

The College of Wooster

Open Works

Senior Independent Study Theses

2019

Exploring Topics of the Art Gallery Problem

Megan Vuich

The College of Wooster, mvuich19@wooster.edu

Follow this and additional works at: <https://openworks.wooster.edu/independentstudy>

Recommended Citation

Vuich, Megan, "Exploring Topics of the Art Gallery Problem" (2019). *Senior Independent Study Theses*. Paper 8534.

This Senior Independent Study Thesis Exemplar is brought to you by Open Works, a service of The College of Wooster Libraries. It has been accepted for inclusion in Senior Independent Study Theses by an authorized administrator of Open Works. For more information, please contact openworks@wooster.edu.

© Copyright 2019 Megan Vuich



EXPLORING TOPICS OF THE ART
GALLERY PROBLEM

INDEPENDENT STUDY THESIS

Presented in Partial Fulfillment of the
Requirements for the Degree Bachelor of Arts in
the Department of Mathematics and Computer
Science at The College of Wooster

by
Megan Vuich

The College of Wooster
2019

Advised by:

Dr. Robert Kelvey

Abstract

Created in the 1970's, the Art Gallery Problem seeks to answer the question of how many security guards are necessary to fully survey the floor plan of any building. These floor plans are modeled by polygons, with guards represented by points inside these shapes. Shortly after the creation of the problem, it was theorized that for guards whose positions were limited to the polygon's vertices, $\lfloor \frac{n}{3} \rfloor$ guards are sufficient to watch any type of polygon, where n is the number of the polygon's vertices. Two proofs accompanied this theorem, drawing from concepts of computational geometry and graph theory.

This paper explains the Art Gallery Problem along with its two most famous proofs. Certain methods of polygon partitioning, which can be found in both proofs, are also discussed. Finally, extensions to the problem involving subsets of polygons and guards, such as mobile guards, orthogonal polygons, and polygons with holes, are briefly examined. The paper concludes with a cursory glance at extensions that have only begun to be considered.

Acknowledgements

I would like to thank my advisor, Professor Kelvey, for his guidance and help throughout the whole process. I also want to thank the Wooster Math department, which is full of some of the most supportive teachers I have ever come across. It was a joy to be taught by them.

I am so very grateful to my parents for their constant encouragement, and to my favorite brother Sam, who doesn't particularly like math but promised to read my IS if I mentioned him in the acknowledgements (a promise is a promise, Sam. All 72 pages).

Finally, I would like to thank my friends, who never complained about how much I stressed over IS. My time at Wooster wouldn't have been the same without you.

Contents

Abstract	iii
Acknowledgements	v
1 Defining the Problem	1
1.1 The Art Gallery Problem	1
1.2 Polygons	2
1.3 Introduction to Graph Theory	4
1.4 Guard's sight	6
2 The Proofs by Chvátal and Fisk	9
2.1 The Watchman Theorem	9
2.2 Necessity and Sufficiency	10
2.3 Chvátal's Proof	11
2.4 Graph Coloring and Fisk's Proof	20
3 Polygon Triangulation	23
3.1 Triangulation Existence Proof	23

3.2	Polygon Regularization	27
3.3	Triangulation of Monotone Polygons	36
4	Extensions of the Problem	43
4.1	Polygons with Holes	43
4.2	Mobile Guards	49
4.3	Orthogonal Polygons	52
4.4	Considering Multiple Extensions	58
5	Stranger Extensions and Further Work	63
5.1	Adding Cameras to Galleries	63
5.2	Mirrors	65
5.3	Conclusion	72

List of Figures

1.1	Parts of a Polygon	4
1.2	Representing the Floor Plan with a Graph	6
1.3	Point and Vertex Guards	7
2.1	Toussaint's Necessity Polygon	11
2.2	n -Triangulations of a Polygon	12
2.3	Examples of Fans	13
2.4	Cutting Polygons	14
2.5	Diagonal d in an Arbitrary Polygon	15
2.6	Fans of a Polygon	16
2.7	Single-Fanned Triangulations	16
2.8	$K = 4$	17
2.9	$K = 5$	18
2.10	$K = 6$	19
2.11	Graph Coloring	21
3.1	Convex Vertices	24

3.2	Internal or External Diagonals	25
3.3	Partitioning from a Non-convex Vertex	25
3.4	Monotone Polygons	28
3.5	The Presence of Interior Cusps	29
3.6	Regularization: Ordering and Checking Vertices	32
3.7	Regularization: Remaining Iterations	35
3.8	Stack Operations	38
3.9	Triangulation: An Example	40
3.10	Triangulation: Another Example	41
4.1	Holes in a Polygon	44
4.2	Holes Sharing Vertices or Edges	45
4.3	Euler's Theorem with Triangulations	47
4.4	Necessity Polygon with Holes	48
4.5	Two Types of Mobile Guards	49
4.6	Closed-Edge Guards on Edges or Diagonals	51
4.7	Examples of Orthogonal Polygons	52
4.8	An Orthogonal Necessity Polygon	53
4.9	Convex Quadrilateralization	54
4.10	An Uncooperative Four-Coloring	56
4.11	Coloring a Quadrilateralization	57
4.12	Open-Edge Mobile Guards on Orthogonal Polygons	59
4.13	Closed-Edge Mobile Guards on Orthogonal Polygons	60
4.14	Orthogonal Polygons with Holes	61
5.1	Toussaint's Polygon with Cameras	64

5.2	Cameras with Limited Range	65
5.3	Mirrors in Galleries	66
5.4	Toussaint's Polygon with a Mirror	67
5.5	An Example of the View of Vertex Guards	68
5.6	The View of Open-Edge Mobile Guards	69
5.7	View of Vertex Guards with Mirrors	70
5.8	Considering Edge Lengths with Mirrors	71

Chapter 1

Defining the Problem

1.1 The Art Gallery Problem

Say you are approached by the curator of your local art gallery. He has an assignment for you: the gallery has recently been broken into, so it is your job to increase security. You must hire enough guards to station in the gallery so that every square foot of the floor plan can be watched by a guard. Of course, the gallery wants to minimize costs, so the less guards you need to hire, the better. What is the least number of guards you will need to guard the entire gallery, and where should they be stationed?

Depending on the details, this problem can become quite complicated. What does the floor plan look like? Are there pillars or large sculptures that can block the guards' view? Must the guards remain stationary, or are they allowed to patrol? Is the gallery on multiple floors, and how are these floors connected? Are there any any mirrors on the walls? Before we become overwhelmed, let's consider this problem in its simplest form: The Art Gallery

Problem, as originally posed by Victor Klee [11]:

Proposition 1. *Consider the floor plan of any art gallery. If guards can be stationed anywhere in the gallery at fixed posts, but with the ability to turn around, how many guards are needed for every point in the gallery to be watched by a guard?*

This problem first made its appearance in 1973 at a conference at Stanford University. Here, Vasek Chvátal, a mathematics professor, asked Victor Klee, another mathematician, for a challenging geometric problem. In response, Klee proposed the problem of finding the minimum number of guards sufficient to watch the interior of an art gallery room with n walls. Klee required that every guard must remain at a fixed post, the walls of the gallery must be straight, and every inch of wall space must be watched by at least one guard [8]. Chvátal would give his formal answer to the problem in 1975, in the form of his Watchman Theorem. In the following decades, many people would enjoy expanding this problem; some added more conditions to the guard's abilities, while others specified the shape of the rooms in question. Several new proofs of the original theorem were produced using areas of math outside of geometry. Even computer scientists spent time developing the problem as a way of studying partition algorithms [11]. It is no surprise that countless extensions of this problem are still being explored today.

1.2 Polygons

It can be reasonable to suggest that any floor plan can be modeled as a polygon.

It is usual practice to define polygons as a collection of vertices connected to each other by line segments, all lying in the Euclidean plane, \mathbb{R}^2 . However, in consideration of this problem, we are also concerned with points inside the polygon, and not just its outline. Therefore, we will define a polygon as the closed region of a plane, bounded by the vertices and edges of the traditional definition. V is considered to be the set of all vertices in polygon P , so that $V = \{v_1, v_2, \dots, v_n\}$. Likewise, E is the set of all edges found in P , where $E = \{v_1v_2, v_2v_3, \dots, v_{n-1}v_n, v_nv_1\}$. This collection of vertices and edges will be referred to as the boundary of polygon P , or δP . In set notation, this is:

$$\delta P = E \cup V$$

For a *simple polygon*, none of the line segments of E will cross over each other. It will be assumed that all polygons discussed in the following chapters will be simple (See fig. 1.1).

The boundary δP divides \mathbb{R}^2 into two regions: an unbounded region called the *exterior* of the polygon, and the bounded region called the *interior* of the polygon [11]. This interior will be the second part of our polygon definition.

Definition 1. A **polygon P** with n sides is the boundary of P , δP , and all points that lie inside of this boundary, also known as $Int(P)$. Hence, $P = Int(P) \cup \delta P$.

It is quite possible that a floor plan may need to be represented by multiple polygons. This could be necessary if the gallery consists of multiple stories of a building, or is contained in non-adjoining rooms (or adjoined rooms that are separated from closed and opaque doors, where guards cannot patrol between rooms).

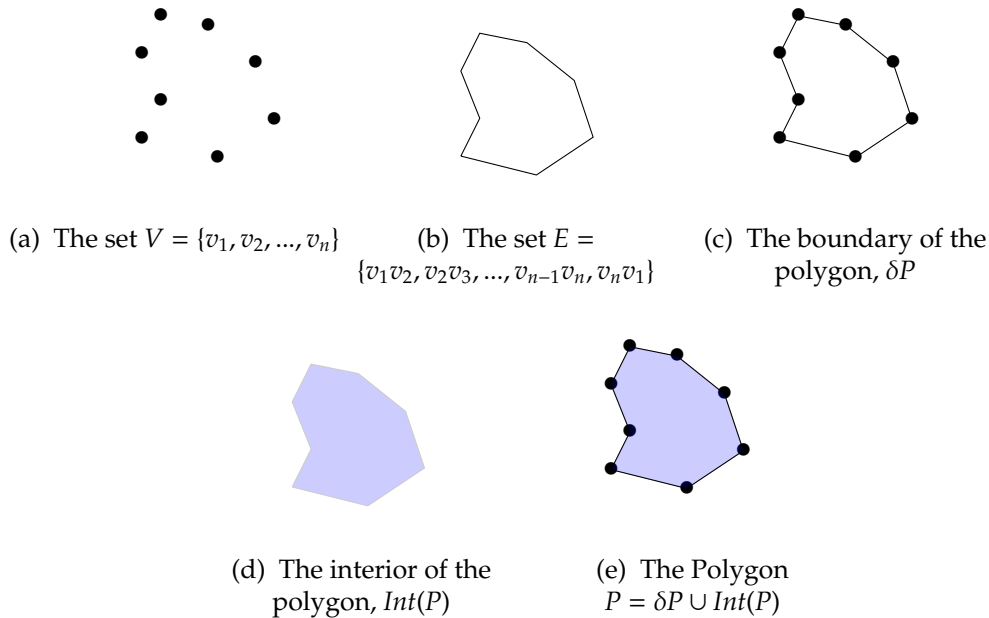


Figure 1.1: Parts of a Polygon

A Polygon as defined for use in the Art Gallery Problem; usually the interior is not included in this definition.

As one more requirement for the simple polygons that we will consider, we do not want more than two adjacent vertices of the polygon to be colinear, or lying on the same edge. Having nonadjacent colinear vertices would make polygon partitioning a more difficult process, and partitioning is essential to proofs that will be mentioned later.

1.3 Introduction to Graph Theory

Instead of using polygons, we can choose to describe art gallery floor plans as graphs.

A *graph* is a mathematical model that displays relationships between objects. In this case, our graphs would show points in the Euclidean plane (or

floor plan), with two points being related if a wall goes from one point to the other point.

Definition 2 ([13]). A **graph** is a set of *nodes* or *vertices*, called $V = \{v_1, v_2, v_3, \dots, v_n\}$, and a set of edges $E = \{e_1, e_2, \dots, e_n\}$, where each edge is a line segment connecting two nodes in V . Two vertices that share an edge are said to be **adjacent**.

Art gallery graphs will always be simple graphs, meaning that any two vertices will only be connected by at most one edge, and no vertex is connected to itself by an edge (this is because we are assuming that all walls in the gallery are straight). These graphs will also be *connected graphs*, meaning that by traveling on the edges of the graph, any one vertex can be reached by any other vertex. To compare this with polygonal representations of floor plans, a graph is equivalent to the boundary δP .

Different graph terminology can correspond with different types of galleries. The most simple styles of one-room floor plans, where there are no pillars, sculptures, or any structures besides walls that can block one's view, can be represented by cycles. A *cycle* is a type of connected graph where each node is connected to exactly two other nodes, so that you can travel from any node back to itself by following edges, in a path that takes you to every other node in the graph.

Finally, the regions that are bounded by a graph's vertices and edges are called the *faces* of the graph. For simple floor plans, as considered in the original Art Gallery Theorem, each graph will only have two faces: the unbounded region outside the cycle, and the bounded region inside. These




Figure 1.2: Representing the Floor Plan with a Graph

Graphs can also be used to represent gallery floor plans, with nodes and edges being equivalent to the vertices and edges of their polygonal counterparts. The face found inside the graph (Face 1) is also identical to the interior of its corresponding polygon.

faces are comparable to a polygon's exterior and interior (See fig. 1.2).

1.4 Guard's sight

Whether you use polygons or graphs to describe your gallery, the mathematical interpretations of their guards are identical. A guard can be shown as a point placed anywhere in the polygon or graph, including on the boundary of P , or on an edge or vertex of the graph G . These are assumed to be standing against the gallery's wall.

A guard's *sight* is any part of the floor plan that the guard can see from their position, allowing for the guard to turn in place. The extent of the guard's vision and the height of the walls is not taken into account; as long as there are no walls blocking their view, these guards' lines of sight go on indefinitely.

Definition 3. A guard g 's sight S is the set of all points $s \in P$ such that the line segment between g and s lies completely in P .

The set S is therefore a subset of P , and the art gallery is completely guarded when the union of every guard's sight is equal to P .

The guards can be categorized by where they are located in the gallery. Generally, we consider three types of guards: *vertex guards*, *edge guards*, and *point guards*. Their names are self-explanatory; vertex guards can only be placed at the vertices of the polygon (or nodes of the graph), edge guards can be placed at any point of δP , and point guards can be placed anywhere in the polygon, in the interior as well as on the boundary.

Much of the work done on the Art Gallery Problem has only focused on vertex and edge guards, as there are less potential points of placement to consider. For example, the main theorem of the problem, identifying how many guards are sufficient to guard any general polygon, specifies that we are only considering vertex guards. You can see that vertex guards are types of edge guards, since the vertices of a polygon are elements of δP , and edge guards are types of point guards, as $\delta P \subset P$.

In some situations, it will take more vertex guards than general point guards to completely watch a polygon, as shown in Figure 1.3.




Figure 1.3: Point and Vertex Guards

This polygon only requires a single point guard to be fully covered, but more than one guard is needed if only vertex guards can be used.

A polygon is *convex* if for any two points contained in the polygon, the line

segment between these points lies completely in the polygon. From this definition, it quickly follows that any convex polygon needs only one guard to be thoroughly watched.

Theorem 1 ([1]). *One guard is sufficient to guard any convex polygon P , and this guard can be placed at any point in P .*

Proof. Let P be a convex polygon. By definition of convexity, any two points in P have a line segment between them that lies entirely in P . Therefore, by placing a guard g at any point in P , every point in P is visible to g , by definition of g 's sight. Hence, one guard is always sufficient to watch convex polygons, when placed anywhere in the polygon. \square

We will show in a later chapter that any non-convex polygon can be partitioned into several convex shapes: specifically, triangles.

Conclusion

We now have all the basic knowledge required to dive into some theorems of the Art Gallery Problem, starting with the first and arguably most important, the Watchman Theorem. This provides a number for the guards that is sufficient to fully watch any simple polygon. This theorem is accompanied by two proofs, which both rely on concepts of polygons and graphs that we have reviewed in this chapter. Continuing further, we will explore how this theorem came to be, and how topics we have already covered paved the way for us to make such an all-encompassing statement about any floor plan.

Chapter 2

The Proofs by Chvátal and Fisk

As mentioned in the previous chapter, the Art Gallery Problem took shape when mathematician Vasek Chvátal requested a geometric problem from his colleague Victor Klee in 1973. It did not take Chvátal long to provide an answer to the problem, complete with a proof; his combinatorial theorem on the subject was published in 1975. A few years later, another mathematician named Steven Fisk outlined his own proof of Chvátal's theorem, so succinct that when it was published it only took up a single journal page [11].

Interestingly enough, these two proofs, both elegant in their own ways, rely on different areas of mathematics to prove the same theorem. In this chapter, the theorem to the Art Gallery Problem, in the form of Chvátal's Watchman Theorem, is presented, along with these two earliest proofs.

2.1 The Watchman Theorem

First, let us look at the Watchman Theorem that Chvátal produced in 1975.

Theorem 2. [Chvátal's Watchman Theorem [11]] $\lfloor \frac{n}{3} \rfloor$ vertex guards are occasionally necessary and always sufficient to watch an entire polygon with n edges.

It is important to note that Chvátal, and later Fisk, only considered vertex guards in this theorem; allowing for point guards or edge guards can sometimes result in the need for less guards (see Figure 1.3). Both mathematicians also implied that this theorem holds true for all simple polygons, but are in no way saying that subsets of simple polygons will require this same amount (subsets of simple polygons that require less than $\lfloor \frac{n}{3} \rfloor$ guards will be discussed in later chapters).

2.2 Necessity and Sufficiency

Two key phrases from Chvátal's Theorem are "occasionally necessary" and "always sufficient". We will find these phrases in most theorems relating to the Art Gallery Problem.

Definition 4. An amount of guards in an art gallery is **necessary** if the gallery cannot be fully watched with one less guard.

Definition 5. An amount of guards is **sufficient** in an art gallery for a certain type of floor plan if all floor plans of this type can be fully watched by this number of guards.

Generally, we only need one example of a floor plan that *needs* a certain number of guards, while it must be proven that any arbitrary floor plan that fits the description of the theorem can be sufficiently watched with this

number of guards. The necessity example acts as a lower bound for the number of guards; for example, consider the most common floor plan used to show that $\lfloor \frac{n}{3} \rfloor$ is the least number of guards needed for any simple polygon, Toussaint's Necessity Polygon (fig. 2.1).




Figure 2.1: Toussaint's Necessity Polygon

This polygon needs $\frac{n}{3}$ guards to be fully watched. This will hold even if more prongs are added to this polygon, as you would have to add three more sides to get another prong.

This comb-shaped polygon has nine walls and three prongs, and each prong must be watched by its own guard. One of the guards must also watch the base of the polygon, and so must be placed at either point a or b . Even so, a guard can only fully see one prong from either of these points. Therefore, this polygon requires $\frac{n}{3}$ guards [1]. If we were to try to prove that less guards are needed for simple polygons, this floor plan would become a counterexample, and our proposed theorem would not hold.

2.3 Chvátal's Proof

Chvátal first assumed that any polygon can be partitioned into triangles, or *triangulated*. To do this, we would add internal diagonals between the vertices of the polygon, until no more can be added without diagonals or edges

crossing. This would result in a *triangulation* of the polygon.

Definition 6. An n -**triangulation** of an n -sided polygon P is a graph G with n vertices such that one of G 's faces is bound by P , and each of the remaining faces is bound by a triangle in the interior of P [4].




Figure 2.2: n -Triangulations of a Polygon

A polygon and two of its corresponding triangulations. Note that triangulations are not always unique.

Why Chvátal's assumption is correct, and how we could go about triangulating any polygon, will be discussed in the next chapter.

Chvátal was not just interested in any triangulation, however; he was specifically interested in fans.

Definition 7. A **fan** is a collection of triangles that share a vertex, called the *center* of the fan.

In the previous chapter it was proven that a guard placed at any point in a convex n -sided polygon will cause the polygon to be fully guarded (theorem 1). Because the center of a fan acts as a vertex for all triangles included in the fan, placing a guard at the center of the fan will cause the fan to be fully watched.




Figure 2.3: Examples of Fans

Chvátal wanted to show that any triangulation of a polygon can be partitioned into $\lfloor \frac{n}{3} \rfloor$ fans. It would then follow that placing a guard at the center of each fan in the triangulation allows all the triangles to be watched, and therefore the whole polygon is watched by $\lfloor \frac{n}{3} \rfloor$ guards [11].

Throughout the proofs of the following lemma and Chvátal's theorem, we will be considering an arbitrary polygon P with $n \geq 6$ vertices, that is triangulated into graph G . P and G will have their vertices labelled $0, 1, \dots, n - 1$, in such a way that vertex i is always adjacent to the vertices $i - 1$ and $i + 1$, and the vertex labelled 0 is adjacent to 1 and $n - 1$. One fan will be removed from G , creating two triangulations, where G' is the remaining portion of G , and G_f is the removed fan. *Removing* a fan means that we "cut" G at one of its diagonals, creating two distinct graphs that share at least one edge and two vertices (see fig. 2.4).

It will not do to simply remove a triangle from G ; two of the triangle's vertices will still be a part of G' . This means that two out of three times, the fan center that G_f was originally a part of will also be a vertex in G' , meaning that G' still has $\lfloor \frac{n}{3} \rfloor$ fans. G_f is now considered its own fan, so combining these two graphs means that G has $\lfloor \frac{n}{3} \rfloor + 1$ fans, which is more fans than we need.




Figure 2.4: Cutting Polygons

When a polygon P is “cut” along one of its diagonals, we now recognize two triangulations P_1, P_2 that both contain the vertices v_1, v_2 and the edge between them. Though P_2 has five vertices, we would say that we removed three vertices from P_1 , as v_i and v_j are still vertices of P_1 .

We want to remove one fan exactly, so at least three vertices must be removed from G . Therefore, we must find a diagonal in G that cuts off a minimum of four edges, or three vertices, so that G' has $n - 3$ vertices or less. This results in a triangle being completely separate from G' . Because every triangle was a part of a fan in G , one of the removed triangle’s vertices would be the center of a fan. Removing the center from G' means it now has $\lfloor \frac{n}{3} \rfloor - 1$ fans. On the other hand, we only want to remove one fan from G , so no more than five vertices should be cut off [11]. We will now prove that every polygon has a diagonal that allows us to do this.

Lemma 1. *For any n -triangulation with $n \geq 6$, there exists a diagonal d that cuts off exactly four, five, or six edges, or three, four, or five vertices, respectively.*

Proof. Let d be a diagonal of the triangulation G that cuts off a minimum number of edges, so long as d cuts off at least four edges. Let k be this

minimum number, so that $k \geq 4$. Also, let d 's endpoints be at 0 and k .

Because G is a triangulation, d is a part of the triangle $(0, t, k)$, where t is some vertex such that $0 < t < k$ (see fig. 2.5). Then, the diagonal $(0, t)$ cuts off t edges from G . By the minimality of d , no diagonal exists that cuts off 4, 5, ..., $k - 1$ edges. Therefore, because $t < k$, it must follow that $t \leq 3$. Likewise, the diagonal (t, k) cuts off $k - t$ edges, so $k - t \leq 3$ by similar logic. It follows from these inequalities that $k \leq 6$.

Hence, there will always exist some diagonal d which cuts off exactly four, five, or six edges of G , when G has $n \geq 6$ vertices.




Figure 2.5: Diagonal d in an Arbitrary Polygon

In any triangulation, there will exist a diagonal d that cuts off four, five, or six edges (adapted from [11]).

□

Theorem 3. ([8]) *A triangulated polygon with n sides can be partitioned into m fans, where $m \leq \lfloor \frac{n}{3} \rfloor$.*

Proof. First, we will establish base cases for polygons with $n = 3, 4, 5$ sides, since $n \geq 3$ for all polygons. If vertex labels are not considered, there is only one way to triangulate three, four, or five-sided polygons, and these




Figure 2.6: Fans of a Polygon

This twelve-sided polygon can be partitioned into four fans, meaning that four guards are sufficient for this particular polygon.

triangulations are fans themselves [8]. The theorem therefore holds for these polygons, as the number of fans in these polygons, m , is:

$$\left\lfloor \frac{n}{3} \right\rfloor = \left\lfloor \frac{3}{3} \right\rfloor = \left\lfloor \frac{4}{3} \right\rfloor = \left\lfloor \frac{5}{3} \right\rfloor = 1.$$




Figure 2.7: Single-Fanned Triangulations

Polygons with three, four, or five vertices and sides can always be triangulated into a single fan, whether they are convex or not.

Let us now assume an inductive hypothesis: any triangulation with $n - 3$ or less vertices can be partitioned into $\left\lfloor \frac{n-3}{3} \right\rfloor = \left\lfloor \frac{n}{3} \right\rfloor - 1$ fans, given that $n \geq 6$.

Given the previous lemma, we know that there is a diagonal in G that can partition G into G' and G_f , where G_f is a polygon with either five, six, or seven vertices, with $k = 4, 5$, or 6 being the number of edges that are now no longer a

part of G' . We will consider the cases for each value of k :

Case 1. $k = 4$. G_f would therefore have five vertices. As shown in the base cases, this polygon is always able to be triangulated into a single fan (See fig. 2.8). Hence, G can be partitioned into $\lfloor \frac{n}{3} \rfloor - 1 + 1 = \lfloor \frac{n}{3} \rfloor$ fans [11].




Figure 2.8: $K = 4$

The polygon G_f can always be triangulated into one fan, no matter where the center of the fan is located.

Case 2. $k = 5$. G_f would then have six vertices. This only allows for $t = 2$ or 3 , since $k - t \leq 3$ and $t \leq 3$. We can assume without loss of generality that $t = 2$. We already know that the diagonals $(0, t)$ and (t, k) exist in G_f , and we now have the the 4-vertex polygon, formed by vertices $t = 2, 3, 4$, and $k = 5$, to triangulate. There are two options remaining (See fig. 2.9):

Case 2.1. Let the diagonal $(3, 5)$ be present in G_f . Then, G_f would have two fans, as there are no vertices in G_f that are found in all triangles present. Let T be the triangle that contains the vertices $0, t = 2$, and $k = 5$, and consider the triangulation G_T to be the union of G' and T . G_T has $n - 3$ vertices, so by our hypothesis will have $\lfloor \frac{n-3}{3} \rfloor = \lfloor \frac{n}{3} \rfloor - 1$ fans. Because one of the fans of G_T contains T , there must exist a fan center at one of the vertices of T , either $0, 2$, or 5 . If the fan is centered at vertex 2 , then it is only possible for T to be the

only triangle in the fan, and therefore, this fan can be centered at any of T 's vertices, so without loss of generality assume it is at 0 or 5.

Case 2.1.1. Let the fan center be at vertex 0. We could then add the triangle $(0, 1, 2)$ to this fan, and G_T will still have $\lfloor \frac{n}{3} \rfloor - 1$ fans. The remaining four-vertex polygon is triangulated into a single fan, so G has $\lfloor \frac{n}{3} \rfloor$ fans in total.

Case 2.1.2. Let the fan center be at vertex 5. We could extend the fan to include the two triangles $(3, 4, 5)$ and $(2, 3, 5)$, leaving the triangle $(0, 1, 2)$ as its own fan. The number of fans in G would therefore total $\lfloor \frac{n}{3} \rfloor$.

Case 2.2. Let the diagonal $(2, 4)$ be present in G_f . Then, G_f can be considered a single fan with fan center at $t = 2$, since every triangle in G_f contains the vertex $t = 2$. Therefore, G can be partitioned into $\lfloor \frac{n}{3} \rfloor - 1 + 1 = \lfloor \frac{n}{3} \rfloor$ fans.




Figure 2.9: $K = 5$

No matter where a center of a fan is placed, G_f can be partitioned into two fans or less, and if G_f has two fans, at least one can be considered as part of G' ([11]).

Case 3. $k = 6$, and G_f would have seven vertices. Because of the minimality of d , it is only possible for $t = 3$, otherwise there would exist a diagonal in G that cuts off four vertices and is not d . G_f is therefore composed of the triangle $(0, t = 3, k = 6)$, and two four-sided polygons $(0, 1, 2, t)$ and $(t, 4, 5, k)$.

Without loss of generality, after triangulation two of the vertices of the triangle $(0, t, k)$ will be fan centers, and $T = (0, t, k)$ will be in one of these fans. It can be seen that at least one of these fans has a center at 0 or k , and can therefore be grouped into G' , as 0 and k are also vertices in G' . Four cases stem from this scenario, as there are two four-sided polygons with two ways of triangulating each of them.




Figure 2.10: $K = 6$

Given that the triangle T must exist in this manner, no matter how the 7-gon is triangulated, it can be partitioned into two or less fans. Because G' has $\lfloor \frac{n}{3} \rfloor - 2$ fans, this results in the entire polygon having $m \leq \lfloor \frac{n}{3} \rfloor$ or less fans ([11]).

Case 3.1. The diagonals $(0, 2)$ and $(4, k)$ exist. Then, two fans can be made with centers at 0 and k . Without loss of generality, assume T is in the fan centered at k , and this fan can be added to G' . Hence, G' now has $\lfloor \frac{n}{3} \rfloor - 1$ fans, and adding the remaining fan gives the entire polygon $\lfloor \frac{n}{3} \rfloor$ fans.

Case 3.2. The diagonals $(0, 2)$ and $(t, 5)$ exist in G_f . Therefore, two

fans can be made from G_f with centers at 0 and t , and the fan centered at 0 can be added to G' . Then, the total number of fans in the polygon are the ones in G' , $\lfloor \frac{n}{3} \rfloor - 1$, plus the remaining fan at t .

Case 3.3. The diagonals $(1, t)$ and $(4, k)$ exist in G_f . Then, two fans can be formed from G_f , centered at t and k . The fan centered at k can be included in G' , giving G' $\lfloor \frac{n}{3} \rfloor - 1$ fans. With the remaining fan at t , the entire polygon P has $\lfloor \frac{n}{3} \rfloor$ fans.

Case 3.4. The diagonals $(1, t)$ and $(t, 5)$ exist in G_f . Therefore, G_f can be considered a single fan centered at t . Since G' has at most $\lfloor \frac{n}{3} \rfloor - 1$ fans by the inductive hypothesis, the largest possible number of fans P can contain is $\lfloor \frac{n}{3} \rfloor$.

Therefore, given our inductive hypothesis and base cases, any triangulation with n vertices can be partitioned into $\lfloor \frac{n}{3} \rfloor$ fans. If a guard is placed at each fan center in the triangulation, then it follows that the corresponding polygon can be fully watched with $\lfloor \frac{n}{3} \rfloor$ vertex guards. \square

2.4 Graph Coloring and Fisk's Proof

When Steve Fisk developed his proof of the Watchman Theorem a few years later, he used a concept called graph coloring, or specifically, vertex coloring. *Graph coloring*, or *vertex coloring*, is the process of assigning colors to the vertices of a graph so that no adjacent vertices share the same color. These colorings are not always unique, as shown in fig. 2.11.




Figure 2.11: Graph Coloring

This graph is 3-colorable, since it is composed of K_3 graphs. Placing a guard at every vertex colored with the least common color will result in the polygon being guarded by at most $\lfloor \frac{n}{3} \rfloor$ guards.

Definition 8. A graph G is k -colorable if it requires at most k colors for its vertices to be properly colored. The *chromatic number* of G is the smallest value of k such that G is k -colorable [12].

Fisk also includes the use of triangulation in his proof. If we consider the graph of a triangle, this graph is said to be a *complete graph* on three vertices, as any vertex in this graph is adjacent to all other vertices of the graph. Generally, a complete graph on n vertices is labeled as K_n , so a triangle can be labeled as K_3 .

Proposition 2. A complete graph K_n is n -colorable.

Since every two vertices in a complete graph are adjacent to each other, no two vertices in a complete graph can share the same color, meaning that the graph must be colored with a number of colors equal to the number of its vertices. It is therefore assumed that a triangulation, which is composed of multiple K_3 graphs, has a coloring with three colors as well [5].

Proof. (Fisk's Watchman Proof, [5]) Let P be a simple polygon with n vertices, and G one of its corresponding triangulations. Let G be colored so that it only

uses colors 1, 2, and 3. The number of vertices with color 1, color 2, and color 3 will add up to n , by G being 3-colorable. Partitioning n into three groups will cause one of two things to happen; either each group has $\frac{n}{3}$ items, or at least one group will have greater than $\frac{n}{3}$ items, leading to one group having less than $\frac{n}{3}$ items. Therefore, at least one color will appear less than or equal to $\lfloor \frac{n}{3} \rfloor$ times in the coloring (we can consider the floor of this number since this number will always be an integer). Without loss of generality, let color 1 be the color for which this is the case, where m is the number of vertices colored color 1. It is known that every triangle in G will have a vertex with color 1, color 2, and color 3. Then, if a guard is placed at each vertex that has color 1, the polygon can be fully watched with $m \leq \lfloor \frac{n}{3} \rfloor$ guards. \square

Conclusion

In this chapter we have looked at the first theorem of the Art Gallery Problem, along with the two most well-known methods of proving it. Chvátal, the same mathematician who theorized the Watchman claim (theorem 2), showed that any polygon can be partitioned into fans of triangles, and any polygon with n vertices will have at most $\lfloor \frac{n}{3} \rfloor$ fans, with one guard for each fan. Fisk conjectured that a triangulated polygon is three-colorable, and placing guards at the vertices of the least-occurring color results in $\lfloor \frac{n}{3} \rfloor$ guards as well. Both of these proofs infer that any polygon can be transformed into a triangulation, but is this always true? Our next step is to explore this question, and if it is found to be true, to determine efficient ways of performing triangulation on any polygon.

Chapter 3


Polygon Triangulation

In previous chapters, it was mentioned that both Chavátal's and Fisk's Watchman proofs assumed that any simple polygon can be triangulated. In this chapter, it will not only be shown why this assumption is correct, but several algorithms that can turn any simple polygon into a triangulation will also be presented.

3.1 Triangulation Existence Proof

First, we will show that it is possible to find a triangulation of any simple polygon P . This proof relies on the presence of convex vertices in the polygon; luckily, every polygon has at least one convex vertex.

Definition 9. The vertex v_c in the polygon P is a *convex vertex* if the interior angle formed by v_{c-1} , v_c , and v_{c+1} is 180° or smaller.



(a) An example of a polygon with only convex vertices. It happens that all convex polygons will have only convex vertices.

(b) This polygon has one vertex that is not convex, v .

Figure 3.1: Convex Vertices

Theorem 4. Any n -sided polygon can be triangulated into $n - 2$ triangles by adding $n - 3$ diagonals to the polygon.

Proof. Consider a polygon with vertices $n = 3$. This is a triangle already, hence there is no triangulation required.

Now consider the inductive hypothesis that any polygon with $n - 1$ vertices or less is able to be triangulated into $n - 3$ triangles, given that $n \geq 4$.

Let's say we have a simple polygon P with n vertices. By definition, there must exist at least one convex vertex in this polygon. Label one of these vertices v_2 , with its two neighbors v_1 and v_3 . Now, consider adding a diagonal d whose endpoints are v_1 and v_3 (See fig. 3.2).

Case 1: d is completely internal. The diagonal d lies entirely in the interior of P , so this diagonal can be added to the polygon.

The polygon is now separated into two smaller polygons, P_1 and P_2 , with P_2 being the triangle with vertices v_1, v_2, v_3 . P_2 is already triangulated, and P_1 now has $n - 1$ vertices, since it no longer contains v_2 . Therefore, by the inductive hypothesis, P_1 can be triangulated into $n - 3$ triangles. Considering




Figure 3.2: Internal or External Diagonals

(a) The diagonal d is internal, as it lies completely in the interior of polygon P . (b) For this polygon, v_2 is not convex, making d external. A different diagonal will therefore be used to partition P .

P_1 and P_2 together, it can be concluded that P can also be triangulated into $n - 3 + 1 = n - 2$ triangles.

Case 2: d is not internal. Then, consider the vertex x that is not adjacent to v_2 , but lies the closest distance to v_2 (where x is not the same vertex as v_1 or v_3). Draw a diagonal from v_2 to x .




Figure 3.3: Partitioning from a Non-convex Vertex

Since v_2 is not convex, the diagonal that partitions P into two smaller polygons must be drawn from the closest non-adjacent vertex to v_2 , which is x .

The polygon P is now separated into smaller polygons P_1 and P_2 ,

assuming there are no holes in P . Without loss of generality, assume that v_1 is a vertex of P_1 only, and therefore v_3 is a vertex of P_2 only. Therefore, the number of vertices in P_1, P_2 is less than or equal to $n - 1$, so by the inductive hypothesis they can both be triangulated. Hence, P can also be triangulated.

□

Now that it has been shown that any polygon can be triangulated, let's explore a method that can transform any simple polygon, P , into a triangulation. To better reference the individual vertices of the polygon and their proximity to each other, we will place P in the xy -coordinate plane. Then, every vertex v can be defined as a point in \mathbb{R}^2 :

$$v = (x_v, y_v).$$

Let us also define a function $h : V \rightarrow \mathbb{R}$, that isolates the y -coordinate, also referred to as the height, of a vertex of P ,

$$h(v) = y_v.$$

Doing this allows us to order the vertices by the value of their respective y -coordinates; for any $v_j \in V$, there exists $v_i, v_k \in V$ such that $h(v_i) \leq h(v_j) \leq h(v_k)$. Now, every polygon will have at least one maximum vertex, and at least one minimum vertex.

Definition 10. We call a vertex v_i in a polygon P **maximal** when $h(v_i) \geq h(v)$ for all $v \in V$.

Definition 11. We call a vertex v_j in a polygon P **minimal** when $h(v_j) \leq h(v)$ for all $v \in V$.

A method of triangulation that we will explore will require us to first transform a simple polygon into one or more *monotone polygons*, and then convert these monotone polygons into triangulations.

Now that we know it is possible to divide polygons into triangles, we can use several different algorithms that allow us to complete triangulation efficiently, or to even allow computers to do the work for us.

Definition 12. A polygon is *monotone* if, when placed on the x-y plane, the boundary of the polygon can be described as two chains of vertices and edges, such that the vertices of each chain can be labelled in such a way that the y-values of the vertices along each chain are strictly increasing or strictly decreasing [9].

3.2 Polygon Regularization

The process of converting a simple polygon into a monotone polygon is sometimes referred to as *polygon regularization*. Through this method, diagonals are added to the polygon to partition it into several smaller polygons, so that in the context of each polygon, each vertex besides the maximum and minimum vertices will be *regular*.

Definition 13. A vertex v of a polygon P is said to be a *regular vertex* if:

1. v is adjacent to a vertex $u \in P$ such that $h(v) < h(u)$, given that v is not the maximum vertex of P , and,

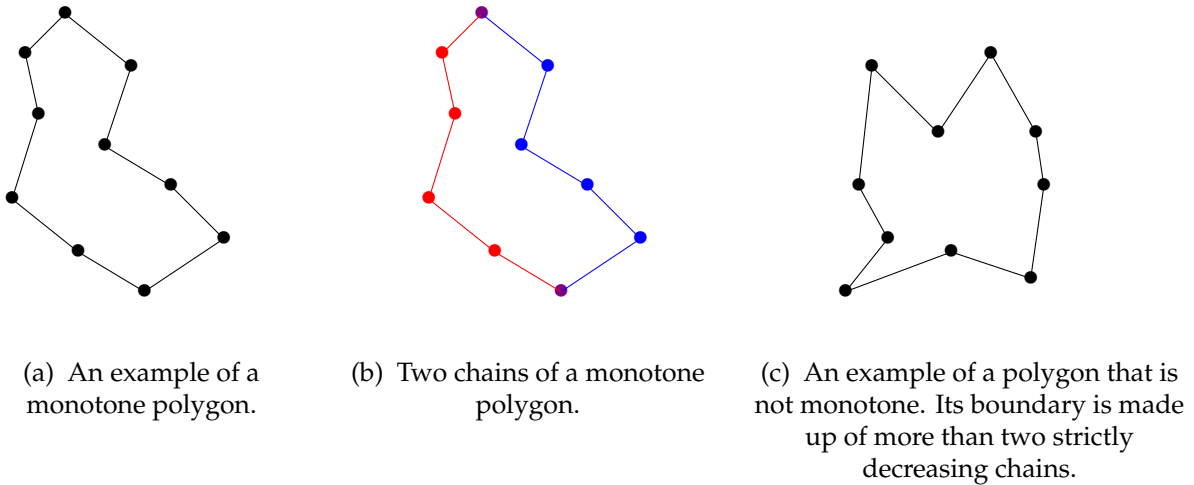


Figure 3.4: Monotone Polygons

For a polygon to be monotone, it must be possible to separate the vertices of the polygon into two chains, where the heights of the vertices in each chain are either always decreasing or always increasing.

2. v is adjacent to a vertex $w \in P$ such that $h(v) > h(w)$, given that v is not the minimum vertex of P .

If all the vertices in a polygon are regular besides the maximum and minimum vertices, then there is no possibility of the polygon possessing *interior cusps*.

Definition 14. A vertex v_k of polygon P is an *interior cusp* of P if the interior angle of v_k is greater than 180° , and one of the following is true for the adjacent vertices of v_k , v_{k-1} and v_{k+1} :

1. $h(v_k) < h(v_{k-1})$ and $h(v_k) < h(v_{k+1})$.
2. $h(v_k) > h(v_{k-1})$ and $h(v_k) > h(v_{k+1})$.

Figure 3.4c has two interior cusps, and the presence of these interior cusps causes this polygon to not be monotone.

Theorem 5. Let Q be an arbitrary simple polygon with n vertices, q_1, q_2, \dots, q_n , appearing in clockwise order around the boundary of Q . Let q_1 have the largest y -coordinate of the vertices of Q ($h(q_1) > h(q_i)$, $i = 1, 2, \dots, n$), and let q_k have the smallest y -coordinate of Q 's vertices ($h(q_k) < h(q_i)$, $i = 1, 2, \dots, n$).

Then, the following statements about Q are equivalent:

1. No vertex of Q is an interior cusp.
2. Q is monotone.
3. All vertices of Q are regular, excluding the maximum vertex and minimum vertex of Q .

Proof. (1 \rightarrow 2) Assume that Q is not monotone. Then, one of the two chains that form from q_1 to q_k is not strictly decreasing.




Figure 3.5: The Presence of Interior Cusps

(a) If q_{c+1} lies to the right of q_c , then the interior angle of q_c is greater than 180° , meaning that an interior cusp is at q_c . (b) With q_{c+1} to the left of q_c , if q_{c+2} has a smaller y -coordinate than q_{c+1} , an interior cusp is at q_{c+1} . (c) It follows that for some q_i between q_c and q_k , an interior cusp must exist at q_i , or else there cannot be a chain connecting q_1 to q_k through q_c .


Let q_c , where $1 < c < k$, be the first vertex on this chain where $h(q_{c+1}) > h(q_c)$. To prevent q_c from being an interior cusp, q_{c+1} must lie to the left of q_c , otherwise the interior angle at q_c would be greater than 180° . Now, to prevent q_{c+1} from being an interior cusp, q_{c+2} must also have a greater y -coordinate than q_{c+1} , and so on. It must continue in this way until one of two things happens (see fig. 3.5):

1. There exists a vertex q_{c+j} , where $c + j < k$, such that $h(q_{c+j}) > h(q_{c+j-1})$ which causes q_{c+j} to be an interior cusp.
2. The path will eventually need to connect to q_k , which has the lowest y -coordinate by definition. Therefore, if no interior cusp has been found in the vertices q_1 to q_{k-2} , then q_{k-1} is an interior cusp, since $h(q_k) < h(q_{k-1})$ and $h(q_{k-1}) > h(q_{k-2})$.

Then, an interior cusp has been found in Q . By contradiction, Q is therefore monotone [9].

(2 \rightarrow 3) Because Q is monotone, the boundary of Q can be divided into two chains, with vertices whose y -coordinates are in strictly decreasing order from q_1 to q_k .

Assume that there exists a vertex in Q , besides q_1 and q_k , that is not regular. Without loss of generality, let this vertex be q_j , where $1 < j < k$, and let $h(q_{j-1})$ and $h(q_{j+1})$ be greater than $h(q_j)$. If both q_{j-1} and q_{j+1} are included in the same chain, the chain would not be strictly decreasing. Therefore, partitioning the vertices of Q into strictly decreasing chains based on the vertices' y -coordinates would require at least four chains.



This would result in Q not being monotone, by definition. Therefore, by contradiction, all of Q 's vertices, excluding the maximum and minimum vertices, are regular.

(3 \rightarrow 1) Let the vertices of Q be regular, excluding vertices that are maximal or minimal, $q_{max} = q_1$ and $q_{min} = q_k$. For all vertices besides q_{max} and q_{min} , by the definition of regular vertices, each vertex is adjacent to exactly one vertex with a greater y -coordinate value than itself, and exactly one vertex with a smaller y -coordinate value than itself. Therefore, no of these vertices can be interior cusps, by definition.

Consider the vertex q_{max} . By definition, $h(q_{max}) > h(q)$, for all vertices $q \in Q$. Therefore, the greatest interior angle that q_{max} could have is 180° , which disqualifies q_{max} from being an interior cusp. With similar logic, q_{min} cannot have an interior angle greater than 180° , so q_{min} cannot be an interior cusp.

Hence, if all the vertices of Q are regular, maximal, or minimal, Q cannot have an interior cusp. □

From the above theorem, we know that partitioning a simple polygon so that all of the polygon's original vertices are regular will result in several monotone polygons. The following list is one process of partitioning an arbitrary polygon with diagonals so that all the polygon's vertices, save its

maximal and minimal vertices, are regular in the context of the newly formed polygons [9]. We will use the non-monotone polygon from fig. 3.4c to illustrate the steps of this regularization, with each step shown in the figures 3.6 and 3.7.

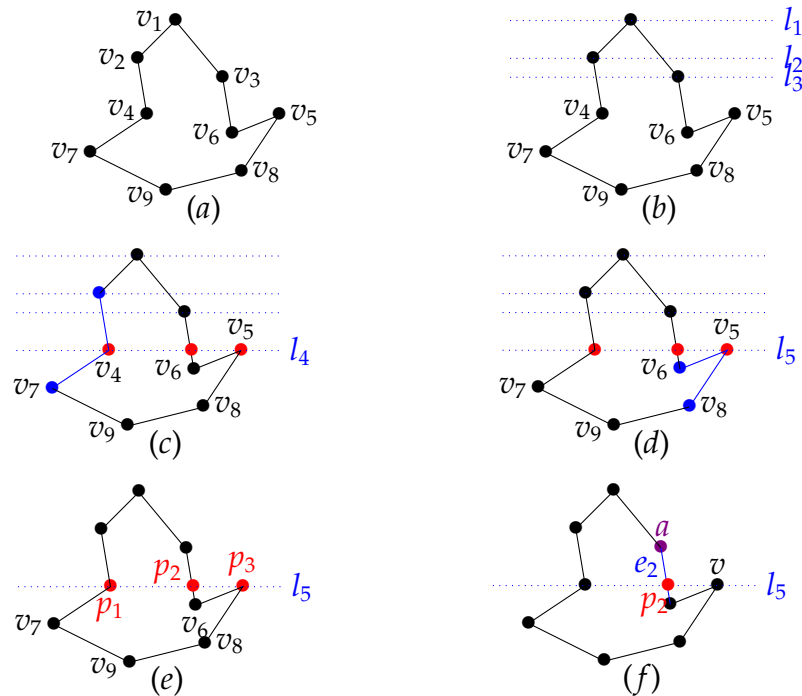


Figure 3.6: Regularization: Ordering and Checking Vertices


The vertices of a non-monotone polygon are ordered by height, and horizontal lines are drawn through each vertex of the polygon. If an l_i intersects the polygon's boundary more than two times, then v_i can potentially be a non-regular vertex. Above, the line that intersects v_4 and v_5 crosses P 's boundary three times. Since v_4 is regular, we know that v_5 must not be, so we continue through the algorithm in consideration of v_5 .

1. Sort vertices of polygon P by decreasing y -coordinates, resulting in


$$v_0, v_1, \dots, v_n.$$

2. For each vertex, starting at v_0 , draw a horizontal line l through v_i .
 - If l intersects P 's boundary one or two times, move on to the next vertex.
 - If not, check to see if v_i is regular. If it is, continue on to vertex v_{i+1} .
 - If v_i is not regular, label the points where the polygon's boundary intersects l as p_1, p_2, \dots . Note that there will exist a p_i such that $v = p_i$.
3. For all p_i on line l , if p_i does not lie on a vertex of the polygon, label the edge it lies on as e_i .
 - If v_i fails regularity by not having an adjacent vertex with a larger y -coordinate, mark the endpoint of e_{i-1} that lies above l as a , and the endpoint of e_{i+1} that lies above l as b , if these edges exist.
 - Otherwise, mark the endpoints of e_{i-1} and e_{i+1} that lie below l as a and b , respectively.
4. Consider the polygon P' with vertices p_{i-1}, v, p_{i+1}, a , and b , skipping the vertices that do not exist.
 - If the edge between a and b intersects edges of P , label the vertex of P that lies inside P' , and has the closest y -coordinate value to v , with u .
 - Otherwise, choose a or b to be marked as u .
5. Add an edge between v and u and return to step two, now considering the vertex v_{i+1} .
6. Once this process has been completed for v_n , remove all edges that were added in the previous steps that lie in P 's exterior.


Figure 3.7f shows two polygons that result from regularizing the non-monotone polygon in fig. 3.4c (these polygons are not unique; the non-monotone polygon could have been partitioned in other ways which would result in different monotone polygons). It is shown that these polygons are monotone, as the boundary of each can be partitioned into two chains of vertices with strictly decreasing y -values (shown in red and blue). The resulting P_2 happens to be a triangle and so is fully triangulated already. One way of triangulating a polygon with more than three vertices, like P_1 , is shown in the next section.



- (a) The shaded region shows the polygon formed by p_2, v , and a . (b) An external diagonal is added to P , so v_5 is now a regular vertex. We move on to the non-regular v_6 .



- (c) p_1 and p_3 are not vertices of P , so e_1, e_3, a and b are marked as shown. (d) A polygon of p_1, v, p_3, b and a is formed.



- (e) An interior diagonal between v and u (formerly b) is created, making v_6 regular. (f) The polygon is cut at the internal diagonal created, forming two monotone polygons.

Figure 3.7: Regularization: Remaining Iterations

The completion of the regularization process for the polygon in fig. 3.6. The process leaves us with two or more monotone polygons, which will then need to be triangulated.

3.3 Triangulation of Monotone Polygons

The next and final step will be to triangulate each of the monotone polygons formed through the regularization process. Several algorithms have been created that allow us to triangulate with ease; one of note was created in 1977 by Garey, Johnson, Preparata, and Tarjan [11].

This algorithm mentions reflex vertices; these are simply vertices that are not convex.

Definition 15. A vertex v in a polygon P is considered to be *reflexive* if the interior angle of v is greater than 180° .

Algorithm 1: Triangulation of a Monotone Polygon [11, 6]

```

1 Push  $p_0$ .
2 Push  $p_1$ .
3 for  $i = 2$  to  $n - 1$  do
4   if  $p_i$  is adjacent to  $v_0$  then
5     while  $t > 0$  do
6       Draw diagonal  $p_i \rightarrow v_t$ .
7       Pop.
8     end
9     Pop.
10    Push  $v_t$  (from when the while loop began).
11    Push  $p_i$ .
12  else if  $p_i$  is adjacent to  $v_t$  then
13    while  $t > 0$  and  $(v_{t-1}, v_t, p_i)$  is not reflex do
14      Draw diagonal  $p_i \rightarrow v_{t-1}$ .
15      Pop.
16    end
17    Push  $p_i$ .
18  end
19 end

```

We will demonstrate how this algorithm works on the polygon that was

regularized in the previous section: fig. 3.7. Because the result of that regularization was more than one monotone polygon, we will only use one of these as our example. For the sake of its usefulness, let's consider P_1 , as P_2 is already a triangle.

A Triangulation Example

First, the vertices must again be ordered by the y -coordinate. Let p_0 be the vertex with the greatest height, and p_{n-1} be the vertex with the smallest height (when n is the number of vertices in the polygon).

A *stack* is used in this algorithm to help keep track of vertices. Stacks are objects usually employed in computer programming; they are defined as ordered collections of items, and items can only be added or removed from the stack at one end, which is referred to as the *top of the stack*. The act of adding an item to the top of a stack is called *pushing* an item, and the act of removing an item from the top of the stack is called *poping* the item. In this algorithm, the items of the stack will be vertices, with the vertex at the bottom of the stack labeled v_0 and the vertex at the stack's top labeled v_t . In the following figures, the stacks will be represented as towers of vertices next to the polygons, with the top and bottom labeled with v_t and v_0 , respectively. The vertices of the polygon will be added or removed from the stack, depending on their attributes.

We will run through the algorithm with each vertex p_i , paying careful attention to what vertices p_i is adjacent to, as well as whether or not p_i is reflexive. For the example polygon we used in the regularization algorithm,




Figure 3.8: Stack Operations

An example stack showing how pushing an item onto the stack places the item at the stack's top. An item can only be popped if it is at the top of the stack; for example, if we wanted to pop a , we would first need to pop b , d , and c , in that order.

we will first push p_0 and p_1 onto the stack (see fig. 3.9). Then, we will determine if p_2 is adjacent to v_0 or v_t , which refer to p_0 and p_1 , respectively. Since p_2 is adjacent to p_0 , the algorithm instructs that we draw a diagonal from p_1 to p_2 . Then, both p_0 and p_1 are removed from the stack, and p_1 and p_2 are pushed onto the stack, with v_0 referring to p_1 now. We then consider p_3 .

For each iteration of this algorithm, p_i will always be adjacent to v_0 , for this particular polygon. We will now look at a different polygon where we have to consider p_i being adjacent to v_t instead.

Another Triangulation Example

In fig. 3.10, we see that p_2 is adjacent to v_t , which is p_1 , during this iteration. Therefore, we consider the interior angle formed by the vertices p_0 , p_1 , and p_2 . Because this angle is reflexive, we simply push p_2 onto the stack and move on to p_3 . This vertex is adjacent to v_0 . Therefore, as long as $t \neq 0$, we add a diagonal between v_t and p_3 . We do this by adding diagonals between p_2 and p_3 , and between p_1 and p_3 , and remove p_2 and p_1 from the stack. Then, since $v_t = v_0$, we remove p_0 from the stack and push p_2 , as it was v_t at the beginning of this iteration, and p_3 (fig. 3.10e).

We have come to the end of this iteration of the algorithm, so we start again in consideration of p_4 . The vertex p_4 is adjacent to p_2 , which is v_0 on the stack, and at the moment $t = 1$, since there are $t + 1$ vertices on the stack. We will therefore draw a diagonal between p_4 and p_3 , and pop p_3 from the stack. Now that $t = 0$, we will pop p_2 , as it is the last vertex left on the stack, and push p_3 and then p_4 , so that $p_3 = v_0$ and $p_4 = v_t$.

We now come to p_5 , which is found to be adjacent to both v_0 and v_t . This is usually a sign that the polygon has been fully triangulated, and in some cases a line is added to the algorithm to end the iteration early when this happens. In algorithm 1, one more iteration than necessary is allowed instead, which causes us to draw a superfluous line between vertices already connected by an edge. For our particular example, this extra diagonal is drawn between p_4 and p_5 , as this algorithm first considers if p_i is adjacent to v_0 , then v_t . Either method results in a triangulated polygon; which algorithm to use is up to the user's discretion. O'Rourke, in his evaluation of this particular algorithm, mentions that it has an $O(n)$ overall time complexity [11].




Figure 3.9: Triangulation: An Example

The first steps of the triangulation of the monotone polygon from fig. 3.7. For each iteration, we find that p_i is adjacent to v_0 . To see the other case of the algorithm in action, check out fig. 3.10.




Figure 3.10: Triangulation: Another Example

Another example of the iterations of the triangulation algorithm with a monotone polygon. The iteration in consideration of p_5 is not shown, as it only results in a superfluous diagonal between p_4 and p_5 .

Conclusion

This chapter has shown that it is possible to turn any simple polygon into a collection of triangulations, and has provided methods for doing so. Using a two-step process, a simple polygon can first be transformed into a set of monotone polygons, and then each of these can be triangulated using an algorithm. Therefore, combined with the watchmen proofs from chapter 2, we have shown how any arbitrary polygon need only have $\lfloor \frac{n}{3} \rfloor$ guards at most to be fully watched. However, not all polygons are the same - certain characteristics of different subsets of floor plans can cause the necessary number of guards to rise or fall. What aspects of a floor plan cause the needed number of guards to drop? What type of floor plans are the most difficult to watch? How can the guards' abilities affect the manpower required to fully watch a polygon? These questioned will be explored in the next chapter.

Chapter 4

Extensions of the Problem

Since the theorems for simple polygons in general have been copiously explored in the past decades, some have become curious about adding conditions to the Art Gallery Problem. By considering only certain subsets of floor plans or types of guards, new findings on lower sufficiency and necessity numbers specifically for these subsets have emerged. Some extensions have sought to make the problem more realistic, giving guards new abilities and focusing on floor plans that would not cause architects to recoil in horror. This chapter includes three extensions that take on this goal, and all three end up with results quite different from those for general polygons.

4.1 Polygons with Holes

Of course, one would expect art in an art gallery, and not all art can be hung on the walls. In cases where art galleries, or any floor plans, have sculptures, pillars, support walls, and other such features that can block one's view, we

consider the concept of polygons with holes.




Figure 4.1: Holes in a Polygon

A polygon P is said to have holes when one or more polygons lie fully inside of P 's interior. The presence of holes in a polygon can restrict the sight of the polygon's guards, shown in red.

Definition 16. A *polygon with holes* can be considered a polygon P where all edges and vertices of one or more polygons H_1, H_2, \dots, H_n lie completely inside of P 's interior. The interiors of H_1, H_2, \dots, H_n are not considered to be part of the interior of P . A polygon without holes can be referred to as a *simply-connected* polygon.

Because of how we identify polygons at the beginning of chapter one (see fig. 1.1), a polygon P which shares one or more vertices with a hole H_i is still considered a polygon with holes; the nodes and vertices of H_i would not be considered as part of P 's boundary. This must be so as we have stated earlier that we are only considering simple polygons when contemplating the Art Gallery Problem. However, our definition of a polygon allows us to ignore a hole in a polygon P if there is an edge connecting a vertex of H_i to a vertex of P (see fig. 4.2). If H_i and P share an edge in this way, then the vertices and edges of H_i can instead be considered as part of the boundary of P .




Figure 4.2: Holes Sharing Vertices or Edges

Polygons that lie inside polygon P , but share vertices with P , are still considered holes. However, if a potential hole and its polygon are connected by an edge, the vertices and edges of that hole can be considered as part of the P 's boundary.

In 1982, Joseph O'Rourke used this definition of simple polygons, and the concept of holes and polygons sharing edges, to show that simple polygons with holes can also be triangulated. The method of doing so is very similar to the triangulation existence proof found in theorem 4 of chapter three.

Theorem 6. *A polygon P with h holes and n vertices, where the vertices on the holes of P are included in n , can be triangulated.*

Proof. We will use induction in terms of the value h .

Base Case: $h = 0$. Theorem 4 states that any simple polygon without holes can be triangulated.

Inductive Hypothesis: A polygon with $h' \leq h$ holes and $n' < n$ vertices can be triangulated.

Inductive Step: Because it is impossible to form a planar polygon with only reflex vertices, there must exist a convex vertex v_2 in P 's boundary with adjacent vertices v_1 and v_3 . Consider adding a diagonal d whose endpoints are v_1 and v_3 .

Case 1: d is internal. The diagonal d then partitions the polygon into two

polygons P_1 and P_2 , where n_1, n_2 are the number of vertices for P_1, P_2 , and h_1, h_2 are the number of holes in P_1, P_2 . It follows that $n_1, n_2 < n$ and $h_1, h_2 \leq h$. By the inductive hypothesis, P_1 and P_2 can be triangulated.

Case 2: d lies partially in the exterior. Then, either v_2 is reflexive, or v_2 is convex, but there exists at least one vertex that lies closer to v_2 than v_1 and v_3 . This vertex, x , is either on the boundary of P , or on the boundary of a hole.

Case 2.1: x lies on δP . By placing a diagonal from x to v_2, P then has been partitioned into polygons P_1 and P_2 , where n_1 and n_2 are the number of vertices in P_1 and P_2 and h_1 and h_2 are the number of holes. By the inductive hypothesis, it is possible to triangulate P_1 and P_2 , and so P is able to be triangulated as well.

Case 2.2: x lies on the boundary of a hole. Then, consider the diagonal from x to v_2 . With this diagonal, P has $n + 2$ vertices (traveling along the boundary of P requires v_2 and x to be passed over twice) and $h - 1$ holes. Since P is still a polygon, P still has at least one convex vertex, so consider the placement of more diagonals until P satisfies the conditions for the inductive hypothesis, or $h = 0$. □

The number of triangles that result from partitioning a polygon with holes in such a way will vary from the number that a simply-connected polygon triangulation has. Remember that triangulations can be treated as graphs, so several results from graph theory are at our disposal: most notably, Euler's theorem.

Theorem 7 (Euler's Theorem). *For any planar triangulation T , $V - E + F = 2$,*

where V is the number of vertices in T , E is the number of edges, and F is the number of faces, or enclosed regions, partitioned by T .




Figure 4.3: Euler's Theorem with Triangulations

An example of Euler's Theorem applied to a triangulation, though the theorem holds for any planar graph.

With this equation, we can find the number of triangles that a triangulation with n nodes and h holes will have. The number of triangles, t , will be $t = F - 1 - h$, since we cannot count the interiors of holes or the exterior of the polygon. Naturally, $V = n$, and $E = \frac{3t+n}{2}$, as we count each edge of each triangle and every boundary edge, then divide by two since we counted every edge twice.

$$\begin{aligned}
 V - E + F &= 2 \\
 n - \frac{3t+n}{2} + h + t + 1 &= 2 \\
 n - t + 2h &= 2 \\
 t &= n + 2h + 2
 \end{aligned}$$

O'Rourke argued that because a simply-connected polygon has $n - 2$ triangles in its triangulation, and requires $\lfloor \frac{n}{3} \rfloor$ guards, then for any polygon

with h holes and $n + 2h - 2$ triangles in its triangulation, it is sufficient to have $\lfloor \frac{n+2h}{3} \rfloor$ guards. However, no polygons have been found yet showing that $\lfloor \frac{n+2h}{3} \rfloor$ vertex guards are necessary, leading some to believe that this number can be lowered.




Figure 4.4: Necessity Polygon with Holes

An example of a polygon that requires $\lfloor \frac{n+h}{3} \rfloor$ guards. This polygon has 24 vertices and 3 holes, and 9 guards are required for the polygon to be fully watched. This is true both in consideration of point guards and vertex guards.

In later years, mathematicians have found that $\lfloor \frac{n+h}{3} \rfloor$ guards are sufficient and sometimes necessary for polygons with holes, when these guards are allowed to be placed anywhere in the polygon as opposed to only on the vertices (these are point guards, not necessarily vertex guards). Hoffman, Kaufmann, and Kriegel, who proved theorem 8 in 1991, conjecture that $\lfloor \frac{n+h}{3} \rfloor$ vertex guards are needed for polygons with holes as well.

Theorem 8. ([7]) *It is sufficient and sometimes necessary to watch a polygon with n vertices and h holes with $\lfloor \frac{n+h}{3} \rfloor$ point guards.*

4.2 Mobile Guards

It is reasonable to assume that the number of guards needed to watch any floor plan will be reduced if the amount of area each guard can watch is increased. To do this, we can assume that guards no longer have to remain stationary.

Definition 17. A guard is considered *mobile* if it is able to move along an edge or diagonal in a polygon. A *mobile guard* who is able to patrol along an edge e or a diagonal d can watch a point p in polygon P if p can be seen by some point on e or d .




Figure 4.5: Two Types of Mobile Guards

(a) It is not possible for a single stationary guard to watch this entire polygon. (b) An open-edge mobile guard is able to watch all points in the polygon that can be seen from any point in its patrol path, excluding the endpoints. This results in the guard not being able to peek around corners. (c) A closed-edge mobile guard can see all points visible from any point on its patrol path, including the endpoints. As shown, one closed-edge mobile guard is sufficient to fully watch this polygon.

This classification of guards can be further divided by considering *open-edge mobile guards* and *closed-edge mobile guards*. The technicalities of

closed-edge mobile guards have been explored for almost as long as the art gallery problem has been around; the question of open-edge mobile guards was first introduced in 2011 by mathematician Viglietta [2].

Definition 18. A point p in polygon P can be seen by an *open-edge mobile guard* if p can be seen from some interior point on the edge or diagonal that the guard patrols. The endpoints of the edge or diagonal are included in a *closed-edge mobile guard's* patrol path.

Unlike other extensions of the art gallery problem, using triangulations in consideration of mobile guards is not as helpful as it is for vertex guards. By the definition of a guard's sight (as seen in definition 3), by placing a guard on an interior diagonal of a triangulation, we can only conclude that this guard can watch the two triangles that include this diagonal. Other methods must therefore be used to find the sufficient number of mobile guards for any polygon.

It is also important to make a distinction between edges and diagonals when it comes to a mobile guard's patrol path. A floor plan that requires its guards to patrol along walls will need more guards than a floor plan whose guards can patrol between any two vertices.

The difference between open and closed-edge mobile guards is equivalent to whether or not the guard can peek around the corners at the ends of their patrol route. Apparently, this ability has a sizable effect on a guard's effectiveness, as it has been found that open-edge mobile guards are only as useful as vertex guards.

Theorem 9. ([2]) *Any monotone polygon with n vertices is sufficiently watched by*

$\lfloor \frac{n}{3} \rfloor$ open-edge mobile guards, and sometimes this number is necessary.

The proof of this theorem is nearly identical to the theorem that a triangulated polygon can be partitioned into $\lfloor \frac{n}{3} \rfloor$ fans; in fact, lemma 1 is explicitly used in this proof as well. Because we know that any polygon has a diagonal that cuts off exactly four, five, or six edges, we can first assume that a polygon with less than n vertices requires less than $\lfloor \frac{n}{3} \rfloor$ open-edge mobile guards, and then prove that any polygon with five, six, or seven edges needs only one open-edge mobile guard. Using induction, we can then conclude that $\lfloor \frac{n}{3} \rfloor$ is sufficient for any monotone polygon [2].




Figure 4.6: Closed-Edge Guards on Edges or Diagonals

A polygon with seven vertices that requires two edge guards but only one diagonal guard. This is therefore a necessity polygon for closed-edge mobile guards that are restricted to edges ($\lfloor \frac{n+1}{4} \rfloor = \lfloor \frac{8}{4} \rfloor = 2$). Allowing guards to travel between non-adjacent vertices means that less guards are needed, which is the case in this example

$$\left(\lfloor \frac{n}{4} \rfloor = \lfloor \frac{7}{4} \rfloor = 1 \right).$$

Meanwhile, closed-edge mobile guards are generally more effective than their open counterparts.

Theorem 10. ([11]) *For any monotone polygon P with n vertices, $\lfloor \frac{n+1}{4} \rfloor$ closed-edge mobile guards that patrol strictly on edges are always sufficient and sometimes necessary to watch P , while $\lfloor \frac{n}{4} \rfloor$ closed-edge mobile guards that are allowed to travel*

between any two vertices are sufficient to guard P , and are also sometimes necessary.

A necessity polygon for theorem 10 was already shown in fig. 4.6, where the seven-vertex shape needs two edge guards when no diagonal guards are allowed, and one diagonal guard if that guard is allowed to patrol between any two vertices in the polygon.

4.3 Orthogonal Polygons

Although one can argue that art galleries have allowance to take on much more unique appearances, most galleries are structurally very similar to any other building - namely, they are rectangular. Most galleries have ninety degree corners and strictly parallel or perpendicular walls, just like the majority of floor plans you would come across. In the context of the Art Gallery Problem, we will represent these traditional floor plans with orthogonal polygons.




Figure 4.7: Examples of Orthogonal Polygons

All orthogonal polygons have right angles exclusively. The polygon in (c) can be rotated so that its edges are horizontal and vertical, making it orthogonal as well.

Definition 19. A polygon P is *orthogonal* if every interior angle of P is 90° or

270° . When placed in the xy -plane, P can be positioned so that all edges are horizontal or vertical. These polygons can also be referred to as *rectilinear* or *traditional*. It is always true that orthogonal polygons have an even number of vertices.

As shown in chapter 3, these types of polygons are still able to be triangulated, and like any other polygon, $\lfloor \frac{n}{3} \rfloor$ guards are sufficient to fully watch any orthogonal polygon. However, it can be shown that we can do even better than that.

Theorem 11. $\lfloor \frac{n}{4} \rfloor$ guards are always sufficient and sometimes necessary to fully watch any simple orthogonal polygon.




Figure 4.8: An Orthogonal Necessity Polygon

This comb-shaped polygon is an orthogonal interpretation of Toussaint's necessity polygon from chapter two. This polygon requires $\lfloor \frac{n}{4} \rfloor$ guards, as it has sixteen vertices, and needs one guard to watch each of its four prongs [11].

The polygon in fig. 4.8 shows that $\lfloor \frac{n}{4} \rfloor$ guards are necessary for some orthogonal polygons. Showing that this number is also sufficient for all polygons of this type is very similar to the proof of the Watchman theorem in chapter 2. However, instead of showing that these polygons can be triangulated, it is instead shown that any orthogonal polygon can be partitioned into convex quadrilaterals. If you thought the triangulation existence proof was tedious, you will not like this one any better; we will leave O'Rourke to explain the logic behind it in his 1988 book [11].




Figure 4.9: Convex Quadrilateralization

An example of an orthogonal polygon that has been partitioned into convex quadrilaterals. Convex quadrilateralization is not always unique.

Now that we can assume any orthogonal polygon can be convexly quadrilateralized, we will show that any orthogonal polygon can be fully watched by $\lfloor \frac{n}{4} \rfloor$ guards. First, let us look at two small theorems that will aid us in this proof.

Theorem 12. *A convex quadrilateralization of a simple orthogonal polygon has $\frac{n-2}{2}$ quadrilaterals.*

Proof. Let Q be a convex quadrilateralization of orthogonal polygon P . For each quadrilateral in Q , choose two non-adjacent vertices and draw a diagonal between them (this is always possible as the quadrilaterals are convex). Then, for each quadrilateral, there now exists two triangles, and we have transformed Q into a triangulation T . By a previous theorem, we know that every triangulation of a polygon with n vertices has $n - 2$ triangles (see theorem 4). Therefore, because we used one quadrilateral of Q to form every two triangles in T , there are $\frac{n-2}{2}$ convex quadrilaterals in Q . \square

Theorem 13. *Two convex quadrilaterals within a quadrilateralized polygon P can share at most 2 vertices.*

Proof. Let Q_1 and Q_2 be distinct convex quadrilaterals.

Case 1: For the purpose of contradiction, assume that quadrilaterals Q_1 and Q_2 share four vertices. Then, it must be the case that $Q_1 = Q_2$, but Q_1 and Q_2 are distinct. Therefore, they cannot share four vertices.

Case 2: For contradiction, assume Q_1 and Q_2 share three vertices, and let these vertices be v_1, v_2 , and v_3 . Without loss of generality, assume v_2 is adjacent to both v_1 and v_3 , and therefore v_1 is not adjacent to v_3 by definition. Because $v_1, v_3 \in Q_1, Q_2$, the line segment between v_1 and v_3 must lie completely inside Q_1 and Q_2 , as these quadrilaterals are convex. This can only occur if the line segment lies completely on the boundary of Q_1 and Q_2 , and v_2 must lie on the boundary as well. However, this would make v_1, v_2 , and v_3 colinear, which is not allowed by our definition of polygons. Hence, Q_1 and Q_2 cannot share three vertices.

We can therefore conclude that Q_1 and Q_2 must share two vertices or less. □

To show that $\lfloor \frac{n}{4} \rfloor$ are sufficient, we will adapt Fisk's proof from chapter two and use a little bit of induction to show that an orthogonal polygon is four-colorable, and then place guards at every instance of a vertex with a least-occurring color.

It is almost enough to simply say that a quadrilateralization is four-colorable by the Four Color Theorem. However, this theorem only states that any planar graph (as in any graph that can be drawn in such a way that its

edges do not cross over each other) can be colored with four colorings, but not every possible four-coloring can be used to position guards.

Note that because these quadrilaterals are convex, one guard placed at any point inside this shape causes it to be fully watched. This includes the vertices of the quadrilateral. However, it is possible to color a quadrilateralization with four colors, but in such a way where some quadrilaterals do not have vertices of all possible colors.




Figure 4.10: An Uncooperative Four-Coloring

This four-coloring of a quadrilateralization is not helpful for placing guards. Placing a guard at the blue or pink vertices results in the polygon not being fully guarded; placing guards at the green or orange vertices results in a number of guards greater than $\lfloor \frac{10}{4} \rfloor = 2$.

Therefore, we can show that it is possible to four-color a quadrilateralization so that each quadrilateral has one vertex of every color.




Figure 4.11: Coloring a Quadrilateralization

Using induction to four-color a quadrilateralization. Note that a four-coloring is not always unique, but it is required that each quadrilateral has a vertex of each of the four colors. In (e), both blue and yellow appear the least amount of times, so guards can be positioned at either color vertices (at blue in (f)) to have the polygon be fully watched.

First, we only consider one quadrilateral from the quadrilateralization. We then take four colors and assign each one to one of the vertices of the quadrilateral. Now, we are able to consider one of the quadrilaterals adjacent (sharing an edge) to this first one. We know that two of this second quadrilateral's vertices are shared with the first, so their colors are already determined. For the two remaining vertices, we color them with the two colors not used on this quadrilateral yet. Therefore, both quadrilaterals have all four colors on their vertices. In fig. 4.11, we see that quadrilateralizations with one quadrilateral require $\lfloor \frac{4}{4} \rfloor = 1$ guard, as a quadrilateral is convex.

Assigning guards to the color that is least occurring in a vertex-colored quadrilateralization will result in it being fully watched. In the worst case, where $n = 4$ and only one quadrilateral exists, the number of guards necessary is strictly $\frac{n}{4}$. However, in most situations quadrilaterals share vertices, so it is highly likely that guards are able to watch more than one quadrilateral at once, making the sufficiency number for any orthogonal polygon $\lfloor \frac{n}{4} \rfloor$.

4.4 Considering Multiple Extensions

The extensions mentioned in this chapter, and many more, can be considered simultaneous to each other. For example, we can think about orthogonal polygons that are surveyed with mobile guards exclusively.

Theorem 14. *A simple orthogonal polygon can be sufficiently watched with $\lfloor \frac{n}{4} \rfloor$ open-edge mobile guards, or $\lfloor \frac{3n+4}{16} \rfloor$ closed-edge mobile guards [2, 3].*

In fig. 4.12 we can see that each prong of the polygon must have its own




Figure 4.12: Open-Edge Mobile Guards on Orthogonal Polygons

A necessity polygon for open-edge mobile guards on orthogonal polygons. This polygon has 32 vertices, and must be watched by at least eight guards ($\lfloor \frac{32}{4} \rfloor = 8$).

guard, since open-edge mobile guards are unable to peek around corners.

A useful property when placing open-edge mobile guards is that there are four types of edges in orthogonal polygons (considering the position of the polygon's interior in relation to these edges). Note that an edge can be horizontal or vertical; if an edge is horizontal, the interior of the polygon can exist above or below this edge. If the edge is vertical, the interior exists to the left or right of the edge. This gives us the four types of edges: upper, lower, right, and left. By the Pigeonhole Principle, at least one of these types must make up one fourth or less of the edges of the polygon (consider the contradiction of a polygon with n sides having all four types of edges occur more than $\frac{n}{4}$ times at once). Placing open-edge mobile guards at the least-occurring type of edge of a polygon causes it to be fully watched. In fig. 4.12, each type of edge occurs eight times, so placing guards at any of the edge types will be sufficient (the figure shows the guards placed at all of the right edges).

Meanwhile, closed-edge mobile guards are more powerful as they are able

to peek around corners.




Figure 4.13: Closed-Edge Mobile Guards on Orthogonal Polygons

Here, an orthogonal with 26 vertices is fully watched by $\lfloor \frac{3(26)+4}{16} \rfloor = \lfloor \frac{82}{16} \rfloor = 5$ closed-edge mobile guards.

For one more combination of extensions to mull over, let us examine orthogonal polygons with holes.

For many years after the creation of the Art Gallery Theorem, it was conjectured that if there are only one or two holes in an orthogonal polygon, $\lfloor \frac{n+h}{4} \rfloor$ guards would still be sufficient, just as it would be for an orthogonal polygon without holes. In the 1990's, the mathematician Hoffman proved that this number held for any number of holes, so long as the guards were point guards [7].

The previous conjectures about vertex guards were replaced by a theorem in the 1980's, as they were proven true, but there is still little to be said about orthogonal polygons with more than two holes in general if we are restricted to vertex guards.




Figure 4.14: Orthogonal Polygons with Holes

The polygon on the left is sufficiently covered by $\lfloor \frac{12}{4} \rfloor = 3$ guards. The polygon on the right, with two holes, can be sufficiently guarded with $\lfloor \frac{n+h}{4} \rfloor$ guards ($\lfloor \frac{20+2}{4} \rfloor = 5$ guards).

Conclusion

Though a result has been found overall for the guarding of all general polygons, varied results can be discovered when one focuses on certain subsets of polygons or guards. In this chapter, we have shown the results for three of the more well-known extensions, though there are many more to examine. In the next chapter, we will quickly look over some more strange extensions, along with some developments of the problem that I wanted to personally look into.

Chapter 5

Stranger Extensions and Further Work

In the previous chapter, we looked at extensions with already existent results. The Art Gallery Problem is such a large topic, however, that there are still extensions of it that have yet to be considered. If I had more time to spend on this topic, the following are some of the extensions I would like to explore.

5.1 Adding Cameras to Galleries

In some cases, an art gallery may look into using technology to partially automate the gallery's security. If the gallery is willing to sacrifice a guard's ability to physically confront thieves, they may choose to swap guards for cameras.

We will also model cameras as points in the polygon; stationary, but able to rotate 360° just as with guards. The live feed from the camera can be viewed

from a device on a guard's person, or from control monitors, which are also points or locations inside the polygon. Control monitors will mainly be useful with mobile guards.

Now, assuming that the gallery wants at least one guard present to watch cameras, we can replace the rest of the guards with cameras if we want to.

Proposition 3. *For any simple polygon with n vertices, it is always sufficient and sometimes necessary for this polygon to be fully watched with one vertex guard with a monitor, and $\lfloor \frac{n}{3} \rfloor - 1$ cameras.*




Figure 5.1: Toussaint's Polygon with Cameras

If the guard (shown in black) has access to the cameras (shown in magenta), then only this one guard is sufficient to watch this polygon.

In regards to non-stationary guards, if a mobile guard has a device on their person to watch cameras, he can see the camera's line of vision at any point in their patrol. If the camera's feed is viewed by a control monitor, then this monitor must exist on a point in a mobile guard's patrol, and the guard can only access the monitor when the guard is also at the control monitor's location.

It may be possible to get more interesting results from this extension by restricting the capabilities of the cameras. This can be done by limiting the number of cameras that can be accessed from a single control monitor or guard's device, or by establishing a maximum distance that a camera can be

from a control monitor in order to access its view (see fig. 5.2). For example, if there are five cameras in a floor plan, but a guard is only able to access the view of four from a control monitor, then more than one control monitor is required. Another solution would be for a guard to be used instead of the inaccessible camera.




Figure 5.2: Cameras with Limited Range

Suppose that camera one (c_1) can only be viewed by control monitor one (m_1), and camera two can only be viewed by monitor two. Then, either a mobile guard can be placed so that it patrols between the two monitors (as shown in the above figure), or two vertex guards are placed at the monitor's positions.

Given more time, I would be interested in seeing how the above limitations would affect the number of sufficient or necessary guards, of both the vertex and mobile variety.

5.2 Mirrors

We can imagine that some walls of a gallery, or some edges of the polygon representing the floor plan, are mirrors (for simplicity's sake, any mirror we consider will make up an entire edge).

If a guard can see the two ends of a mirror, then the guard is able to see every point that the mirror reflects (acting as an open-edge mobile guard on the mirror's wall).

The tricky issue about considering mirrors is the angle of reflection. We will adhere to the rule that a mirror's angle of reflection always equals its angle of incidence.

Definition 20. A point p in a polygon can be watched by a guard g using a mirror on edge e , whose endpoints are e_1 and e_2 , when there exists a point m on the mirror such that line segments gm and pm can be drawn such that gm and pm are completely interior to the polygon, and the angle $gme_1 = pme_2$.




Figure 5.3: Mirrors in Galleries

In parts (a) and (b), the point p can be seen by the guard g , since the point m on the mirror is visible to g . Figure (c) shows the area of the polygon a vertex guard can see from a certain place; figure (d) shows the visible parts of the polygon for the same vertex guard if a mirror is placed on a certain edge.

When considering this type of problem, we will want to establish a few conditions for mirrors and the guards who use them. As stated before, we will only look at mirrors that span the entire edge they are attached to. Also, to ignore polygons with useless mirrors, at least one guard in each floor plan

must be able to see at least one point of the mirror (in a fully watched polygon, it follows that every point of the mirror can be seen by one or more guards, as its edge is a part of the polygon).

In fig. 5.3, the polygon can be fully watched by one guard if it is placed in a certain spot. Now, let's consider a variation of this polygon that must be guarded by more than one guard.

For Toussaint's necessity polygon, a guard placed at any vertex can only see one prong, meaning that each prong needs its own guard.




Figure 5.4: Toussaint's Polygon with a Mirror

Despite the presence of a mirror, any vertex guard in Toussaint's necessity polygon can watch only one prong of the polygon, though in some cases the mirror allows the guard to watch portions of the second and third prong.

Considering only one mirror in the same polygon, it is still only possible for a guard to fully watch one prong, although in some cases they can see parts of other prongs (see fig. 5.4). Therefore, the necessary number of guards is still $\lfloor \frac{n}{3} \rfloor$, even with the presence of one mirror.

Because of this, we would like to make the following conjecture about the number of necessary vertex guards in a gallery with a single mirror:

Conjecture 1. For any simple n -sided polygon that has one edge which

functions as a mirror, $\lfloor \frac{n}{3} \rfloor$ guards are always sufficient and sometimes necessary to fully guard the polygon.

We can use Toussaint's Necessity Polygon (fig. 5.4) as a necessity polygon for this case, as even with the presence of the mirror, three guards are still required to watch the nine-sided shape. In consideration of sufficiency, we already know that $\lfloor \frac{n}{3} \rfloor$ vertex guards are sufficient to guard any simple polygon (theorem 2), so the presence of a mirror in the polygon is trivial.




Figure 5.5: An Example of the View of Vertex Guards

The view of a single vertex guard at each vertex in an L-shaped polygon. The polygon is able to be fully-watched from two of the vertices.

If conjecture 1 is true, is there any point of using mirrors? Let us look at the polygon in fig. 5.5, fig. 5.6, and fig. 5.7, and consider the guard's sight for each vertex of the polygon (or edge, for open-edge mobile guards). For both vertex and open-edge mobile guards, there are six options for the guard's position, and one can notice that there are corresponding positions that cause the same portion of the polygon to be viewed. This follows from theorem 9 in chapter four that says vertex guards and open-edge mobile guards have the same power (as they have the same requirements for guards).




Figure 5.6: The View of Open-Edge Mobile Guards

The view of an open-edge mobile guard from each edge of an L-shaped polygon. The polygon is able to be fully watched by this guard from two of its edges. Comparing these polygons to those in fig. 5.5, we can see that the same portions of the polygons can be watched with vertex guards, and so corroborates the fact that $\lfloor \frac{n}{3} \rfloor$ of both vertex guards and open-edge mobile guards are sufficient for any general polygon.

Considering all positions for mirror guards, however, gives us more than six possibilities, as there are multiple options for placing a mirror that is viewable but non-adjacent to the guard. In some positions, this guard can see just as much as its corresponding vertex and open-edge mobile guard. However, when a mirror is placed on a certain edge, the guard can see a larger part of the polygon (the added area is shown in fig. 5.7 in blue). This leads us to believe that for any position of a vertex guard, there is an optimal position for a corresponding mirror that causes the guard to see more than without it.

Conjecture 2. In terms of the area that a guard can see, the power of the combination of a vertex guard and a mirror is greater than or equal to the power of a vertex guard or an open-edge mobile guard.




Figure 5.7: View of Vertex Guards with Mirrors

The view of a single vertex guard at each vertex in an L-shaped polygon, with the help of a single mirror. For each guard, there are two or more edges where a mirror can be placed so that it is not on an edge adjacent to the guard, and so that the guard can see at least one point of the mirror. Depending on where the mirror is positioned, the guard either sees as much as if they did not have a mirror, or an additional area of the polygon is added to their sight (pictured in blue).



Figure 5.8: Considering Edge Lengths with Mirrors

In the first image, one guard and one mirror are sufficient to fully watch the polygon. If some of the polygon's edges are lengthened, this is no longer the case.

Given more time, I would like to look into how to determine the optimal position for a mirror, given the position of a guard in a polygon. I would also be interested in finding out if there is a type of polygon where the presence of mirrors is most useful; that is, where mirrors add the most area to a guard's sight.

From the polygons earlier in this section, it would be reasonable to say that the length of the edges of the polygon can have a role in the efficacy of a mirror (see fig. 5.8). Looking further into this extension, it may be useful to place weights on the edges of polygons, or their respective graphs, to keep track of edge length.

5.3 Conclusion

Now we have briefly explained the Art Gallery Problem and some of its more well-known results. We have gone into detail about the Watchman theorem, and how the two most popular proofs of it work. Finally, we explored some extensions of the problem, both with results and without. As this chapter has shown, there are many more undiscovered topics in the Art Gallery Problem to interest us. We conclude by expressing excitement about what results may come from this problem in the future.

Bibliography

- [1] Martin Aigner, Gunter M. Ziegler, and Ohio Library and Information Network. *Proofs from The Book*, pages 231–234. Springer, Berlin, 4th edition, 2010.
- [2] Antonio L. Bajuelos, Santiago Canales, Gregorio Hernández, Mafalda Martins, and Ines Matos. Some Results on Open-Edge and Open Mobile Guarding of Polygons and Triangulations. *The Computer Journal*, 58(1):160–171, 2015.
- [3] Iliana Bjorling-Sachs. Edge Guards in Rectilinear Polygons. *Computational Geometry*, 11(2):111–123, October 1998.
- [4] V. Chvátal. A Combinatorial Theorem in Plane Geometry. *Journal of Combinatorial Theory, Series B*, 18(1):39–41, 1975.
- [5] Steve Fisk. A short proof of Chvátal’s Watchman Theorem. *Journal of Combinatorial Theory, Series B*, 24(3):374, 1978.
- [6] Michael R. Garey, David S. Johnson, Franco P. Preparata, and Robert E.

- Tarjan. Triangulating a Simple Polygon. *Information Processing Letters*, 7(4):175–179, 1 1978.
- [7] Frank Hoffmann, Michael Kaufmann, and Klaus Kriegel. The Art Gallery Theorem for Polygons with Holes. In [1991] *Proceedings 32nd Annual Symposium of Foundations of Computer Science*, pages 39–48, Oct 1991.
- [8] Ross Honsberger. *Mathematical Gems II*, volume no. 2. Mathematical Association of America, Washington, 1976.
- [9] D. T. Lee and F. P. Preparata. Location of a Point in a Planar Subdivision and its Applications. *Proceedings of the eighth annual ACM symposium on Theory of computing - STOC 76*, 1976.
- [10] T. S. Michael and Val Pinciu. Guarding the Guards in Art Galleries. *Math Horizons*, 13(3):22–25, 2006.
- [11] Joseph O’Rourke. *Art Gallery Theorems and Algorithms*, volume 3. Oxford University Press, New York, 1987.
- [12] Fred S. Roberts and Barry Tesman. *Applied Combinatorics*. Chapman & Hall, 2 edition, 2009.
- [13] Alan Tucker. The Art Gallery Problem. *Math Horizons*, 1(2):24–26, 1994.
- [14] Pawel Zyliński. Orthogonal Art Galleries with Holes: A Coloring Proof of Aggarwal’s Theorem. *The Electronic Journal of Combinatorics*, 13, March 2006.