh, "Approximation algorithms for art gallery problems," in *Proc. of Canadian Information Processing Society Congress*, 1987, pp. 429–434.

[9] G. Hernandez-Penalver, "Controlling guards (extended abstract)," in *Canadian Conference on Computational Geometry*, 1994.

[10] J. Kahn, M. Klawe, and D. Kleitman, "Traditional galleries require fewer watchmen," *SIAM Journal on Algebraic and Discrete Methods*, vol. 4, no. 2, pp. 194–206, 1983.

[11] M. I. Karavelas, C. D. Toth, and E. P. Tsigaridas, "Guarding curvilinear art galleries with vertex or point guards," *Computational Geometry: Theory and Applications*, vol. 42, no. 6–7, pp. 522–535, 2009.

[12] M. Knapek, R. Oropeza, and D. J. Kriegman, "Selecting promising landmarks," in *Proc. IEEE International Conference on Robotics and Automation*, 2000, pp. 3771–3777.

[13] D. T. Lee and A. K. Lin, "Computational complexity of art gallery problems," *IEEE Transactions on Information Theory*, vol. 32, no. 2, pp. 276–282, 1986.

[14] A. Lubiw, "Decomposing polygonal regions into convex quadrilaterals," in *Proc. Symposium on Computational Geometry*, 1985, pp. 97–106.

[15] M. Mata, J. M. Armingol, A. de la Escalera, and M. A. Salichs, "A visual landmark recognition system for topological navigation of mobile robots," in *Proc. IEEE International Conference on Robotics and Automation*, 2001, pp. 1124–1129.

[16] D. Meyer-Delius, M. Beinhofer, A. Kleiner, and W. Burgard, "Using artificial landmarks to reduce the ambiguity in the environment of a mobile robot," in *Proc. IEEE International Conference on Robotics and Automation*, 2011, pp. 5173–5178.

[17] T. S. Michael and V. Pinciu, "Art gallery theorems for guarded guards," *Computational Geometry: Theory and Applications*, vol. 26, no. 3, pp. 247–258, 2003.

[18] U. Nehmzow, "Animal and robot navigation," *Robotics and Autonomous Systems*, vol. 15, no. 1–2, pp. 71–81, 1995.

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

[20] C. F. Olson, "Selecting landmarks for localization in natural terrain," *Autonomous Robotics*, vol. 12, no. 1, pp. 201–210, 2002.

[21] J. O'Rourke, "An alternate proof of the rectilinear art gallery theorem," *Journal of Geometry*, vol. 21, no. 1, pp. 118–130, 1983.

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

[23] A. Ranganathan, E. Menegatti, and F. Dellaert, "Bayesian inference in the space of topological maps," *IEEE Transactions on Robotics*, vol. 22, no. 1, pp. 92–107, 2006.

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

[25] C. Yap, "Parallel triangulation of a polygon in two calls to the trapezoidal map," *Algorithmica*, vol. 3, no. 1–4, pp. 279–288.