A Tight Bound For Point Guards in Piece-Wise Convex Art Galleries

Javier Cano* Csaba D. Tóth† Jorge Urrutia‡

September 18, 2010

Abstract

We consider the problem of guarding curvilinear art galleries. A closed arc a joining two points, p and q, in the plane is called a *convex arc* if the curve obtained by joining a with the line segment pq encloses a convex set. A piece-wise convex polygon P with vertices $v_0, \ldots, v_{n-1}$ is the region bounded by a set $\{a_0, \ldots, a_{n-1}\}$ of n convex arcs with pairwise disjoint interiors such that a_i joins v_i to v_{i+1}, addition taken $\bmod n$, each of them *convex with respect to the interior* of P. A piece-wise convex art gallery is the connected region bounded by a piece-wise convex polygon. We show that $\lceil \frac{n}{2} \rceil$ point guards are always sufficient in order to guard a piece-wise convex art gallery. This bound is best possible.

1 Introduction

Let $V = \{v_0, v_1, \ldots, v_{n-1}\}$ be a set of n points in the plane together with a set of straight-line segments $E = \{e_0, e_1, \ldots, e_{n-1}\}$ with pairwise disjoint relative interiors such that the endpoints of e_i are v_i and v_{i+1}, $i = 0, \ldots, n-1$, addition taken $\bmod n$. Let A be the region bounded by the closed curve obtained by joining the elements of E. We call A an *art gallery*, E and V will be called the edge and vertex sets of A, respectively. A set $G \subseteq A$ of points, called guards jointly monitor A if for any point $p \in A$ there is a point $q \in G$ such that the line segment pq lies in A. A classic problem in computational geometry is finding a minimum set of guards for a given art gallery. In the 1970s, Chvátal [6] proved that $\lfloor \frac{n}{3} \rfloor$ guards are always sufficient and sometimes necessary to guard any art gallery with n vertices. Since then many variations of this problem have been studied; see [18, 17, 21] for a detailed reference on art gallery problems. Applications areas of this kind of problems include robotics [12, 22], motion planning [14, 16], computer vision and pattern recognition [2, 19, 20, 23], computer graphics [5, 15], CAD/CAM [3, 7], and wireless networks [8]. Recently Karavelas, Tsigaridas, and Tóth [11] generalized the art gallery problem to *curvilinear art galleries*, where the edges in E are arbitrary Jordan arcs with pairwise disjoint relative interiors, rather than line segments. In general, the minimum number of guards for a curvilinear art gallery cannot be bounded in terms of the number of vertices. If, however, we restrict the arcs to be *convex*, the number of guards is bounded by a function of the number of vertices. A Jordan arc a_i between points v_i and v_{i+1} is *convex* if the closed curve containing a_i and the line segment v_iv_{i+1} enclose a convex region C_i of the plane; see Figure 1 (left).

*Posgrado en Ciencia e Ingeniería de la Computación, Universidad Nacional Autónoma de México, D.F. México, j_cano@uxmcc2.iimas.unam.mx

†Department of Mathematics, University of Calgary, cdtoth@math.ucalgary.ca

‡Instituto de Matemáticas, Universidad Nacional Autónoma de México, D.F. México, urrutia@matem.unam.mx

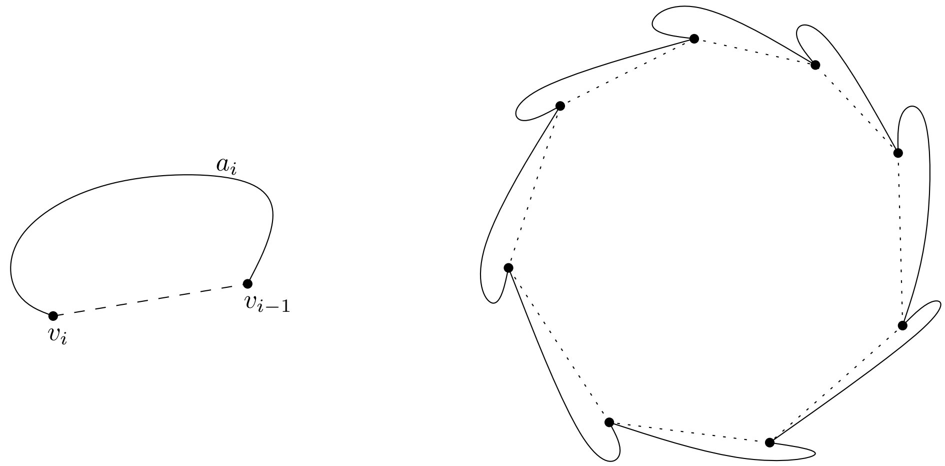

Figure 1: Left: A convex arc. Right: A piece-wise convex polygon that needs $\lceil \frac{n}{2} \rceil$ point guards.

Notice that each convex arc has a convex and a reflex side. A is a *piece-wise convex art gallery* if its interior lies on the convex side of each arc a_i; see Figure 2 (left). In the remainder of this paper, *curvilinear art gallery* will always refer to a piece-wise convex curvilinear art gallery.
Karavelas, Tsigaridas, and Tóth [11] proved that $\lfloor \frac{2n}{3} \rfloor$ vertex guards (that is, guards restricted to be on the vertices of A) are always sufficient and sometimes necessary to guard any curvilinear art gallery with $n \geq 2$ vertices; they also proved that $\lceil \frac{n}{2} \rceil$ point guards (that is, guards anywhere inside A) are sometimes necessary. Cano, Espinosa and Urrutia [4] proved that $\lfloor \frac{5n}{8} \rfloor$ point guards are always sufficient to guard a curvilinear art gallery with $n \geq 2$ vertices. Recently, Karavelas [10] showed that $\lfloor \frac{2n+1}{5} \rfloor$ edge guards (that is, guards allowed to move along edges of the gallery) are always sufficient to guard any curvilinear art gallery with n vertices, and $\lfloor \frac{n}{3} \rfloor$ edge guards are sometimes necessary. In this paper, we prove that any curvilinear art gallery with n vertices can be guarded with at most $\lceil \frac{n}{2} \rceil$ point guards.

Theorem 1. *Let A be a piecewise-convex curvilinear art gallery with n vertices. Then $\lceil \frac{n}{2} \rceil$ guards are always sufficient and sometimes necessary to guard A.*

The proof is based on a convex decomposition of a piece-wise convex art gallery with n vertices into $n+1$ convex cells. We partition the cells into $\lceil \frac{n}{2} \rceil$ sets, each of which can be monitored by a single guard lying on their common boundary. We use a special convex decomposition (discussed in Section 3) in which every convex cell has at least two vertices of the gallery on its boundary. Such a decomposition can be constructed by a technique reminiscent of that of Al-Jubeh *et al.* [1].
It is easy to see that the upper bound of $\lceil \frac{n}{2} \rceil$ in Theorem 1 is best possible. For every $n \geq 3$, there is a curvilinear art gallery with n vertices that requires $\lceil \frac{n}{2} \rceil$ point guards. A construction due to Karavelas, Tóth and Tsigaridas [11] is shown in Figure 1 (right). To close our paper, we give a simpler proof for the sufficiency of $\lfloor \frac{2n}{3} \rfloor$ vertex guards for curvilinear art galleries.

1.1 Notation

We introduce some notation. Let σ be a simple polygonal closed curve in the plane. For two vertices $x, y \in \sigma$ we denote by $[x, y]$ the counterclockwise path that starts with x and ends at y. Clearly, $[x, y] \cup [y, x] = \sigma$ and $[x, y] \cap [y, x] = \{x, y\}$. Let P be a polygonal path. If the number of vertices in P is even, then we call P an *even path*, and an *odd path* otherwise. Analogously, we call a cycle with an even (resp., odd) number of vertices an *even cycle* (resp., *odd cycle*).

2 Convex Decompositions

Let A be a curvilinear art gallery. A *convex decomposition* of A is a finite set C of closed convex regions with pairwise disjoint interiors, called *cells*, such that their union is A. The boundary of every cell in C consists of some straight line segments lying in A and some convex arcs contained in the boundary of A. We define the *vertices* of C to be the endpoints of the straight line segments on the boundaries of cells in C. Every vertex of C is either a vertex of the art gallery or a *Steiner vertex*, which lies in the interior of A or in the relative interior of some arc a_i. The *edges* of C are the portions of these line segments between consecutive vertices of C. See Figure 3 for an example. We denote by $\delta(C)$ the graph formed by the edges and vertices of C. We allow any possible straight line arc a_i to be a (degenerate) cell in C, in this case a_i is also an edge of C lying on the boundary of the degenerate cell.
We define the *dual graph* $D(C)$ of a convex decomposition C as the graph whose vertices are the cells of C, two of which are adjacent if and only if their boundaries intersect. Observe that the cells incident to any Steiner vertex of $\delta(C)$ form a clique in $D(C)$; see Figure 2.

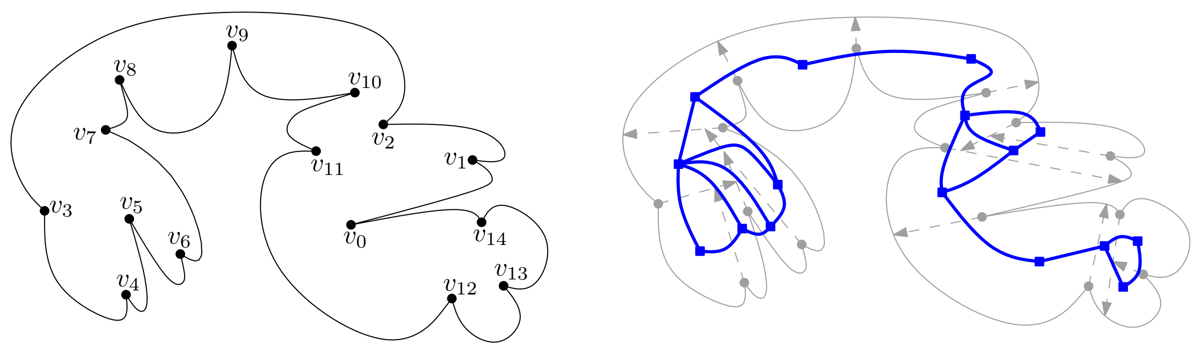

Figure 2: Left: A piece-wise convex art gallery with 15 vertices. Right: The dual graph of a convex decomposition of a curvilinear art gallery.

Normal decompositions. To prove our main result, we construct a family of convex decompositions of A into $n+1$ cells. A convex decomposition C of A with $n+1$ cells is called *normal* if the edges of $\delta(C)$ can be directed so that we obtain a directed graph $\vec{\delta}(C)$ satisfying the following three conditions:

1. the vertices of A have out-degree 1,
2. every vertex v of C located in the interior of A has out-degree 1,
3. every vertex of C in the relative interior of an edge a_i of A has out-degree 0.

Standard convex decompositions. For a curvilinear art gallery A, we can easily construct a special normal decomposition in which every edge lies on a directed segment emitted by one of the vertices. For every vertex v_i of A, let W_i be the wedge formed by all rays emitted from v_i that partition the (counterclockwise) angle between tangent lines to a_i and a_{i-1} at v_i into two convex angles. (If this angle is already convex, then W_i is the angular domain between the tangents of a_{i-1} and a_i, otherwise it is between the tangents of a_i and a_{i-1}.)

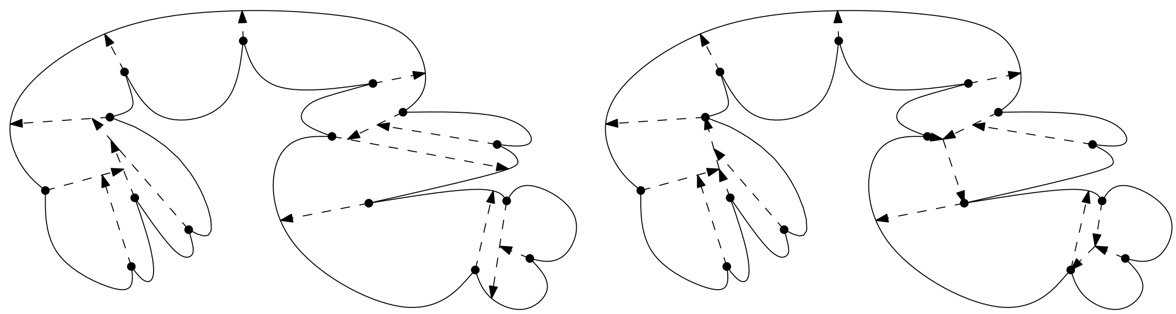

Figure 3: A standard convex decomposition (left) and a good decomposition (right) of a curvilinear art gallery.

Lemma 2. *For every vertex v_i of A, there is a directed segment $\vec{r}_i$ lying in $A \cap W_i$ that connects v_i to another point on the boundary of A, which is not in the relative interior of arc a_i or a_{i+1}.*

Proof. We construct a directed segment $\vec{r}_i$ for every vertex v_i. We distinguish two cases. First suppose that C_{i-1} and C_i (as defined in the introduction) intersect in a single point v_i. Let h be a separating line between such that $a_{i-1} \setminus \{v_i\}$ and $a_i \setminus \{v_i\}$ lie in two different open halfplanes bounded by h. It is clear that h partitions the (counterclockwise) angle between tangent lines to a_i and a_{i-1} at v_i into two convex angles. Shoot a ray from v_i into the interior of A along h, and let $\vec{r}_i$ be the part of such a ray from v_i to the first intersection point with the boundary of A.
Now suppose that C_{i-1} and C_i intersect in several points (including v_i). If a_{i-1} or a_i is a line segment, then let $\vec{r}_i$ be this segment with a direction from v_i to v_{i-1} or v_i. Now suppose that neither a_{i-1} nor a_i is a line segment. Then either the directed segment $\overrightarrow{v_i v_{i-1}}$ lies in C_i or the directed segment $\overrightarrow{v_i v_{i+1}}$ lies in C_{i-1}. Let $\vec{r}_i$ be the initial portion of this directed segment from v_i to the first intersection point with the boundary of A. By construction, the endpoint of $\vec{r}_i$ cannot be in the relative interior of arc a_{i-1} or a_i. $\square$

We construct a normal decomposition for a given curvilinear art gallery as follows. For $i = 0, 1, \ldots, n-1$, draw a directed line segment starting from v_i along a directed segment $\vec{r}_i$ as described in Lemma 2 until it hits the boundary of A or a previously drawn segment. See Figure 3 (left). It is clear that the n directed segments decompose A into $n+1$ convex cells (degenerate cells are possible if a_i or a_{i-1} is a line segment collinear with $\vec{r}_i$). We call any convex decomposition constructed in this way a *standard decomposition* of A. Observe that the directions of segments $\vec{r}_i$ induce a direction on all edges of $\delta(C)$. Let $\vec{\delta}(C)$ denote this directed graph. It is easy to verify now that every standard decomposition is a normal decomposition.

Cyclic and acyclic cells. Typically every cell in a normal decomposition is adjacent to the boundary of A. Some cells, however, may be disjoint from the relative interior of every convex arc a_i, $i = 0, \ldots, n-1$; see Figure 1 (right). Since the out-degree of every vertex of $\vec{\delta}(C)$ on the boundary of such cell is one, the boundary is a directed cycle. We say that a cell in a normal decomposition is *cyclic* if it is disjoint from the relative interior of every arc a_i, $i = 0, \ldots, n-1$, and *acyclic* otherwise.

Good and bad cells. Let C be a normal convex decomposition of A. A cell c of C is called *good* if its boundary contains at least two vertices of A, otherwise c is called *bad.* A convex decomposition of a curvilinear polygon is called *good* if all of its cells are good; see Figure 3 (right). The following observation about standard convex decompositions will be useful.

Observation 1. *Every cell of a standard convex decomposition of A contains at least one vertex of A on its boundary.*

Proof. Let c be a cell in a standard convex decomposition C. Let i, $0 \le i \le n-1$, be the largest index such that the boundary of c contains some portion of the line segment starting from v_i. Since no other edge of c can hit the relative interior of this segment, the segment endpoint v_i also lies on the boundary of c. □

We can now classify the components of $\vec{\delta}(C)$.

Lemma 3. *Let C be a normal decomposition of A. Then every connected component of $\vec{\delta}(C)$ is either a directed tree rooted at a point in the relative interior of an edge of A, or a directed graph with exactly one directed cycle which bounds a cyclic cell.*

Proof. Consider a connected component t of $\vec{\delta}(C)$. Since the out-degree of every vertex of $\vec{\delta}(C)$ is at most one, t contains at most one directed cycle. If t contains no cycle, then it is a rooted tree, and the root has to be a point with out-degree 0; that is, a Steiner point lying in the relative interior of an edge of A. Now suppose that t contains a cycle, say σ. Note that the interior of σ lies in the interior of A, since A is simply connected. It remains to show that σ bounds a single cell in C. Suppose, to the contrary, that at least two cells of C are inside σ. These cells must be separated by some edges of $\delta(C)$ which are not part of σ. None of these edges can start from a vertex of σ, otherwise the out-degree restriction is not satisfied. Hence, at least one of these edges has to start from a vertex of A. This, however, is impossible since the interior of σ lies in the interior of A. We conclude that the interior of σ is a single cell in C. □

Special cells for each component of $\vec{\delta}(C)$. Let t be a connected component of $\vec{\delta}(C)$. We say that a cell $c \in C$ is *incident to t* if the boundary of c contains at least one edge of t. We specify some *special* cells for t. If t contains a cycle σ, then let the cell bounded by σ be special. If t is a directed tree rooted at some vertex x (lying in the relative interior of some arc a_i), then let the two cells incident to x having an arc of a_i on its boundary be special.

3 Constructing a Good Normal Decomposition

In this section we construct a good normal decomposition for a curvilinear art gallery with $n \ge 3$ vertices.

Lemma 4. *Every curvilinear art gallery with $n \ge 3$ vertices has a good normal decomposition.*

Proof. Let C be a standard convex decomposition of A. If C is a good decomposition, then our proof is complete. Otherwise we will *deform* $\vec{\delta}(C)$ continuously into a good decomposition. Our algorithm successively processes every bad cell of C, deforming its boundary until it contains at least two vertices of A. During the deformation, we maintain a normal decomposition and good cells remain good. Specifically, we maintain the following four invariants:

- I1 C is a normal decomposition of A.
- I2 For every edge e of $\vec{\delta}(C)$, there is a vertex v of A such that $\vec{\delta}(C)$ contains a directed path of collinear edges, including e, that either starts from v or ends at v.

I3 If a cell $c \in C$ is incident to a vertex v of A, then c remains incident to v.

I4 If a cell $c \in C$ is cyclic, then it remains cyclic.

Note that Invariants I1 and I2 hold for every standard convex decomposition. Invariant I3 implies that when all bad cells have been processed, we obtain a good decomposition of A.
Consider a bad cell c of C. We process cell c while maintaining invariants I1–I4. We first process all acyclic bad cells and then process all cyclic bad cells as follows.

Processing an acyclic bad cell. Let $c \in C$ be an acyclic bad cell. By Observation 1 and invariant I3, the boundary of c contains exactly one vertex of A, which we denote by v_i. The edges of $\vec{\delta}(C)$ on the boundary of c induce a directed path π in $\vec{\delta}(C)$ which starts at vertex v_i. Since c is acyclic, π ends at a point x in the relative interior of an edge a of A adjacent to v_i. Without loss of generality, we may assume that $a = a_i$, and thus v_{i+1} is the other endpoint of a_i. Refer to Fig. 4.

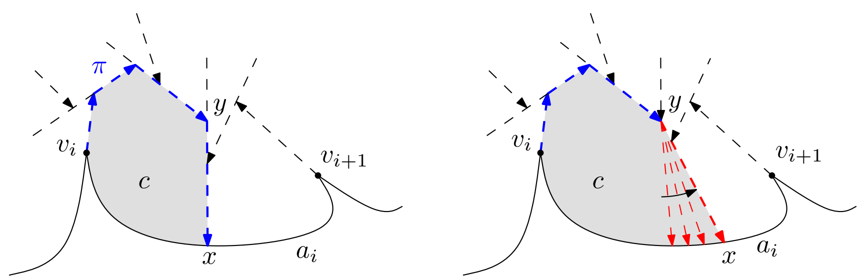

Figure 4: Stretching segment $\vec{yx}$.

Observe that some edges along π may be collinear. Let $e_1, e_2, \ldots, e_k$ be the maximal directed line segments that contain collinear edges of π in this order such that e_1 starts from v_i and e_k ends at x. Let $\overrightarrow{yx} = e_k$. We process c as follows. Move point x continuously along a_i towards v_{i+1} and stretch the directed edge $\overrightarrow{yx}$ until one of the following possibilities arises:

1. We have $k \geq 2$ and $\overrightarrow{yx}$ becomes collinear with e_{k-1}. Then set $k := k - 1$, recompute y, and continue moving x (see Figure 5).

2. We have $x = v_{i+1}$ (see Figure 6, left) or some vertex v_r of A appears in the relative interior of $\overrightarrow{yx}$ (see Figure 6, right).

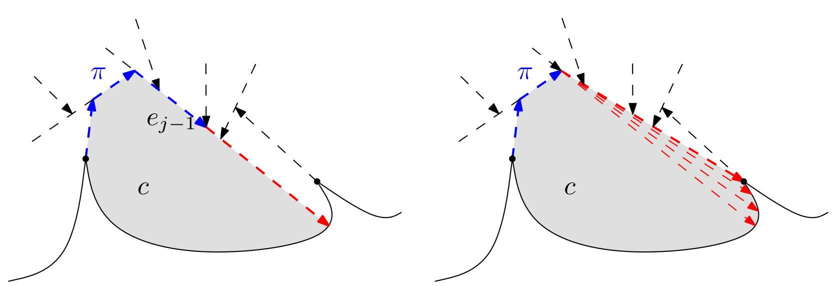

Figure 5: $\vec{yx}$ becomes collinear with e_{k-1}.

While stretching segment $\overrightarrow{yx}$, the edges of $\vec{\delta}(C)$ that hit $\overrightarrow{yx}$ from the opposite side of c are continuously shortened, and the edges and Steiner vertices completely swept by $\overrightarrow{yx}$ disappear. If at the beginning of processing cell c, point x is adjacent to another bad acyclic cell c', and the outgoing edge of v_{i+1} is shortened to a single point when we move x to v_{i+1}, then add a new directed edge $\overrightarrow{v_{i+1}v_i}$ (which effectively decomposes cell c into two good cells). This completes the description of the processing of cell c. The process terminates, since at each step, either k is decremented or c becomes a good cell.

We show next that invariants I1–I4 are maintained. First we show that c remains convex. The first stopping rule guarantees that c has convex angles at every internal vertex of path π. If $y = v_i$, then cell c remains convex at v_i, since $\overrightarrow{yx}$ connects two points of the convex arc a_i. It is clear that the cells on the opposite side of $\overrightarrow{yx}$ remain convex. The only case when a cell c' can disappear is when c' is a bad cell incident to v_{i+1}, we move x to v_{i+1}, and the outgoing edge of v_{i+1} is shortened to a single point. In this case, however, we add a new outgoing edge at v_{i+1}, and split c into two good cells, thereby restoring a normal decomposition. Invariant I2 continues to hold for all edges of $\vec{\delta}(C)$ that we do not modify. The edges along $\overrightarrow{yx}$ do not satisfy I2 *during* the continuous motion. At the end of the process, $\overrightarrow{yx}$ contains a vertex of A, and so I2 becomes true for all edges along $\overrightarrow{yx}$. It is easy to verify that invariants I3 and I4 are maintained.

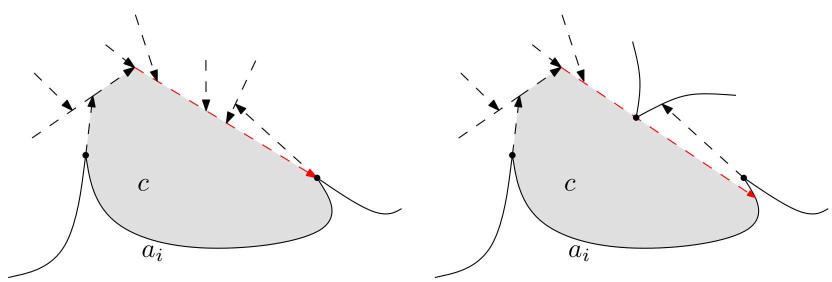

Figure 6: $\vec{yx}$ hits a vertex of A.

Also observe that if $\vec{\ell}$ hits v_{i+1}, a cell c' on the opposite side of $\overrightarrow{yx}$ may become cyclic, see Figure 6 (left).

Processing a cyclic bad cell. Let c be a cyclic bad cell of C. The boundary of c is a directed cycle π in $\vec{\delta}(C)$. Refer to Fig. 7. By Observation 1 and invariant I3, its boundary contains some vertex v_i of A. Some edges along π may be collinear. Let $e_1, e_2, \ldots, e_k$ be the maximal directed line segments that contain collinear edges of π in this order such that e_1 starts from v_i and e_k ends at v_i. Note that $k \geq 3$, and let $\overrightarrow{yx} = e_{k-2}$. We process c as follows. By invariant I2, $\vec{\delta}(C)$ contains a directed path through e_{k-1} and starting or ending at some vertex w of A. Since the directed path passing through e_{k-1} bends at the endpoint of e_{k-1}, there is a collinear directed path *from* w through e_{k-1}, including point x. Let $\vec{\ell} = \overrightarrow{wx} \subset \vec{\delta}(C)$. Move point x continuously along $\vec{\ell}$ towards w and stretch the directed edge $\overrightarrow{yx}$ until one of the following possibilities arises:

1. We have $k \geq 4$ and $\overrightarrow{yx}$ becomes collinear with e_{k-3}. Then set $k := k - 1$, recompute y, and continue moving x.

2. We have $x = w$ (see Figure 7, right) or some vertex v_r of A appears in the relative interior of $\overrightarrow{yx}$.

This completes the description of the processing of a cyclic cell c. The process terminates, since at each step, either k is decremented or c becomes a good cell.

We show next that invariants I1–I4 are maintained. The first stopping rule guarantees that c has convex angles at every vertex of cycle π. It is clear that the cells on the opposite side of $\overrightarrow{yx}$ remain

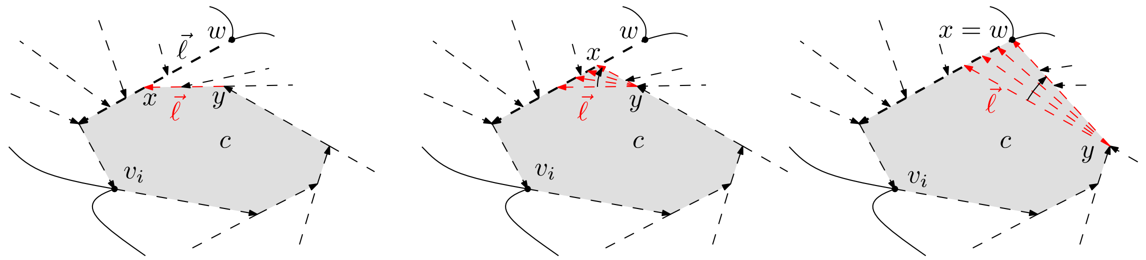

Figure 7: Deforming a cyclic cell.

convex, and cannot disappear. Invariant I2 continue to hold for all edges of $\vec{\delta}(C)$ that we do not modify. Similar to the processing of acyclic cells, the edges along $\overrightarrow{yx}$ do not satisfy I2 *during* the continuous motion. At the end of the process, $\overrightarrow{yx}$ contains a vertex of A, and so I2 becomes true for all edges along $\overrightarrow{yx}$. It is easy to verify that invariants I3 and I4 are maintained. □

4 The Dual Graphs of Good Normal Decompositions

Let C be a good normal decomposition of A. Since C is fixed, we will refer to $\vec{\delta}(C)$ simply as $\vec{\delta}$. Let t be a connected component of $\vec{\delta}$ adjacent to at least three cells of C. Let $D(t)$ be the subgraph of $D(C)$ induced by the cells of C incident to t. In this section, we prove several important properties of $D(t)$. We begin with an easy observation.

Observation 2. *Every vertex of $D(C)$ has degree at least 2.*

Proof. Let $c \in C$ be a convex cell. Clearly the degree of every vertex is at least the number of edges of $\vec{\delta}$ on its boundary, and every cell is adjacent to at least one edge of $\vec{\delta}$. Suppose that c is adjacent to exactly one edge e of $\vec{\delta}$. Since c is good, both endpoints of e are vertices of A. Let $\vec{e} = \overrightarrow{uv}$. Since v has out-degree 1, there is another edge, say e', that starts from v, and lies between some cells c_1 and c_2. Since v is on the boundary of c, it is adjacent to both c_1 and c_2 in $D(C)$. □

The following lemmas are the key to our result.

Lemma 5. *$D(t)$ contains a cycle that passes through all acyclic cells adjacent to t.*

Proof. Recall that by Lemma 3, t is either a rooted tree or it contains a directed cycle bounding a cyclic cell of C.
We construct a cycle $H_a(t)$ in $D(t)$ as follows. Walk around the boundary of A starting from an arbitrary point. We say that the walk *encounters* a cell c if the walk traverses an arc on the boundary of c (rather than either passing through only one vertex on the boundary of c or none at all). Relabel the cells represented by vertices of $D(t)$ along the boundary of A to $c_1, \ldots, c_k$ in the order in which they are encountered in this walk. The order is well defined: if the walk encounters a cell c_i twice, say at arcs γ_1 and γ_2, then the portion of the boundary of A between γ_1 and γ_2 is separated from t by cell c_i, and cannot encounter any other cell adjacent to t. Let $H_a(t) = (c_1, \ldots, c_k)$. It is clear that consecutive cells in $H_a(t)$ are adjacent in $D(t)$. That is, $H_a(t)$ is a simple cycle in $D(t)$, which passes through all acyclic cells adjacent to t, as required. □

Lemma 6. *Let c be a special cell adjacent to t, and assume that c is not adjacent to any other component of $\vec{\delta}$. Then there is a vertex $v(c)$ of A incident to c such that $v(c)$ is incident to two more cells $c_1, c_2 \in C \setminus \{c\}$ which are consecutive in $H_a(t)$.*

Proof. Suppose first that t is a directed tree (see Figure 8). Since c is special, the root x of t is on the boundary of c. Suppose that the part of t lying on the boundary of c is the directed path π from vertex v_i to x. Since c is not adjacent to any other component of $\vec{\delta}$, the root x lies on a convex arc of A incident to v_i. However, cell c is good, and so the boundary of c contains at least one more vertex of A, $v(c)$, which is an internal vertex of path π. Since $v(c)$ is an internal vertex of π, there are two cells, say $c_1, c_2 \in C \setminus \{c\}$, whose boundaries each contain $v(c)$ and some initial part of a convex arc incident to $v(c)$. By construction, c_1 and c_2 are adjacent in the cycle $H_a(t)$. Now suppose that c is a cyclic cell, bounded by a cycle σ of t. Let $v(c)$ be an arbitrary vertex of A along σ. Let $c_1, c_2 \in C \setminus \{c\}$ be the cells whose boundaries each contain $v(c)$ and some initial part of a convex arc incident to $v(c)$. Again, c_1 and c_2 are adjacent in the cycle $H_a(t)$. □

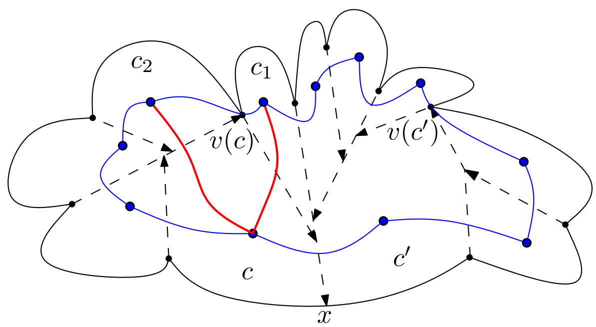

Figure 8: Two cells c_1, c_2, adjacent to c in $D(t)$.

Corollary 7. *$D(t)$ is Hamiltonian.*

Proof. If t is a directed tree, then $H_a(t)$ is a Hamiltonian cycle of $D(t)$ by Lemma 5. If t has a cycle, then cycle $H_a(t)$ passes through all acyclic cells, but misses one cyclic cell c. By Lemma 6 there are two consecutive cells, c_1 and c_2, in $H_a(t)$ that are both adjacent to c in $D(t)$. By removing the edge c_1c_2 from $H_a(t)$ and connecting c with c_1 and c_2 we obtain a Hamiltonian cycle in $D(t)$. □

In the remainder of this paper, we denote by $H(t)$ the Hamiltonian cycle constructed in the proof of Corollary 7.

Basic cycles. Let γ be a simple cycle in graph $D(t)$. We define region R_γ in the plane as the union of the cells in γ. Observe that region R_γ is simply connected if and only if the cells in γ do not enclose any cyclic cell $c \notin \gamma$. We say that γ is a *basic cycle* of $D(t)$ if R_γ is simply connected; see Figure 9 (left). In particular, $H(t)$ is a basic cycle, and if t is a tree, then every simple cycle in $D(t)$ is basic. We denote by $D(t, \gamma)$ the subgraph of $D(t)$ induced by the vertices of γ.

Lemma 8. *Every basic cycle γ in $D(t)$ with $k \geq 3$ cells contains three consecutive cells incident to a vertex of $\vec{\delta}$.*

Proof. Label the cells in γ counterclockwise by $c_0, c_1, \dots, c_{k-1}$ along the boundary of R_γ. If c_i and c_j, $i+1 < j$, are adjacent in $D(t)$, then $\gamma' = (c_i, c_{i+1}, \dots, c_j)$ is called a *sub-cycle* of γ, addition taken $\bmod k$, Figure 9 (right). Every sub-cycle γ' is a basic cycle, since $R_{\gamma'} \subset R_\gamma$ contains no cell in its interior. It is enough to show that γ has a sub-cycle γ' of 3 cells: the common boundary between the three consecutive cells in γ' meets, since R_γ is simply connected.

Let γ' be the smallest sub-cycle of γ. By relabelling the cells if necessary, we may assume that $\gamma' = (c_0, c_1, \ldots, c_i)$. If $i = 2$, then our proof is complete. Assume that $i \geq 3$. The boundary between c_0 and c_1 is a (possibly degenerate) line segment s. Since $R_{\gamma'}$ is simply connected, one endpoint of s must be incident to some other cell c_j in γ'. Hence $\gamma'' = (c_0, c_1, \ldots, c_j)$ is a strictly smaller sub-cycle of γ, contradicting the minimality of γ'. □

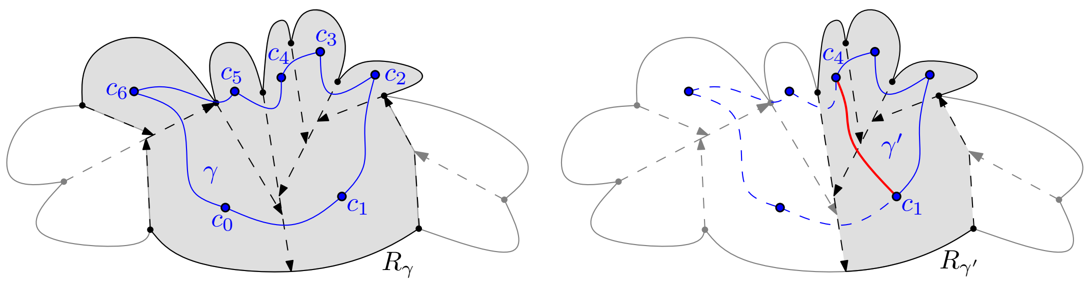

Figure 9: Left: A basic cycle γ. Right: A sub-cycle γ' of γ.

Lemma 9. *For every basic cycle γ in $D(t)$ with $k \geq 3$ cells, graph $D(t, \gamma)$ has a clique cover of size $\lfloor k/2 \rfloor$ such that the cells in each clique can be guarded from a single point.*

Proof. If γ is an even cycle, then it has a perfect matching of size $\lfloor k/2 \rfloor$, which is a desired clique cover, so we are done. Suppose that γ is odd. By Lemma 8, γ contains three consecutive cells incident to a common vertex of $\vec{\delta}$. This triple together with a perfect matching on the remaining $k - 3$ vertices of γ is a desired clique cover of size $\lfloor k/2 \rfloor$. □

5 Constructing a Guard Set

Proof of Theorem 1. Let A be a curvilinear art gallery with $n \geq 3$ vertices. Fix a good normal decomposition C of A. As we noted before, each vertex of $\vec{\delta}$ corresponds to a clique in the dual graph $D(C)$. To show that A can be guarded by at most $\lceil \frac{n}{2} \rceil$ point guards, it is enough to show that $D(C)$ has a clique cover of size at most $\lceil \frac{n}{2} \rceil$ such that each clique is induced by some vertex of $\vec{\delta}$, and so the convex cells in each of these cliques can be guarded from a single point. In the remainder of the proof we describe an algorithm for constructing a clique cover of $D(C)$ having this property and size at most $\lceil \frac{n}{2} \rceil$.

We define a graph Γ on the connected components of $\vec{\delta}$. Two connected components t and t' of $\vec{\delta}$ are adjacent in Γ if and only if there is a cell $c \in C$ adjacent to both of them; see Figure 10. Notice that all the components of $\vec{\delta}$ incident to a cell $c \in C$ induce a clique in Γ. Relabel the components of $\vec{\delta}$ by $t_1, \ldots, t_k$ according to a breadth-first search traversal of Γ. Notice that this labelling has the property that every t_m is adjacent to at most one cell which is adjacent to some previous component t_i with $i < m$, otherwise A would not be simply connected. Let $n(t_m)$ be the number of cells in $D(t_m)$. Note that $n(t_m) > 1$.

We construct a clique cover $\mathcal{G}$ of $D(C)$ as follows. Initially, let $\mathcal{G} = \emptyset$. Our algorithm runs in k iterations. In iteration $m = 1, 2, \ldots, k$, we process graph $D(t_m)$ and compute a set $\mathcal{G}_m$ such that the cliques in $\cup_{i=1}^{m} \mathcal{G}_i$ cover all but at most one cells in $D(t_m)$. We may leave at most one cell in $D(t_m)$ uncovered provided that it is contained in $D(t_j)$ for some $j > m$ (which will be processed later).

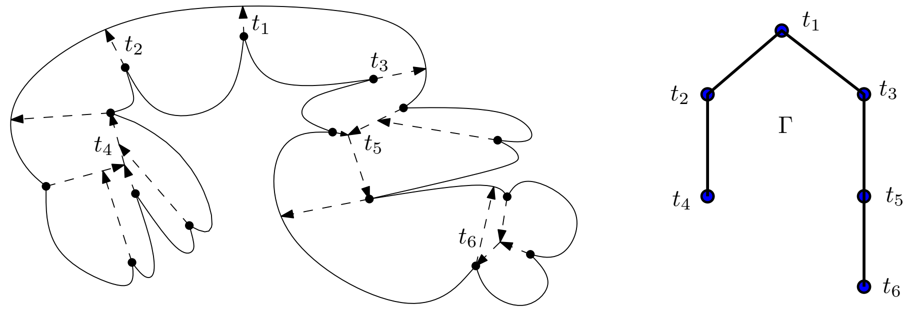

Figure 10: A good normal decomposition of a curvilinear art gallery, and the corresponding graph Γ.

Recall that for every t_m, at most one cell of $D(t_m)$ is contained in a previously processed $D(t_i)$, $i < m$. This cell may or may not be covered by a clique in $\mathcal{G}_m$. Accordingly, at the beginning of the m-th iteration two cases may arise:

Case a: No cell of $D(t_m)$ has been covered in any previous iteration. We proceed as follows: If $n(t_m) = 2$, then $D(t_m)$ is a clique of size $1 = n(t_m)/2$. If $t_m \geq 3$, by Corollary 7 and Lemma 9, $D(t_m)$ has a desired clique cover of size $\lfloor \frac{n(t_m)}{2} \rfloor$.

Case b: Exactly one cell of $D(t_m)$, say c, was covered in a previous iteration. If $n(t_m) = 2$, then let $\mathcal{G}_m = \emptyset$. Then one cell in $D(t_m)$ is still uncovered. By Observation 2, the uncovered cell is adjacent to some other component t_j with $j > m$, which will be processed later. In the remainder of the proof, we assume $n(t_m) \geq 3$. Suppose that cell $c \in D(t_m)$ is already covered. We will distinguish several subcases. In each subcase, we partition $D(t_m) \setminus \{c\}$ into subgraphs that are cliques induced by a vertex of $\vec{\delta}$, even paths, basic cycles, and at most one singleton (a cell adjacent to a component t_j, $j > m$). A perfect matching of an even path of length ℓ is a clique cover of size $\ell/2$. By Lemma 9, a basic cycle of size ℓ has a clique cover of size $\lfloor \ell/2 \rfloor$. This guarantees that we obtain a desired clique cover $\mathcal{G}_m$ of size at most $\lfloor \frac{n(t_m)-1}{2} \rfloor$. We continue with the details.
If $n(t_m)$ is odd, then $H(t_m) \setminus \{c\}$ is an even path. If $n(t_m)$ is even, then several sub-cases arise depending on whether $D(t_m)$ has a cyclic cell or not.

Case b1: $D(t_m)$ has a cyclic cell $c_1 \in D(t_m)$. Note that $c \neq c_1$, since the cyclic cell is adjacent to t_m only. Since c_1 is a good cell, its boundary contains at least two vertices of A. By Lemma 6, each vertex of A on the boundary of c_1 is incident to two consecutive cells in the cycle $H_a(t_m)$. Therefore there are two pairs of consecutive vertices, c_2, c_3 and c_4, c_5 in counterclockwise order along $H_a(t_m)$ (with possibly $c_3 = c_4$ or $c_2 = c_5$) such that $\{c_1, c_2, c_3\}$ and $\{c_1, c_4, c_5\}$ are cliques, each of which can be guarded from a vertex of A. See Figure 11.
Partition the cycle $H_a(t_m)$ into paths $[c_3, c_4]$ and $[c_5, c_2]$. Suppose without loss of generality that $c \in [c_5, c_2]$. Clearly $[c_5, c_2] \setminus \{c\}$ is the union of two (possibly empty) paths, which we denote by p_2 and p_5 such that $c_2 \in p_2$ and $c_5 \in p_5$ respectively. Note that either p_2 or $p_2 \setminus \{c_2\}$ is even; denote this path by p_2'. Similarly either p_5 or $p_5 \setminus \{c_5\}$ is even, and denoted by p_5'. Since c_1 is adjacent to c_2, c_3, c_4, c_5, the graph $D(t_m) \setminus (\{c\} \cup p_2' \cup p_5')$ has a spanning cycle, that contains c_1, and so it is a basic cycle.

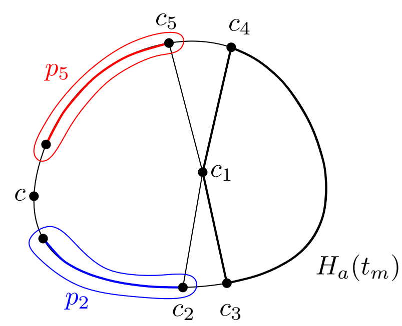

Figure 11: $[c_5, c_2] \setminus \{c\}$ is the union of paths p_2 and p_5.

Case b2: $D(t_m)$ has no cyclic cell. Let c_1 and c_2 be the special cells adjacent to t_m. We distinguish three subcases depending on whether c_1 and c_2 are adjacent to any other component of $\vec{\delta}$ or we can apply Lemma 6:

Case b2.1: Both c_1 and c_2 are adjacent to some other components of $\vec{\delta}$. Recall that $H(t_m)$ is a Hamiltonian cycle of $D(t_m)$ in which c_1 and c_2 are consecutive cells. Since $n(t_m)$ is even, $H(t_m) \setminus \{c\}$ is an odd path. First suppose that c is a special cell of t_m, say $c = c_1$. Then $H(t_m) \setminus \{c_1, c_2\}$ is an even path, and we leave c_2 uncovered. Now suppose that c is not a special cell of t_m. Then $H(t_m) \setminus \{c, c_1\}$ or $H(t_m) \setminus \{c, c_2\}$ is the union of two even paths. Suppose without loss of generality that this happens for $H(t_m) \setminus \{c, c_1\}$, and we leave c_1 uncovered.

Case b2.2: Exactly one of c_1 or c_2 is adjacent to some other component of $\vec{\delta}$. Assume without loss of generality that c_1 is adjacent to no other component of $\vec{\delta}$. By Lemma 6, there are two consecutive cells, c_3 and c_4, along $H(t_m)$ such that $\{c_1, c_3, c_4\}$ is a clique which can be guarded from a single point. The edge c_1c_3 splits the cycle $H(t_m)$ into two cycles, which we denote by say H_1 and H_2 respectively such that $H_1 \cap H_2 = c_1c_3$. We may assume without loss of generality that $c_4 \in H_1$ and $c_2 \in H_2$; see Figure 12. Now we have:

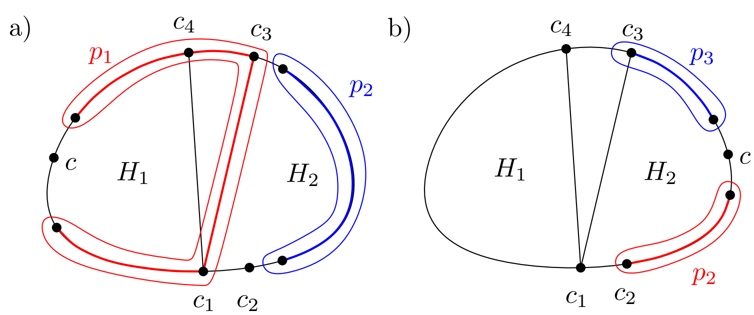

Figure 12: Illustrations for **Case b2.2.1** and **Case b2.2.2**.

Case b2.2.1: $c \in H_1$. Then the cells of $D(t_m) \setminus \{c, c_2\}$ lie on two paths: $p_1 = H_1 \setminus \{c\}$ and $p_2 = H_2 \setminus \{c_1, c_2, c_3\}$. See Figure 12(a). We leave c_2 uncovered. If p_1 is even, then p_2 is even, too, and we have a desired partition of $D(t_m) \setminus \{c\}$. If p_1 is odd, then p_2 is odd, too. By construction, edge c_4c_1 is a triangular chord of p_1. We obtain an even path p'_1 from p_1 by replacing edges c_1c_3 and c_3c_4 with the edge c_1c_4. We obtain an even path p'_2 from p_2 by appending c_3 to it.

Case b2.2.2: $c \in H_2$. Then $H_2 \setminus \{c, c_2\}$ is the union of two (possibly empty) paths, which we denote by p_2 and p_3 such that $c_2 \in p_2$ and $c_3 \in p_3$ respectively; see Figure 12(b). Note that either p_2 or $p_2 \setminus \{c_2\}$ is even; denote this path by p_2'. Similarly either p_3 or $p_3 \setminus \{c_3\}$ is even, and denoted by p_3'. Since c_1 is adjacent to c_3 and c_4, the graph $H_1 \setminus p_3'$ has a spanning cycle H_1' which is a basic cycle. The even paths p_2' and p_3', basic cycle H_1', and possibly leaving cell c_2 as a singleton, we have a desired partition of $D(t_m) \setminus \{c\}$.

Case b2.3: Neither c_1 nor c_2 is adjacent to any other component of $\vec{\delta}$. By Lemma 6, there are two consecutive cells, c_3 and c_4, along $H(t_m)$ such that $\{c_2, c_3, c_4\}$ is a clique which can be guarded from a single vertex $v(c_2)$. Similarly, there are two consecutive cells, c_5 and c_6, along $H(t_m)$ such that $\{c_1, c_5, c_6\}$ is a clique which can be guarded from a single vertex $v(c_1)$. We distinguish two subcases depending on whether the vertices $v(c_1)$ and $v(c_2)$ are distinct:

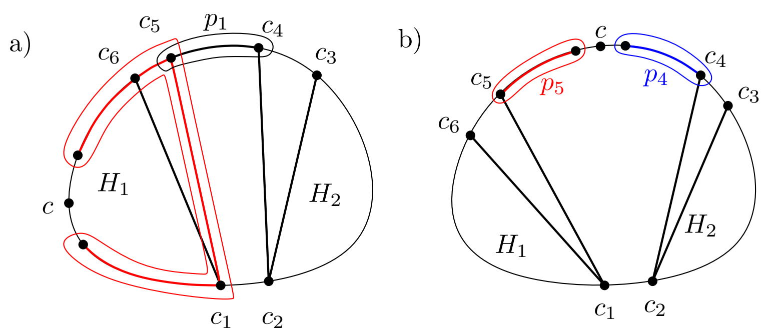

Figure 13: Illustration for **Case b2.3.1**.

Case b2.3.1: $v(c_1) \neq v(c_2)$. Suppose without loss of generality that c_3, c_4, c_5, c_6 are in counterclockwise order, with possibly $c_4 = c_5$; see Figure 13. Let $p_1 = [c_4, c_5]$ along $H(t_m)$. Let $H_1 = [c_5, c_1] \cup c_1c_5$ and $H_2 = [c_2, c_4] \cup c_4c_2$ be two interior disjoint cycles of $D(t_m)$; see Figure 13.

Suppose first that $c \in H_1$ (we can argue analogously if $c \in H_2$). Let $p_2 = H_1 \setminus \{c\}$ be a path. We partition $D(t_m) \setminus \{c\}$ into two even paths and a basic cycle. If p_2 is even, then set $p_1 = p_1 \setminus \{c_3\}$, otherwise set $p_2 = (p_2 \setminus \{c_3\}) \cup c_1c_4$. If p_1 is even, then set $H_2 = H_2 \setminus \{c_5\} \cup c_2c_6$, otherwise set $p_1 = p_1 \setminus \{c_5\}$. We have partitioned $D(t_m) \setminus \{c\}$ into the even paths p_1 and p_2 and basic cycle H_2. Suppose next that $c \in p_1$. Now $p_1 \setminus \{c\}$ is the union of two paths, say p_4 and p_5, such that $c_4 \in p_4$ and $c_5 \in p_5$. As in the above, depending on the parity of p_4 and p_5, we can choose to remove c_4 from p_4 or from H_2, and similarly remove c_5 from p_5 or H_1, obtaining two even paths and two basic cycles.

Case b2.3.2: $v(c_1) = v(c_2)$. This implies that $c_3 = c_5$ and $c_4 = c_6$, and c_1, c_2, c_3, c_4 induce a 4-clique, whose vertices can be guarded from vertex $v(c_1) = v(c_2)$. Denote the 4-clique by q.
Let $H_1 = [c_4, c_1] \cup c_1c_4$ and $H_2 = [c_2, c_3] \cup c_3c_2$ be two cycles of $D(t_m)$; see Figure 14a. Assume that $c \in H_1$ (we can argue analogously if $c \in H_2$). Let $p_1 = H_1 \setminus \{c\}$. If p_1 is even, then p_1 and H_2 is the desired partition of $D(t_m) \setminus \{c\}$. So suppose that p_1 is odd. Notice that $H(t_m) \setminus \{c, c_1, c_2, c_3, c_4\}$ is the union of three paths, two of which are even and the remaining path is is odd. Note that one endpoint of each path is adjacent to a cell in clique q. We can append one cell of q to the odd path, and obtain a partition of $D(t_m) \setminus \{c\}$ into three even paths and a triangle contained in q.

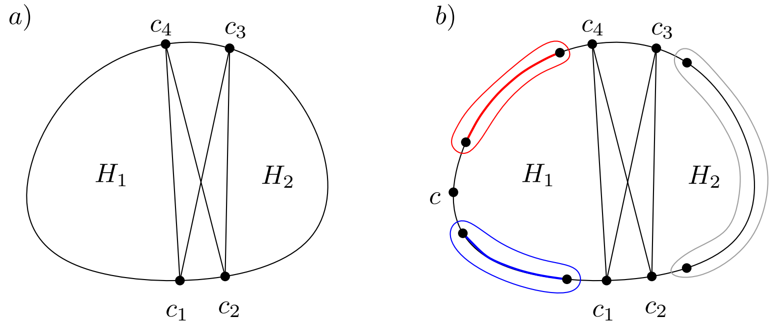

Figure 14: Illustration for **Case b2.3.2.**

For $m = 1, 2, \ldots, k$, we have computed a set $\mathcal{G}_m$ such that $\mathcal{G} = \cup_{m=1}^{k} \mathcal{G}_m$ is a clique cover of $D(C)$. We have $|\mathcal{G}_m| \le \lfloor \frac{n(t_m)}{2} \rfloor$ for every m. Recall that for every component t_m, at most one adjacent cell could be adjacent to another a previous component t_i, $i < m$. It follows that $|\mathcal{G}| \le \lfloor \frac{n+1}{2} \rfloor = \lceil \frac{n}{2} \rceil$. □

6 A Simpler Proof for Vertex Guards

Karavelas, Tóth and Tsigaridas [11] proved that $\lfloor \frac{2n}{3} \rfloor$ vertex guards are always sufficient and sometimes necessary to guard a piece-wise convex curvilinear polygon with $n \ge 2$ vertices. We finish this paper by providing a simpler proof of their result.

Theorem 10 ([11]). *Let A be a piece-wise convex curvilinear art gallery with $n \ge 2$ vertices. Then $\lfloor \frac{2n}{3} \rfloor$ vertex guards are always sufficient and sometimes necessary to guard A.*

Proof. Let A be a curvilinear art gallery with $n \ge 2$ vertices. Label the vertices by $v_0, \ldots, v_{n-1}$ along the boundary of A, addition taken $\mod n$. For any two consecutive vertices v_i and v_{i+1}, let P_i be the shortest path from v_i to v_{i+1} contained in A. Refer to Figure 15. Every path P_i is a simple polygonal chain. Since the boundary of A consists of convex arcs, every vertex of P_i is a vertex of A. Since P_i and the convex arc a_i have the same endpoints, v_i and v_{i+1}, $P_i \cup a_i$ is a simple closed curve. Let R_i denote the simply connected region in the interior of $P_i \cup a_i$. We call R_i the *room* of a_i. Since P_i is a shortest path between v_i and v_{i+1} in A, all internal vertices of P_i are reflex vertices of region R_i.

Since the paths P_i connect consecutive vertices of A, they are pairwise non-crossing, and the rooms R_i are interior disjoint. The paths P_i, $i = 0, 1, \ldots, n-1$, jointly decompose A into simply connected regions, see Figure 15 (right). The regions adjacent to the boundary of A are rooms. We call any other region a *polygonal region*; these are simple polygons bounded by some edges of a path P_i.

Let V be the set of n vertices of A. Consider the decomposition of A into n rooms and possibly some polygonal regions. Triangulate every polygonal region and let E denote the set of edges of all paths P_i, and all edges of the triangulations of the polygonal regions. We define a dual graph T of graph (V, E) as follows. The vertices of T are the triangles in the triangulation of the polygonal regions. Two nodes are adjacent if and only if the corresponding triangles share an edge; that is, if each edge of the dual graph of T corresponds to an edge $e \in E$.

It is not difficult to see that T is a forest. Every edge $e \in E$ decomposes A into two curvilinear art galleries, and so the removal of the dual edge of e disconnects one of the connected components of T. It follows that graph (V, E) has a proper 3-vertex coloring. Fix an arbitrary 3-vertex coloring of (V, E); see Figure 15 (left). The total size of the two smallest color classes is at most $\lfloor \frac{2n}{3} \rfloor$. We show that guards at these vertices jointly monitor the entire art gallery. It is clear that every

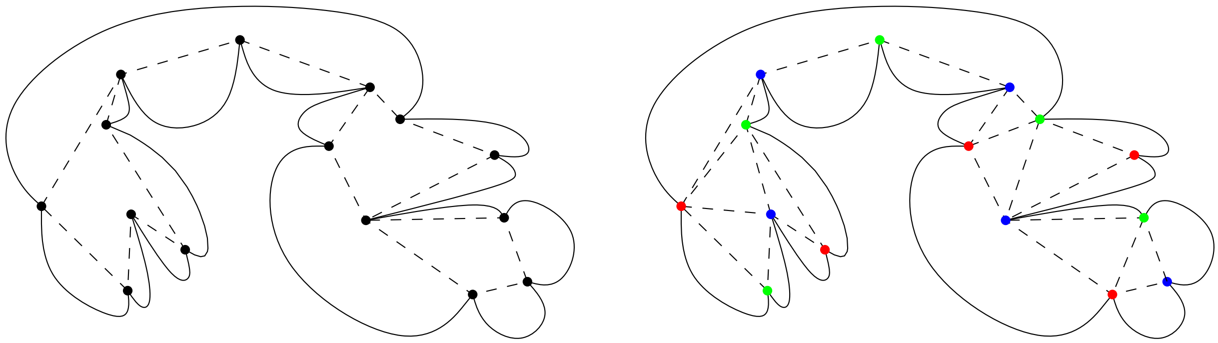

Figure 15: Left: Shortest paths between consecutive vertices of a curvilinear art gallery A. Right: A triangulation of the polygonal regions of A and a 3-coloring of graph (V, E)

triangle in the triangulation of a polygonal region is guarded by vertices in each color class. We show next that every room is guarded by vertices in any two color classes.
We say that a point $p \in A$ *sees all of an edge* $e \in E$ if the triangle spanned by e and p is contained in A. The following claim implies that every point in a room sees both endpoints of some edge in E.

Claim. *Let R_i be a room of of A, and let $p \in R_i$. Then p sees all of some edge e in path P_i.*

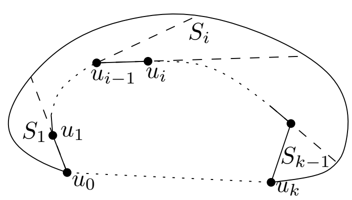

Figure 16: The decomposition of room R_0 into convex cells.

If P_i has exactly one edge e, then the room R_i is convex, and p sees all of e. Suppose that P_i has at least two edges. Suppose that $P_i = (v_i = u_0, u_1, \ldots, u_k = v_{i+1})$. For $j = 1, \ldots, k-1$, extend edge $u_{j-1}u_j$ beyond its endpoint u_j until it hits the convex arc a_i. The extensions decompose R_i into $k-1$ convex cells, each adjacent to a unique edge of P_i. If p lies in the interior of a convex cell, then p sees all of the edge of P_i adjacent to the cell. If p lies on the extension of edge $u_{j-1}u_j$ for some $j = 1, 2, \ldots, k-1$, then p sees all of edge u_ju_{j+1}. This completes the proof of the Claim, and thus the proof of the theorem. □

We conclude by constructing a family of curvilinear art galleries with n vertices, where $n \equiv 0 \mod 3$, that requires at least $\frac{2n}{3}$ vertex guards. A similar construction has been presented in [11]. The construction is based on a pattern formed by three consecutive convex arcs depicted in Figure 17 (left). Let Q be a regular $\frac{n}{3}$-gon, replace every edge of Q by a rotated copy of the three convex arcs as shown in Figure 17 (right). For each triple of consecutive arcs, we can construct three interior-disjoint regions such that each region is seen from only two vertices of the pattern. It now follows that the three regions require at least two vertex guards. Over $\frac{n}{3}$ copies of this pattern, n interior disjoint regions require $\frac{2n}{3}$ vertex guards.

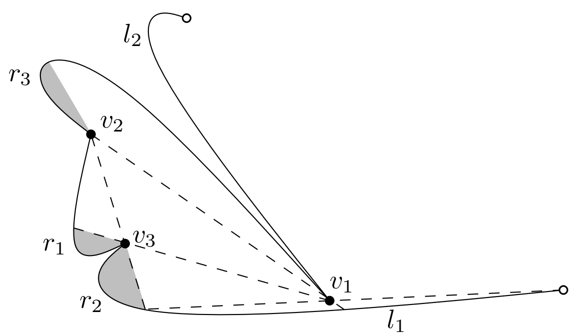

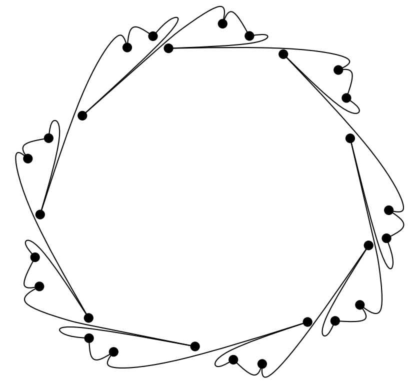

Figure 17: Left: Basic pattern for the lower bound construction. Right: A curvilinear art gallery with 27 vertices that requires 18 vertex guards.

References

[1] M. Al-Jubeh, M. Hoffmann, M. Ishaque, D. L. Souvaine and Cs. D. Tóth, *Convex decompositions with 2-Edge Connected Dual Graphs.* Journal of Combinatorial optimizartion (2010), in print. `DOI: 10.1007/s10878-010-9310-1`

[2] D. Avis and H. El Gindy *A combinatorial approach to polygon similarity*, IEEE Trans. Inform. Theory, IT-2:148–150, 1983.

[3] W. Bronsvoort. *Boundary evaluation and direct display of CSG models.* Computer-Aided Design **20** (1988), 416–419.

[4] J. Cano, J. Espinosa and J. Urrutia *Vigilancia en Galerías de Arte Curvilíneas* Actas de los XIII Encuentros de Geometria Computacional, Zaragoza, España, pp. 59–66, 2009.

[5] B. Chazelle and J. Incerpi. *Triangulation and shape-complexity.* ACM Trans. Graph. **3** (2) (1984), 135–152.

[6] V. Chvátal, *A combinatorial theorem in plane geometry*, J. Combin. Theory Ser B **18** (1975), 39–41.

[7] M. Eo and C. Kyung, *Hybrid show testing scheme for ray tracing.* Computer-Aided Design **21** (1989), 38–48.

[8] D. Eppstein, M. T. Godrich and N. Sitchinava. *Guard placement for efficient point-in-polygon proofs.* In Proc. 23rd Annu. ACM Sympos. Comput. Geom., pp. 27–36, 2007.

[9] S. Fisk, *A short proof of Chvátal's watchman theorem*, J. Combin. Theory Ser B **24** (1978), 374.

[10] M.I. Karavelas *Guarding curvilinear art galleries with edge or mobile guards via 2-dominance of triangulation graphs* Computational Geometry: Theory and Applications **44** (1) (2011), 20–51.

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

[12] R. Kuc and M. Siegel, *Efficient representation of reflecting structures for a sonar navigation model.* In Proc. 1987 IEEE Int. Conf. Robotics and Automation, pp. 1916–1923, 1987.

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

[14] T. Lozano-Pérez and M. A. Wesley. *An algorithm for planning collision-free paths among polyhedral obstacles.* Commun. ACM, **22** (10) (1979), 560–570.

[15] M. McKenna. *Worst-case optimal hidden-surface removal.* ACM Trans. Graphs. **6** (1987). 19–28.

[16] J. S. B. Mitchell. *An algorithmic approach to some problems in terrain navigation.* In D. Kapur and J. Mundy, editors, Geometric Reasoning. MIT Press, Cambridge, MA, 1989.

[17] J. O'Rourke. *Art Gallery Theorems and Algorithms.* The International Series of Monographs on Computer Science. Oxford University Press, New York, NY, 1987.

[18] T. C. Shermer. *Recent results in art galleries.* Proc. IEEE **80** (9) (1982), 1384–1399.

[19] J. Stenstrom and C. Connolly. *Building wire frames for multiple range views.* In Proc. 1986 IEEE Conf. Robotics and Automation, pp. 615–650, 1986.

[20] G. T. Toussaint. *Pattern recognition and geometrical complexity.* In Proc. 5th IEEE Internat. Conf. Pattern Recognition., pp. 1324–1347, 1980.

[21] J. Urrutia, *Art gallery and illumination problems*, In J.-R. Sack and J. Urrutia editors, Handbook of Computational Geometry, pp. 973–1027. North-Holland, 2000.

[22] S. Xie, T. Calvert and B. Bhattacharya. *Planning views for the incremental construction of body models.* In Proc. Int. Conf. Pattern Recognition, pp. 154–157, 1986.

[23] M. Yachida. *3-D data acquisition by multiple views.* In Robotics Research: Third Int. Symp., pp. 11–18, 1986.