THE ART GALLERY PROBLEM IN POLYOMINO CORRIDORS

BY

KIERAN KAEMPEN

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Computer Science
in the Graduate College of the
University of Illinois Urbana-Champaign, 2022

Urbana, Illinois

Adviser:

Professor Jeff Erickson

ABSTRACT

The classical Art Gallery Problem asks for a the smallest set of points, called guards, inside a given simple polygon P, such that every point in P is visible to at least one guard. This problem is known to be computationally hard, even in restricted cases. We consider a special case of this problem, where the input polygon P consists of a path of axis-aligned unit squares joined along edges; we call such a polygon a polyomino corridor. We show that an optimal guard set of a corridor can be computed in linear time if the corridor satisfies certain additional conditions. We also formulate (but do not prove) a natural structural conjecture; if this conjecture is true, an optimal guard set can be found in any corridor in linear time. Finally, we present several related geometric and combinatorial results.

TABLE OF CONTENTS

CHAPTER 1: INTRODUCTION

How many people would you need to place in the Louvre to guarantee that every location in the building could be watched by one of the people? How about the Smithsonian museums, or even your home? This is the type of query that the Art Gallery Problem aims to answer. The general setting of the problem is the following: we are given a polygon P (in some settings, possibly with holes), and we are tasked with finding the smallest set of guards we can place within P such that every point in P is seen by some guard. For two points $p, g \in P$, we say that p is visible to g if the straight line segment joining p and g does not intersect the exterior of P.

The Art Gallery Problem has seen a significant amount of study over the years, from combinatorial, theoretical, and algorithmic perspectives. The problem has been investigated in various different settings, and in each combinatorial upper bounds have been established on the number of guards necessary. Various models of visibility have also been discussed in these settings, as well. However, when using the standard visibility model (that described above), almost all settings that have been investigated theoretically have been found to be NP-hard to solve. This is primarily why the case we're looking at may be interesting. We're looking specifically at polyomino corridors, which we define as a sequence of unit square tiles glued together on edges such that the dual graph of the resulting polyomino is a path. We believe that this class of polygons admits polynomial-time algorithms to calculate a minimum-sized guard set. Past research has shown that with a different visibility model, this problem is relatively easy to solve in linear time. It is also apparent from past research that even slight generalizations of the case that we consider are indeed NP-hard. Together, these factors make the potential of this variant being polynomial-time solvable even more interesting, as it would in some sense tighten the gap between what is NP-hard and what is polynomial.

Our results are primarily separated into two chapters. The first is chapter 3, which contains all of the full results that we've proved. Section 3.2 mostly contains general combinatorial theorems about this variant. Sections 3.3 and 3.4 contain results showing how to optimally guard sub-variants of the problem where the polyominos are long in some sense, and section 3.5 contains results showing how to optimally guard a very specific sub-variant of the problem where the polyominos are jagged and monotone, but regular.

The second is chapter 4, which contains many ideas we've had about the problem that haven't worked out, or haven't worked out yet. Section 4.1 contains our attempts at trying to find a useful structure in optimal guard sets of these polyomino corridors. Section 4.2

contains an algorithm for any polyomino corridor to calculate an optimal guard set that has a particular nice property. Lastly, section 4.3 describes our attempt to improve the lower bounds that we develop.

CHAPTER 2: PAST RESULTS

2.1 GENERAL ART GALLERY

While the Art Gallery Problem has gone through a significant amount of study, its first occurrence was in a paper by Chvátal [6], which proves the so-called *Art Gallery Theorem* - that is, that in any polygon P with n vertices, $\left\lfloor \frac{n}{3} \right\rfloor$ guards are always sufficient and sometimes necessary to guard P. Chvátal's proof is primarily graph-theoretic, and we'll present a short overview of the techniques here.

He first considers a triangulation of the given polygon, and develops the concept of a *fan*. A fan is a triangulated polygon P containing a vertex v such that v is adjacent to all triangles in P. Obviously, a fan can be guarded entirely by a single guard placed at v. Chvátal then proceeds to prove that every triangulated polygon can be partitioned into m fans such that $m \leq \left\lfloor \frac{n}{3} \right\rfloor$. This proof proceeds by induction. The inductive step relies on the existence of an edge in the triangulation that separates the polygon into two other polygons, one of small constant size. Careful casework analysis on the structure of the triangulation in the small polygon, along with the inductive existence of a fan-partition of the larger polygon, yields the desired result.

The partition then immediately implies that $\left\lfloor \frac{n}{3} \right\rfloor$ guards, placed at the "hub" vertices of each of the fans in the partition, is always sufficient to guard the polygon entirely, since it necessarily guards all the triangles, whose union is the polygon itself. This bound is also sometimes necessary, as shown by the crown-shaped polygons that Chvátal presents.

Michael and Pinciu [15] continue on Chvátal's work by presenting a slightly more general version of the Art Gallery Theorem. Their version is as follows: given a polygon P and two sets V^* and E^* from the vertices and edges (respectively) of P, there exists a *constrained* guard set for P of size $\left\lfloor \frac{n+2|V^*|+|E^*|}{3} \right\rfloor$, where every vertex in V^* is a guard and every edge in E^* contains some guard.

This statement is equivalent to the original Art Gallery Theorem when both V^* and E^* are empty. Additionally, this more general statement admits a much nicer proof by induction than the original proof of Chvátal; we will present a short overview of this proof as well.

In contrast to finding a suitable constant-size sub-polygon to separate, the authors only consider an *ear* of the polygon P. This is defined as a sequence of two adjacent edges of P such that the line segment connecting the two distant vertices on the two line segments is an internal diagonal of P. In other words, an ear is a triangle in some triangulation of P such that two adjacent edges of P are edges of the ear.

Let the vertex not adjacent to the interior diagonal on the ear be labeled x. x is then removed from P so that the inductive hypothesis may be applied. In particular, if x is not in V^* and is not incident to an edge in E^*, then the edge opposite x in the ear is added to E^* before applying the inductive hypothesis. This guarantees that the triangle containing x remains guarded by the inductively-produced guard set. If x is in V^*, then x can be included as a guard in the guard set, removing it from V^* before applying the inductive hypothesis. If one of the two edges incident to x is in E^*, then that edge can be deleted from E^*, and x can be added to V^* before applying the inductive hypothesis. Lastly, if both edges incident to x are in E^*, then they can both be removed from E^* and x can be added to V^*, transforming this case into one previously solved. In every case, the weights of the sizes of V^* and E^* and the algebraic particulars of the cases lead to the desired result holding.

Interestingly, both Chvátal's argument and Michael and Pinciu's argument require only placing guards at vertices. It is known that this is not sufficient for optimal solutions to the Art Gallery Problem, but in the cases of these bounds, it apparently is.

2.1.1 Reflex Vertices

Other work has been done based on the nature of the vertices of the polygon P. Obviously a polygon where all vertices are convex (have an internal angle of less than π) can be guarded with a single guard. The difficulty when guarding a polygon comes from the allowance of reflex vertices (those with an internal angle of greater than π), as reflex vertices and the edges adjacent to them can block vision between two otherwise-mutually-visible points in P. A natural follow-up question is whether or not we can come up with similar bounds on the number of guards necessary based on the number of reflex vertices as compared to the number of convex vertices.

Iwerks and Mitchell [12] expand on previous work to complete a comprehensive description of these bounds depending on the number of reflex vertices as compared to the number of convex vertices. In particular, for a polygon P with r reflex vertices and c convex vertices (and with $n = r + c$), the cumulative results are the following:

$$\begin{cases} 1 \text{ guard suffices when } r = 0 \\ r \text{ guards suffice when } r \leq \left\lfloor \frac{c}{2} \right\rfloor \\ \left\lfloor \frac{n}{3} \right\rfloor \text{ guards suffice when } \left\lfloor \frac{c}{2} \right\rfloor < r < 5c - 12 \\ 2c - 4 \text{ guards suffice when } r \geq 5c - 12 \end{cases} \tag{2.1}$$

The contribution from Iwerks and Mitchell is the third case, when r is bounded on either side - the other cases had been proved prior [16] [2]. In particular, since the $\lfloor \frac{n}{3} \rfloor$ bound was already known, the contribution was to give a construction that proved the tightness of that bound give acceptable values of r and c.

The polygons that were dominated by the convex vertices and showed a tight bound in that case were known as *shutter* polygons, and the polygons that were dominated by the reflex vertices and showed a tight bound in that case were known as *pseudotriangle chains*. Iwerks and Mitchell present a procedure by which a given pseudotriangle chain can go through a process of iterated vertex removal (in particular, the removal of reflex vertices) in a certain order.

At every step, the resulting polygon requires exactly $\lfloor \frac{n}{3} \rfloor$ guards to be fully guarded, and the procedure can be stopped at any point, or continued until the polygon reaches a shutter polygon. In other words, for a fixed c, this procedure transforms a pseudotriangle chain with c convex vertices into a shutter polygon with c convex vertices, such that at any step in the process, the number of guards needed to guard the intermediate polygon is exactly $\lfloor \frac{n}{3} \rfloor$. Each reflex vertex removal changes the value of r by only 1, so this procedure essentially shows that for any r that lies in between the two values above, $\lfloor \frac{n}{3} \rfloor$ guards is sometimes necessary.

2.1.2 Complexity

Proceeding beyond the realm of Art Gallery *Theorems*, much work has also been done in the context of the Art Gallery *Problem* - that is, the computational problem of finding an optimally-sized guard set for a given polygon. It is known that the Art Gallery Problem is computationally hard - particularly, that it is NP hard [14]. However, recent work by Abrahamsen et al. [1] has shown the even stronger result that the Art Gallery Problem is $\exists\mathbb{R}$-complete.

For context, this class of problems is related to the *existential theory of the reals* (ETR). ETR essentially asks, given a system of multivariate polynomial equalities and inequalities with real coefficients, whether or not the system can be satisfied by some choice of real numbers for each of the variables. The computational class $\exists\mathbb{R}$ is the set of computational problems that can be reduced to ETR in polynomial time. Abrahamsen et al. prove that every instance of ETR can be reduced to a corresponding instance of the Art Gallery Problem, and vice versa. The reduction is complex, and trying to force vertices to coordinates involving even simple algebraic numbers can lead to polygons with hundreds of vertices. The simplest presented example in the paper is an initial example of a polygon with 50 vertices at integer coordinates that required a guard at an irrational coordinate in any optimal guard set. In fact, the result that the authors showed about $\exists\mathbb{R}$-completeness even applies to polygons whose vertices have integer coordinates.

2.1.3 Algorithms

Nonetheless, despite the computational difficulty of the problem, numerous endeavors have been made to solve the Art Gallery Problem both theoretically and practically. One of the earliest exact algorithms for the problem was presented by Efrat and Har-Peled [7]. In their paper, they tackle both the traditional problem, and also the vertex-restricted version of the problem, where the guards must all be chosen from the vertices of the polygon.

For the vertex-guard case, they present one contribution; namely, a randomized approximation algorithm. The algorithm works by guessing the number of guards needed, then running multiple iterations over the vertices of the polygon P. In each iteration a subset of k vertices is selected at random after weights are assigned to the vertices. Then if the subset doesn't guard some point $q \in P$, then all of the vertices that see the point q have their weights increased if their combined weight is small enough. Obviously if the subset does guard all of P, the algorithm succeeds. The algorithm runs in near-linear expected time, with logarithmic factors involving the size of the optimal vertex-guard set. The returned guard set is an $\mathcal{O}(\log OPT)$-approximation to the optimal guard set, which is an improvement on the previously-known $\log n$-approximations [8].

For the unconstrained case, they present two contributions. The first is an exact algorithm. The problem is translated into a polynomial-sized set of predicates in the first order theory of the reals, and then known algebraic techniques are used to find a solution to these predicates (or return that one does not exist), to check if a solution with k vertex-guards (for some k) exists. This algorithm is exponential in the size of the optimal guard set.

The second contribution in the unconstrained case is another approximation algorithm. This one works very similarly to the approximation algorithm in the vertex-guard case. Instead of vertex guards, it assumes that the guards lie on a grid of points that lies within the polygon P. Using similar techniques, they iteratively refine the weights of the grid points, yielding another $\mathcal{O}(\log OPT)$-approximation, though the OPT here is only the optimum solution with respect to the chosen grid. The algorithm runs in close to the same time, with an extra logarithmic factor based on how small the grid is.

Tozoni et al. [19] focus their attention instead on an algorithm that, while not necessarily guaranteeing termination or a quick runtime, indeed seems to perform well in practice. The algorithm primarily functions by solving discretized versions of the problem over multiple iterations. In so doing, they iteratively build up both a guard set to upper-bound the number of guards needed, and a witness set to lower-bound the number of guards needed. Both geometric techniques and integer linear programming formulations are used within the algorithm. While there is no guarantee that the upper bound and lower bound will meet, the algorithm has shown the ability to solve a large majority of instances with multiple hundreds of vertices, in time much faster than prior algorithms.

More recently, Hengeveld and Miltzow [9] present another practical algorithm that works on polygons with a particular property known as *vision-stability.* A polygon P is considered vision-stable if augmenting or restricting a guard's visibility around a corner by some small positive amount does *not* change the optimal number of guards needed to guard P.

Digression aside, Hengeveld and Miltzow's first algorithm is a one-shot algorithm that works by guessing the vision-stability δ (the amount by which seeing around corners doesn't matter), and then calculating an arrangement of rays originating from reflex vertices based on the value of δ. This arrangement is then used to generate a witness set and a candidate set, which can then be used to build a discretized version of the Art Gallery Problem to be solved by an integer linear program. If the guessed vision-stability is correct for the polygon, this algorithm will return successfully.

In practice this guessing doesn't work particularly well, so the iterative algorithm they develop uses similar ideas over multiple iterations, procedurally refining the arrangement and lowering the supposed vision-stability of P. This iterative algorithm has the same performance guarantees granted by the one-shot algorithm, and in addition performs well in practice. In particular, it led to a noticeable speedup over the algorithm presented by Tozoni et al. [19]. Furthermore, it appears that in some sense, "most" polygons are vision-stable, and thus the algorithm can indeed be applied to them.

2.2 ORTHOGONAL ART GALLERY

The difficulty of the general Art Gallery Problem has led people instead to consider restricted variants of the problem, where the polygons involved are in some sense more well-behaved. One such instance of this is the Orthogonal Art Gallery Problem - in this setting, the polygons are orthogonal, which means every edge is parallel to either the x-axis or the y-axis of a coordinate plane.

Some of the first relevant work was done on the Orthogonal Art Gallery *Theorem* - in particular, Kahn et al. [13] proved that $\lfloor \frac{n}{4} \rfloor$ guards are sufficient to guard any simple orthogonal polygon, which improves Chvátal's previous bound by a factor of about $\frac{3}{4}$. The proof presented by Kahn et al. is substantially more complicated than Chvátal's proof of the original $\lfloor \frac{n}{3} \rfloor$ bound, but we will present some of the overarching techniques here (though not in sufficient detail).

The main result shown by the authors was that every orthogonal polygon P admits a convex quadrilaterization. This is a partition of the interior of P into convex quadrilaterals such that all quadrilateral edges are either edges of P or interior diagonals between vertices of P. The primary technique to show this result relies on the emergence of certain structures in any given orthogonal polygon. A basic structure used is that of *neighboring* edges. These are edges of P that are opposed to each other (either vertically or horizontally), visible to each other (by at least one point), and as close as possible, in the sense that there is no closer opposing edge to either of the neighboring edges. This definition leads to the definition of what's called a *tab*. A tab is a pair of neighboring edges that are connected together by another edge. Tabs are special in that the tab itself, with an interior diagonal drawn between the two unconnected vertices of the opposing edges, must be included in *any* convex quadrilaterization of P.

The authors show that two neighboring edges that aren't a tab, two tabs that oppose each other in a specific way, or a single tab that exhibits a certain behavior, all lead to some inductive structure that can be exploited to show that P has a convex quadrilaterization. They then show that any orthogonal polygon has at least one of these structures, from which it is immediate that any orthogonal polygon does indeed have a convex quadrilaterization.

From this point, the existence of the desired guard set is relatively immediate. The authors consider the graph obtained by taking a convex quadrilaterization of P, along with all of the interior diagonals of the convex quadrilaterals. Through a relatively simple inductive argument from the quadrilaterization, it is shown that the resulting graph is 4-colorable. Then since each convex quadrilateral (with its diagonals) is isomorphic to K_4, each quadrilateral is adjacent to a node with each color. Then choosing the least-represented color leads to a choice of vertices with size at most $\lfloor \frac{n}{4} \rfloor$, that guards each of the quadrilaterals completely (as they are convex and thus guarded by any of their vertices), thus guarding P with the desired number of guards. Note again that similarly to both Chvátal's arguments and to the arguments presented by Michael and Pinciu, this upper bound only requires guards placed on the vertices of the polygon, though this again is also not always sufficient to optimally guard orthogonal polygons.

While nothing yet is known about the $\exists\mathbb{R}$-completeness of the Orthogonal Art Gallery Problem, Schuchardt and Hecker [17] proved that it is indeed still NP-hard. Their first argument was to show that the vertex-restricted Orthogonal Art Gallery Problem was NP-hard. They do this by designing a certain kind of gadget, that is a part of an orthogonal polygon. These gadgets require a certain number of guards within them, and guarding the remaining portion of them leads to long, thin spikes of visibility. When many of these gadgets are appended to a large rectangular region, in addition to other slightly different gadgets at the end of the long visibility spikes, a guarding of the resulting orthogonal polygon simulates a solution to a particular 3SAT instance. The authors show that any given 3SAT instance can be reduced to such an instance of the Orthogonal Art Gallery Problem (in polynomial time, of course), and that solving the instance of 3SAT is equivalent to optimally guarding the polygon. This establishes the NP-hardness of the vertex-restricted Orthogonal Art Gallery Problem. The authors then design slight modifications to their gadgets that allow them to prove in an almost identical way that the unrestricted Orthogonal Art Gallery Problem is NP-hard as well.

Tomás [18] looks specifically at a class of "thin" orthogonal polygons. In particular, she restricts her attention to those orthogonal polygons P where the dual graph of the arrangement obtained by extending edges of P inward is a tree. She proves a few different results; namely, that guarding the entire polygon and guarding the vertices of the polygon are NP-hard, both in the case of vertex-guards, and in the case of boundary guards. For guarding the vertices with boundary guards, her argument depends on certain gadgets, which are subsets of the orthogonal polygon that require a certain number of guards (either boundary guards or vertex guards) within them, such that guarding the remainder of the gadget can be done along a long-distance sight line from elsewhere in the polygon. Assembling many of these gadgets together off of a large rectangle allows for the problem to simulate an instance of the vertex cover problem in graphs. Indeed, any instance of vertex cover can be converted into an instance of guarding the vertices of such an orthogonal polygon in polynomial time, establishing the NP-hardness of the problem. Tomás uses similar gadgets for the problem of guarding the vertices with vertex guards, taking advantage of the long, thin strips of visibility they produce, and instead reduces from the MINIMUM LINE COVERING problem. This problem is as follows: given an arrangement of lines in the plane, find a minimum-sized set of points such that every line contains at least one point from the set. She applies similar techniques as well to the problem of guarding the entire polygon with boundary guards and with vertex guards, in both cases reducing from vertex cover.

2.2.1 Non-Standard Visibility

When restricting our attention to orthogonal polygons and polyominoes, it is also reasonable to consider other visibility models - particularly, *r-visibility* becomes a natural choice. In a polygon P, two points x and y are considered *r-visible* if the minimal axis-aligned rectangle containing x and y lies entirely within P. This is a more restricted definition of visibility, as any two points that are r-visible are visible in the typical sense, but the relation does not necessarily go the other way.

Iwamoto and Kume [11] prove (through an argument about polyominoes) that the Orthogonal Art Gallery Problem with r-visibility in an orthogonal polygon with holes is NP-hard. More on their argument will be discussed in section 2.3. Despite this, however, Worman and Keil [20] actually show that the r-visible Orthogonal Art Gallery Problem is solvable in polynomial time if the input polygon is simple. The techniques they use are complex. They focus instead on the equivalent problem of finding a minimum-size r-star cover of a given orthogonal polygon P. An *r-star* is an orthogonal polygon such that some point within it sees every other point within it by r-visibility. In other words, an r-star is an orthogonal polygon that can be guarded under r-visibility by one guard. They then compute a particular visibility graph based on r-visibility, and find a minimum clique cover of the resulting graph. While normally an NP-hard problem, the authors show that the graph that results from their computations is a *perfect* graph, which are graphs in which the minimum clique cover problem can be solved in polynomial time. Overall, these reductions give them an algorithm to find a minimum r-star cover of the polygon P in $\mathcal{O}(n^{17} poly \log n)$ time.

Investigating more specific variants of the problem, Hoorfar and Bagheri [10] develop a linear time-algorithm for the r-visibility model in the case that the orthogonal polygon P is path-like. In particular in their case, this means that the dual graph of the arrangment derived from extending all vertical edges of P inward is a path. Their algorithm works by partitioning P into multiple more well-behaved subregions, and prove that there is an optimal guard set with the guards placed only in locations between these subregions.

Another variant of the problem is presented by Biedl and Mehrabi [5], in which they discuss the case where the polygon P has bounded treewidth. To define treewidth with respect to an orthogonal polygon, they extend all edges of P inward until they intersect another edge of P. The resulting arrangement can be considered as a graph, where vertices are intersection points between line segments of the arrangement, and edges are the line segments between these intersection points. If this graph has bounded treewidth, then an optimal guard set under r-visibility (among other visibility models) can be calculated in $\mathcal{O}(n)$ time, where hidden constants in the big-O notation depend on the precise value of the treewidth. The primary technique used is to decompose the derived graph into a union of trees, which then simplifies the problem.

2.3 POLYOMINO ART GALLERY

Our final specialization of the Art Gallery Problem is a subcase even of the orthogonal problem - polyominoes. In this setting, the polygons considered are composed of a union of axis-aligned square tiles, all of unit side length, which intersect with each other only along boundaries, such that all vertices of the square tiles lie on lattice points. The square tiles are equivalently referred to as *tiles*, *cells*, or *pixels*.

Biedl et al. [4] do a significant amount of work with combinatorial bounds, hardness results, and algorithmic results for the Polyomino Art Gallery Problem. Their first results are in terms of the Polyomino Art Gallery *Theorem*. In particular, they give bounds in terms of the number of pixels m in the polyomino. The authors show that in any polyomino, $\left\lfloor \frac{m+1}{3} \right\rfloor$ guards are sufficient, and are sometimes necessary. To prove this, they do much of the work with a spanning tree of the dual graph of the polyomino, where the dual is constructed by placing a vertex in every cell and an edge connecting vertices representing adjacent cells. In particular, they construct this spanning tree as a BFS tree with a particular behavior that gives the resulting tree a convenient property. This allows them to represent the polyomino as a sequence of polyominoes, each one differing from the last by a constant-sized group of tiles. Each of these constant-sized groups of tiles can be guarded by $\left\lfloor \frac{m'}{3} \right\rfloor$ guards, where m' is the number of tiles in the group. With some careful analysis, this decomposition leads to the bound above. It is worth noting that, as in the general case and the orthogonal case, the argument for sufficiency of the given bound also produces a guarding with at most that many guards, where all of the guards are on vertices. In general, restricting guards to vertices is also not sufficient to optimally guard polyominoes.

The authors also consider the problem of guarding a polyomino with pixel guards - in this case, a point is seen by a pixel guard if there is straight-line visibility between that point and *any* point within the pixel. Using similar techniques, they show that $\left\lfloor \frac{3m}{11} \right\rfloor + 1$ pixel guards are sufficient, and sometimes necessary to guard any given polyomino, though the cases are substantially more numerous and more complicated than those in the point-guard case.

Biedl et al. also present hardness results, for both the pixel guard case and the point guard case. Particularly, they show that determining if a given polyomino P can be guarded with k pixel guards or with k point guards is NP-hard. Similarly to the arguments presented by Tomás and by Schuchardt and Hecker, the authors create polyomino gadgets, which require a certain number of guards and cause long-distance sight lines to exist that could be used to guard the remainder of the gadget. Instead of reducing directly from a known NP-complete problem, they reduce from a specific instance of the MINIMUM LINE COVERING problem, which was in turn reduced from a general instance of MAX2SAT(2L). MAX2SAT(2L) is the following problem: given a boolean formula in conjunctive normal form such that each clause has at most 2 literals (and each literal is used at most twice), find an assignment of truth values to the literals that satisfies the maximum number of clauses. They use the linearity produced by the gadgets to simulate line intersections, and show that the given instances (and thus the original instance of MAX2SAT(2L)) can be solved whenever the Polyomino Art Gallery Problem can be solved, showing that the problem is indeed NP-hard. The technique used applies equivalently to both the point guard case and the pixel guard case.

Lastly, the authors give polynomial-time algorithms to solve simple cases of the problem under the r-visibility model. The cases they handle are those where the dual graph of the polyomino is a tree or a path. In the tree case, they just invoke the slow polynomial algorithm presented by Worman and Keil [20] that was discussed earlier, mentioning that it is an open problem to give a more efficient algorithm in that case. In the path case, they given a relatively simple greedy algorithm that builds a maximal r-star from the remaining portion of the polyomino. This algorithm runs in linear time and is proved to produce an optimal solution for this visibility model.

As mentioned earlier, Iwamoto and Kume [11] prove that without the assumption of simplicity (more specifically, with the allowance of holes), it is NP-hard to find an optimal guarding of a polyomino, even in the r-visibility model. Their reduction, similarly to the other three mentioned already, relies on the existence of gadgets that must be guarded in a particular way. Their reduction differs, however, in that these gadgets don't produce (and don't need to produce) long-distance visibility lines. Instead, they connect the gadgets together with polyomino strips to form a planar graph, and in so doing can simulate any instance of PLANAR 3SAT. They show that guarding the resulting polyomino with a certain number of guards implies a solution to the instance of PLANAR 3SAT, thus establishing that guarding polyominoes with holes under the r-visibility model is NP-hard.

Alpert and Roldán [3] investigate a type of visibility as yet uninspected with regard to the Polyomino Art Gallery Problem. They consider pixel guards with rook and queen vision - in other words, the same type of visibility that describes the movement of a chess rook or a chess queen. So a rook guard can only see other full pixels in its same row or column, and a queen guard can see those pixels, along with those in the same diagonal. This is a more discretized version of the Polyomino Art Gallery Problem. The authors present two results for each type of visibility - a bounding result and a hardness result.

The bounding results are as follows: for any polyomino P, there is a guarding of P with at most $\left\lfloor \frac{m}{2} \right\rfloor$ rook guards, and a guarding of P with at most $\left\lfloor \frac{m}{3} \right\rfloor$ queen guards, where m is the number of tiles in the polyomino. Additionally, these bounds are indeed sometimes necessary. We will give a short overview of the techniques used to prove these statements. For rook guards, the argument is as simple as "checkerboarding" P - in other words, coloring the tiles with two colors so that adjacent tiles have different colors. Then they place rook guards on whichever color is less represented. Since P is connected, every tile is adjacent to some tile of the other color, and so placing rook guards like this guarantees that every tile is guarded, thus guarding P. A similar argument is applied to queen guards, though it requires a little more care. The tiles are labeled with their ℓ_1 distance (within P) from some particular tile, mod 3. Again, the least-represented value is chosen, and queen guards are placed on all tiles with that value. The sufficiency of this choice comes from the fact that a queen guard sees all tiles within an ℓ_1 distance of 2 from the guard.

The hardness results are the following: the problems of whether a polyomino P can be guarded with k rook guards or guarded with k queen guards are both NP-hard. The reductions used to show this are based off of the same PLANAR 3SAT reductions used by Iwamoto and Kume [11], and use similar techniques to achieve the result.

CHAPTER 3: NEW RESULTS

3.1 DEFINITIONS AND OVERVIEW

We define a *polyomino* as a closed region composed of axis-aligned unit square tiles, whose interior is connected, and such that all the tile vertices lie on lattice points. To simplify our presentation, we consider only simply-connected polyominoes. For a given tile t in a polyomino, we call a tile t' a *neighbor* of t if t' and t share an edge.

A polyomino *corridor* is a simply-connected polyomino with at least two tiles, such that exactly two tiles have only one neighbor, and all other tiles have exactly two neighbors. In other words, the dual graph obtained by replacing each tile with a vertex, and connecting vertices that represent adjacent tiles, is a path. We define a *corner* of a polyomino corridor to be a tile t who either has one neighbor, or who has two neighbors t_1 and t_2 that share adjacent edges of t. We define a *hallway* in a polyomino P to be any maximal axis-aligned rectangle lying within P.

We are primarily concerned with finding a *guarding* of polyomino corridors. For a polyomino corridor P, a point $p \in P$ is *seen by* or *visible to* a point $g \in P$ iff the line connecting p to g lies entirely within P. A *guarding* or *guard set* of P is a set of points $G = \{g_1, g_2, \ldots, g_k\}$ such that for any point $p \in P$, p is visible to some $g_i \in G$. For a given guard set G of P, we consider a tile t to be *totally seen* by a guard $g \in G$ if every point in t is visible to g, and *partially seen* by g if g can see at least one point in t but not every point in t.

We'll now give an overview of the new results presented in this paper. In section 3.2, we will present combinatorial bounds on the number of guards needed based off of the number of times the polyomino corridor turns. Additionally, we present a simple $\frac{5}{2}$-approximation algorithm based on these bounds. In section 3.3, we look specifically at the case where the polyomino is a "shallow stairway" - that is, a monotone polyomino where the horizontal hallways are long and the vertical hallways are length 2, so the polyomino looks like a sequence of long stair steps going upward. We offer a short proof that this case can be guarded optimally in a very simple manner. In section 3.4, we investigate the case where every hallway in the polyomino is long - specifically, where every hallway has length at least 3. Much of the section is spent building up various technical lemmas relating to the structure of any guarding of the polyomino. Eventually, these lemmas lead us to the final result of that section, which is that we can compute a minimum-size guard set of any such long polyomino in $\mathcal{O}(n)$ time. The last of our main results is in section 3.5, where we look at a very specific case, where the polyomino takes the form of a monotone stairway that is steeper than those we consider in section 3.3. Specifically, the polyominoes are monotone, have vertical hallways of length 2, and horizontal hallways that alternate between length 2 and length 3. The main result of this section is a precise description of the minimum number of guards needed to guard any such polyomino. The specificity of the case and the complexity of the solution and proof should serve as indication to why the general problem of optimally guarding polyomino corridors is difficult.

Chapter 4 goes over a few ideas we've had that we were not able to fully prove, but that we think may prove interesting to look into further. Section 4.1 describes our process of conjecturing about the existence of guard sets with a pleasant structure, and the numerous counterexamples that we encountered that led us to continuously refine our conjecture. Section 4.2 details an algorithm that will optimally solve the Art Gallery Problem for polyomino corridors in $\mathcal{O}(n)$ time, assuming the truth of the final conjecture from section 4.1. Lastly, section 4.3 briefly describes our attempts at better refining the combinatorial lower bounds we had developed, and the issues that we encountered in our attempts.

Our contributions differ from past results in that most prior results on this topic have not yielded any "interesting" sets of polygons that admit polynomial time solutions for the general visibility model. It is worth noting that the class of polyominoes we discussed in this paper is exactly one of the classes for which Biedl et al. [4] give a linear-time algorithm to calculate an optimal guard set, but their work is only with the r-visibility model. Similarly, our polyominoes are a subset of the class of orthogonal polygons discussed by Hoorfar and Bagheri [10], for which they also give a linear-time algorithm to calculate an optimal guard set. Again, though, they work in the r-visibility model rather than the "standard" visibility model. Lastly, it is of interest that many of the polyominoes we deal are not vision-stable (as described in the paper by Hengeveld and Miltzow [9]), due to the linearity between edges of the polyominoes, and the placements of guards that those linearities induce. As such, many of the polyominoes in this class would not be optimally solved by the algorithm that Hengeveld and Miltzow develop.

3.2 GENERAL OBSERVATIONS

Theorem 3.1. Given a polyomino corridor P and a guarding G of P, every corner $c \subseteq P$ is totally seen by some $g \in G$.

Proof. Consider some corner c of our polyomino P. If c has two neighbors, then we can consider the two hallways that extend from the corner in question. Without loss of generality, we can assume that the hallways extend up and to the right of our corner. The far ends of

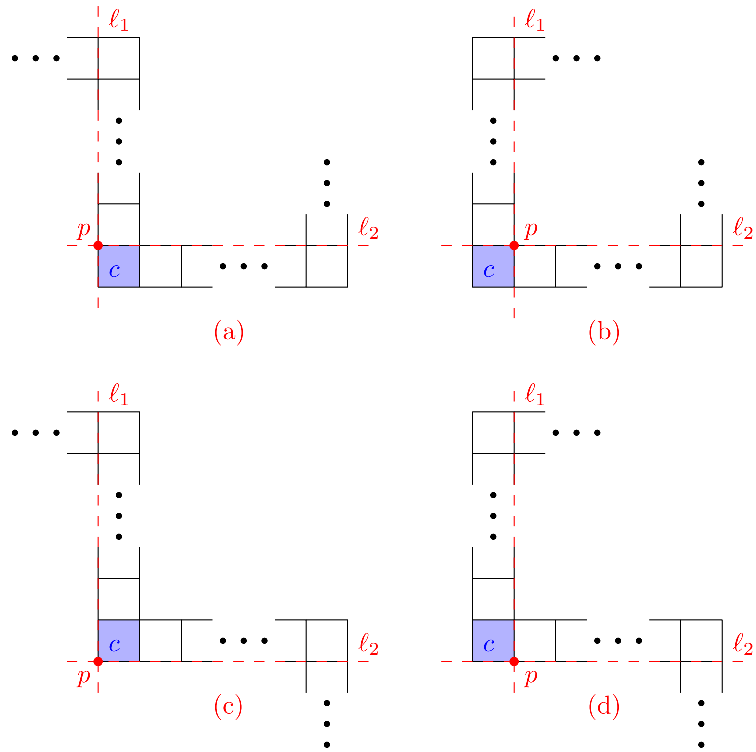

Figure 3.1: The four cases of which direction the hallways extending from a given corner could turn at their other ends

the hallways either terminate or turn in some direction. Figure 3.1 shows the four possible ways that the hallways could turn. Of particular interest is the red point p, which is in the corner in question. Since the polyomino is fully guarded, some guard g sees p. But we see that in order to see p, g must either be to the right (resp. left) of ℓ_1 in cases (a) and (c) (resp. (b) and (d)), or below (resp. above) ℓ_2 in cases (a) and (b) (resp. (c) and (d)). But we see that in any of these cases, this places g in one of the hallways extending from c, from which it totally sees c, as desired. Otherwise, c has one neighbor, in which case one of the previous arguments will apply to the single hallway that extends from c. Note that if any of these hallways terminates, then clearly g is either in that hallway, or in another hallway that doesn't terminate, and still totally sees c. Thus in any case, we have that g totally sees c, and we are done. QED.

Lemma 3.1. Given a polyomino corridor P with m corners, if every point $p \in P$ sees at most k corners simultaneously, then the number of guards needed in a guarding of P is at least $\lceil \frac{m}{k} \rceil$.

Proof. For a guard g, let $\#(g)$ be the number of corners totally seen by g. Then since every corner in P must be totally seen by some guard, we have that $m \leq \sum_{g \in G} \#(g) \leq \sum_{g \in G} k =$

$k|G|$, and thus we see that $|G| \geq \frac{m}{k} \geq \lceil \frac{m}{k} \rceil$, as desired. QED.

Lemma 3.2. In any polyomino corridor P, any point in P can totally see at most five corners.

Proof. Note first that there are four ways that two hallways can intersect. Their intersection is either equal to the entirety of both hallways themselves, equal to a single corner, equal to a single edge between tiles, or equal to a single point. The four intersection patterns are described in Figure 3.2. Note also that as per the arguments in theorem 1, a point totally sees some corner if and only if the point is in one of the hallways incident to the corner. Now consider some point p that sees some number of corners. We can consider the sorted order of corners along our polyomino corridor, and consider the two corners c_1 and c_2 that are totally seen by p and furthest apart in this order. Then since p totally sees c_1 and c_2, p must be in one of the hallways incident to c_1, and one of the hallways incident to c_2. Our point p must lie in the green region in any of the cases. But since c_1 and c_2 are the extremal corners that p can see (since otherwise they would not be maximally far apart under the aforementioned order), we have that all corners seen by p are between c_1 and c_2. It follows that p sees 2 corners in case (a), 3 corners in case (b), 4 corners in case (c), and 5 corners in case (d). Thus, since these are the only ways two hallways can intersect, a given point can only ever see at most 5 corners, as desired. QED.

Corollary 3.1. In any polyomino corridor P with c corners, the number of guards needed to guard P is at least $\lceil \frac{c}{5} \rceil$.

Proof. Lemma 3.1 and lemma 3.2. QED.

This bound has a few uses, but one is to give us an immediate constant-factor approximation algorithm for any polyomino corridor.

Theorem 3.2. There exists a simple greedy algorithm that gives a $\frac{5}{2}$-approximation for the AGP in polyomino corridors.

Proof. We can greedily guard a polyomino corridor P as follows. Consider the set of hallways contained in P, ordered by their appearance along P, and denote them as $h_0, h_1, \ldots, h_{k-1}$. Note that h_i intersects with h_{i+1} in exactly a single corner. We place guards as follows. For every even i, we place a guard g_i in the corner at the intersection of h_i and h_{i+1}. Then g_i guards both h_i and h_{i+1}. Since P is the union of all hallways h_i, we have that this guard set does indeed guard P. The number of guards we've placed is exactly $\lceil \frac{k}{2} \rceil$. Note that the number of corners is exactly one more than the number of hallways, so if P has c corners,

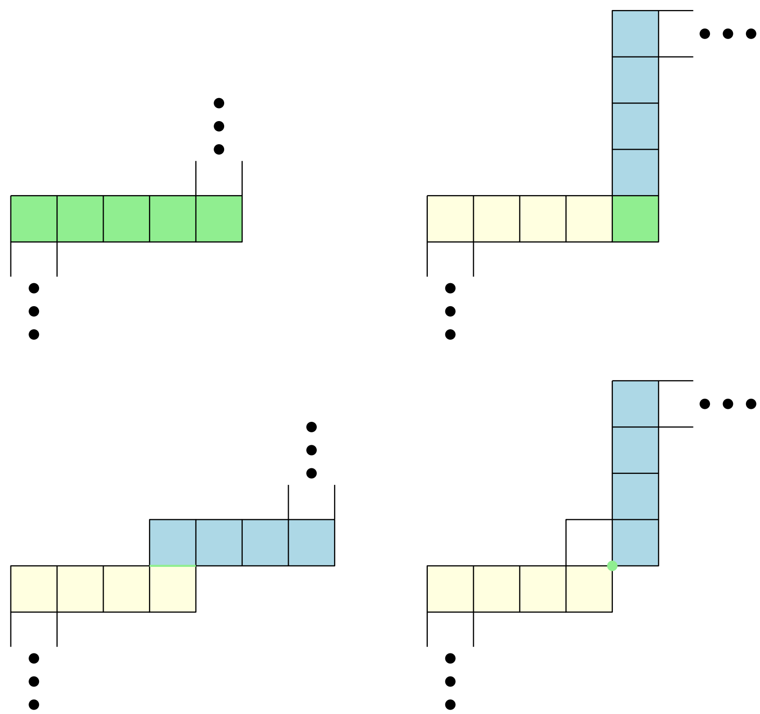

Figure 3.2: The four ways in which two hallways can intersect each other.

$c = k+1$. Then we've placed exactly $\lceil \frac{c-1}{2} \rceil = \lceil \frac{c}{2} - \frac{1}{2} \rceil = \lfloor \frac{c}{2} \rfloor$ guards. Now let OPT be the size of an optimal guard set for P. By corollary 3.1, we have that $OPT \geq \frac{c}{5}$. Then we have that the approximation factor given to us by this algorithm is $\frac{\lfloor \frac{c}{2} \rfloor}{OPT} \leq \frac{\frac{c}{2}}{\frac{c}{5}} = \frac{5}{2}$, so this is a $\frac{5}{2}$-approximation, as desired. Note that depending on how the original polyomino is presented, this algorithm runs in either $\mathcal{O}(n)$ or $\mathcal{O}(n \log n)$, depending on if sorting is necessary. QED.

3.3 SHALLOW STAIRWAYS

In this section, we'll turn our attention to polyomino corridors that take the form of "shallow stairways". These are monotone polyomino corridors such that all vertical hallways have length 2 and all horizontal hallways have length at least 3.

Theorem 3.3. If a polyomino with c corners takes the form of a monotone sequence of horizontal hallways, all with length at least 3, then it can be guarded optimally with exactly $\lceil \frac{c}{4} \rceil$ guards.

Proof. Let P be some polyomino with the aforementioned properties. Note that a point in such a polyomino sees at most 4 corners. Then we see that $\frac{c}{4}$ guards (and thus at least $\lceil \frac{c}{4} \rceil$

guards) are necessary to guard it, by lemma 3.1. Now, we will decompose P into an interior-disjoint union of *lightning bolts*, which are polyominoes composed of two adjacent horizontal hallways. Each lightning bolt can be guarded with one guard, placed on the intersection of the boundaries of the two hallways. Then since we can decompose into $\lceil \frac{c}{4} \rceil$ lightning bolts, we can guard the entire polyomino with $\lceil \frac{c}{4} \rceil$ guards, as desired. QED.

3.4 LONG CORRIDORS ARE LONG

One particular case that we have a solution for is the case where every hallway has length at least 3. We will call such polyominoes *long* polyomino corridors. Additionally, we further classify hallways based on the hallways around them. A given hallway is an *S-hallway* if the two hallways adjacent to it go in opposite directions. A given hallway is a *C-hallway* if the two hallways adjacent to it go in the same direction. For completeness, a hallway that is adjacent to one or fewer other hallways is an *end-hallway*. Lastly, we note that an S-hallway is adjacent to two reflex vertices. If we consider just the interior of the hallway (i.e. the hallway without the two corners), then the line between the two reflex vertices splits this interior into two congruent right triangles. We call each of the triangles defined this way a *half-hallway*. An example of what a half-hallway looks like is shown in figure 3.3.

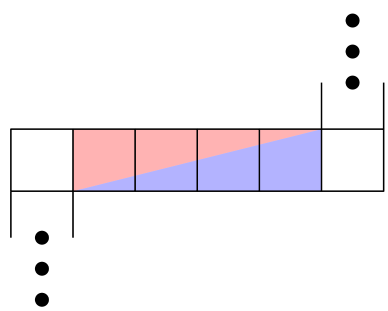

Figure 3.3: An example of two half-hallways, highlighted in red and blue

Theorem 3.4. In any guarding of a long polyomino corridor, every C-hallway h has some guard g such that g totally sees h.

Proof. Let P be a long polyomino corridor, G be a guarding of P, and h be a C-hallway in P. h contains two reflex vertices, which have a segment between them that is a portion of the boundary of P. Let p be some interior point on this segment. Then there is some guard $g \in G$ such that g sees the point p. But we note that in order to see p, g must lie below the line ℓ shown in Figure 3.4. But then g sees the entirety of h, and so we are done. QED.

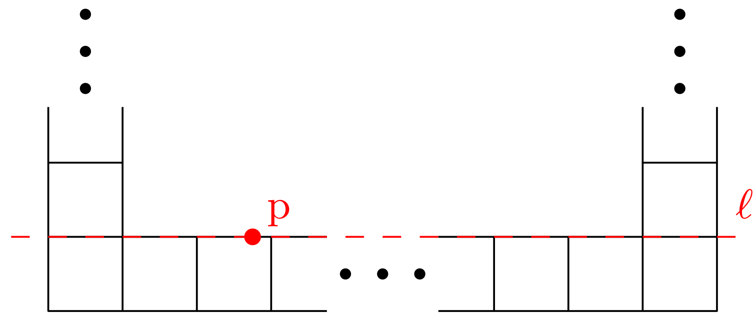

Figure 3.4: A C-hallway and a point on it that must be seen by a guard.

Theorem 3.5. In any guarding of a long polyomino corridor, every S-hallway h either has some guard g such that g totally sees h, or there are two guards g_1 and g_2 such that g_1 and g_2 totally see the half-hallways h_1 and h_2 of h, respectively.

Proof. Let P be a long polyomino corridor, G be a guarding of P, and h be an S-hallway in P. Let ℓ be the line through the two reflex vertices on the boundary of h, and assume that there is no guard g that totally sees h. We will first look specifically at the case where h has a length of 3, and then show that the same argument applies when h has a length longer than 3. Without loss of generality we can assume that the hallway in question is vertical, and we will consider the half-hallway *below* the dividing line, as the argument we use will be symmetrical.

Using Figure 3.5 as reference, we will prove in either case that there must be some guard in the blue shaded region. In particular, we will restrict our attention to the segment s below the point p, and ensure that s is fully guarded. Note that s can't be guarded by any point in h itself, since otherwise that point would totally see h, which would be a contradiction. Thus s has to be guarded just by points in the shaded regions. We look at two different cases, based on the direction of the leftmost vertical hallway.

In the first case, note that for any point on the red portion of the line ℓ, the only point on s that is guarded is p itself. Additionally, for any guard g in the interior of the pink shaded region, there is some point g' on s seen by g that is minimally distant from p. But because g lies above ℓ, g' is not equal to p itself, and so the open segment between g' and p is unguarded. Note that for any finite number of guards in the interior of the pink region, we can find one whose corresponding point g' is minimally distant from p, and then the segment between g' and p is seen by none of the finite number of guards. Thus in order to only use a finite number of guards, there must be some guard in the blue shaded region of the diagram. But then the guard totally sees the lower half-hallway, and so we are done.

In the second case, note that the same argument in case 1 about the shaded pink region

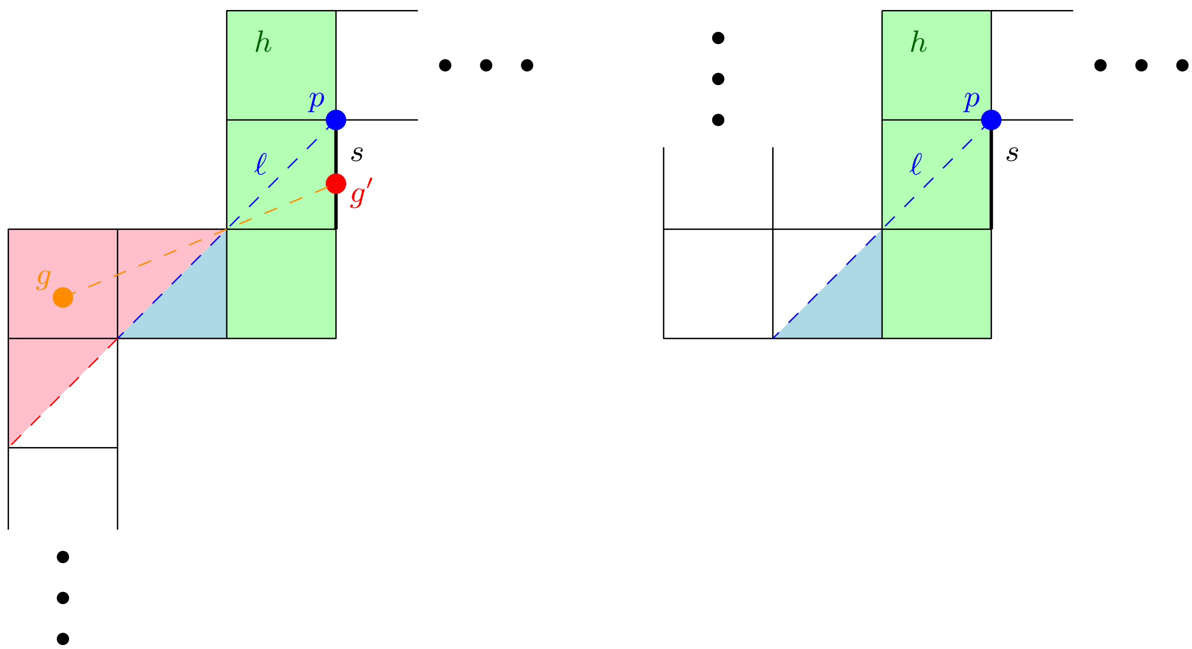

Figure 3.5: A vertical S-hallway and the two cases we consider.

applies, to show that there must be some guard in the blue shaded region, and thus the lower half-hallway is totally seen by some guard.

Now note that if h has a length larger than 3, all cases essentially reduce to the second case. The line ℓ will be steeper, and so will intersect with the interior of the line segment at the bottom of the shown blue shaded regions. Then the same arguments we've made here will apply to that case, and so we are done. QED.

Definition: For a guard g in a long polyomino corridor, let the *restricted visibility region* of g to be the union of corners and half-hallways that are totally seen by g. In other words, anything seen by g that is not a subset of a cell or half-hallway fully guarded by g is removed from g's visibility region. We will also denote this by $vis_r(g)$.

Lemma 3.3. In any guarding G of a long polyomino corridor P, if we restrict all guards in G to only see their restricted visibility regions, then G still guards P.

Proof. Let P be a long polyomino corridor, and let G be a guarding of P. Every cell c is either totally seen by some guard g, or in a half-hallway where each half is totally seen by some guards g_1 and g_2. In the first case, c is not removed from g's visibility region, and in the second case, the portions of c in each of the half-hallways are not removed from g_1 and g_2's visibility regions, preserving that c is guarded by G. QED.

Definition: For any long polyomino corridor P, we can partition P into its half-hallways, corners, C-hallway interiors, and end-hallway interiors. The sets in such a partition intersect only in single points or on line segments. Additionally, if we create a graph by placing a node in the interior of each of these sets, and connecting nodes whose sets border each other along a line segment, the resulting graph will be a path. We will call this partition of P the *canonical decomposition* of P. The elements of the canonical decomposition of P will be referred to as *canonical regions* of P.

Theorem 3.6. For a guard g in a long polyomino corridor P, $vis_r(g)$ is equal to a subset of the canonical decomposition of P, such that the associated graph representation of the subset is connected.

Proof. Let P be a long polyomino corridor, and let g be a guard in P. We will prove that any canonical region in $vis_r(g)$ either contains g, or is adjacent to another canonical region in $vis_r(g)$ that is closer to g. Let c be a canonical region in $vis_r(g)$. Suppose that c does not contain g.

If c is a corner, then as we saw in an earlier proof, g must be in a hallway containing c. In particular, it is in the interior of this hallway, or the other corner incident to the hallway. In either case, g sees the entire hallway, and thus sees its interior. Whether this hallway is an S-hallway, a C-hallway, or an end-hallway, its interior contains a canonical region adjacent to c that is closer to g, and so we are done.

If c is the interior of an end-hallway, then similar to the proof about C-hallways, g must be in this end-hallway. Since g isn't in c, it must be in one of the two corners adjacent to c, and thus whichever corner contains g is closer to g than c, and is totally seen by g, thus it is a canonical region in $vis_r(g)$, as desired.

If c is the interior of a C-hallway, then similarly to the previous case, g must be in one of the two adjacent corners, and so we are done.

Lastly, if c is a half-hallway, there are two more cases. In the first, g sees the entire hallway. Then it sees both the corner adjacent to c, and also the other half-hallway adjacent to c. g either lies within the adjacent corner (in which case we are done), the adjacent half-hallway (in which case we are done), or the other corner in the hallway, in which case the adjacent half-hallway is totally seen by g, and is a canonical region in $vis_r(g)$ that is closer to g than c, and so we are done.

Thus we see that $vis_r(g)$ has the desired structure - considering the two canonical regions that are maximally distant from g on either side of it, we have that there is a sequence of bordering canonical regions in $vis_r(g)$ that get closer to g, until we are in the canonical region that contains g. This gives us a path in the associated graph from one maximally distant region to the other, as desired. QED.

Theorem 3.7. For any guard g in a long polyomino corridor P, $vis_r(g)$ spans at most 4 hallways.

Proof. We will consider cases based on the extremal canonical regions in $vis_r(g)$. Note that these extremal canonical regions can only be corners or half-hallways - this is because if g sees an end-hallway interior or a C-hallway interior, it sees both of the adjacent corners, one of which must be more extremal.

In the case where both of the extremal canonical regions are corners, we have that g must be in the intersection of some hallways containing the two corners. Note also that g cannot be contained in either of these corners, as then it would see past the corner in some direction, contradicting extremality. Thus we have two cases. If the hallways containing the corners are the same hallway, then g lies in the interior of this hallway. Since the two corners are the extremal canonical regions in $vis_r(g)$, we have that $vis_r(g)$ lies entirely in this hallway. In the case that the hallways are different, g lies in the corner at their intersection and sees the two corners at either end. In this case, similarly we have that $vis_r(g)$ lies entirely in the union of these two hallways.

In the case where both of the extremal canonical regions are half-hallways, then refer again to Figure 3.5. Since the half-hallways are extremal, g must not see the entire hallway, and thus it must lie in the blue shaded region for each half-hallway. But we see that for two such regions to possibly intersect, they must be associated with half-hallways whose whole hallways have length 3, and then there are only a few cases we have to consider. See Figure 3.6. We see that in either way the shaded regions intersect, $vis_r(g)$ is contained within either 3 or 4 hallways.

Lastly, we have the case where one of the extremal regions is a half-hallway, and the other is a corner. In this case, we have that g must be in a hallway containing the corner, and also in the associated shaded region. If the hallway containing the corner also contains the shaded region, then since the corner and half-hallway are extremal, $vis_r(g)$ must be contained within only two hallways. On the other hand, the hallway intersects with a single point of the shaded region, and then $vis_r(g)$ is contained in three hallways.

We see that in any case, $vis_r(g)$ is always contained in at most 4 hallways. QED.

Corollary 3.2. For any guard g in a long polyomino corridor P, $vis_r(g)$ contains at most 9 consecutive canonical regions.

Proof. This follows relatively immediately from each of the cases enumerated in the proof of theorem 3.7. We can see two examples explicitly in figure 3.6, where a guard would see 6

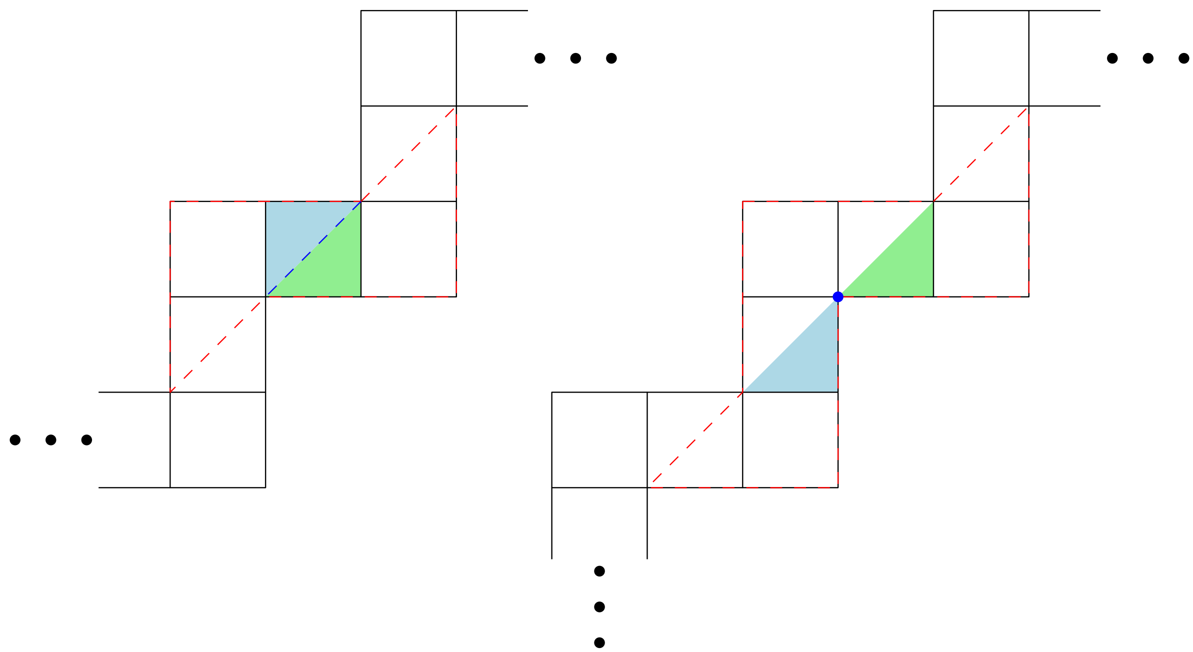

Figure 3.6: The possible cases where both extremal regions are half-hallways

canonical regions on the left (4 half-hallways and 2 corners), or 9 canonical regions on the right (6 half-hallways and 3 corners). QED.

Lemma 3.4. The number of possible configurations of a restricted visibility region is finite.

Proof. From the previous corollary, there can only be at most 9 canonical regions in any region $vis_r(g)$. There are only finitely many ways that we can take a connected sequence of 9 canonical regions, and since a restricted visibility region must take one of those configurations, the number of such possible configurations is finite. Additionally, the complexity of each configuration is finite as well. QED.

Theorem 3.8. There is an $\mathcal{O}(n)$ time dynamic programming algorithm to calculate an optimal guard set for long polyomino corridors.

Proof. Let P be a long polyomino corridor, and consider its canonical decomposition. Traversing along the associated path graph from one of the two starting points, for each region we associate an array element in our DP array A. We also add an additional 0^{th} element before all the others. For $1 \leq i \leq m$ (where m is the number of canonical regions in the decomposition), the value at index i will represent the minimum number of guards needed to guard all canonical regions up to the i^{th} one. The 0^{th} array value will be a sentinel value initialized to 0, and the rest of the values will be initially set to ∞.

We begin iterating through the canonical regions. At region i, we consider all of the finitely many restricted visibility regions that could contain it. To do this, we consider the finitely many sequences of at most 9 canonical regions that contain region i, and check if each one could be a restricted visibility region. This amounts to checking if the union of canonical regions can be guarded with one guard, or equivalently checking if it is star-shaped. O'Rourke [16] notes that this is possible for an arbitrary polygon in $\mathcal{O}(n)$ time. Since the complexity of these regions is constant, we can do this check (and locate a guard if possible) in $\mathcal{O}(1)$ time, so in constant time we can calculate all of the restricted visibility regions that contain canonical region i.

Each of these associated restricted visibility regions r may potentially contain some earlier canonical regions. For each r, we find the largest-index previous canonical region that isn't contained in r. Say this index is k. Then we set $A[i] = min(A[i], A[k] + 1)$. The value $A[k] + 1$ is the number of guards needed to guard r (exactly 1), in addition to all the guards needed to guard everything up to the limit of what r can see. We continue this way over all indices, and the value we return is the value stored in $A[m]$. At each index, we can also store the set of guards as well, building these up as we go to return the optimal guard set at the end.

At each iteration, we need to do only a constant amount of work, since there are a constant number of possible restricted visibility regions to check over. Additionally, it is clear that $m = \mathcal{O}(n)$. Thus, the running time of this algorithm is $\mathcal{O}(m) * \mathcal{O}(1) = \mathcal{O}(n)$, as desired. QED.

3.5 STEEP STAIRWAYS ARE TRICKY

The most general result we've developed so far is the linear-time algorithm for the case where all hallways are long. As it turns out, removing this restriction increases the difficulty of the problem substantially. The primary reason for this is that allowing hallways of length 2 can lead to potential long-distance visibility between a point and a guard, but also while lacking the nice structure we get with shallow stairways. Combined, these factors make finding algorithms and proving algorithmic optimality much more difficult. More will be discussed on this in chapter 4.

However, it may still be worth investigating more subcases that fall under this realm. In particular, we will investigate the case composed of a "stairway" composed of alternating length-2 and length-3 hallways. In other words, a polyomino corridor that is monotone, whose vertical hallways all have length 2, and whose horizontal hallways alternate in sequence between length 2 and length 3. We present the following result:

Theorem 3.9. Let P be a polyomino corridor with c corners that takes the form of a monotone sequence of horizontal hallways, which have lengths alternating between 2 and 3. Suppose the number of horizontal hallways is h (so $h = \frac{c}{2}$). Then the optimal number of guards needed to guard P is given by the following:

$$\begin{cases} \frac{c}{4} & \text{h is even} \\ \left\lceil \frac{c}{4} \right\rceil & \text{P starts and ends with a length-2 hallway} \\ 1 & \text{otherwise, and h} = 1 \\ 2 & \text{otherwise, and h} = 3 \\ \left\lfloor \frac{c}{4} \right\rfloor & \text{otherwise} \end{cases} \tag{3.1}$$

Proof. We will proceed by induction. In particular, we will prove that for the numbers given above, that many guards is both sufficient to guard P, and necessary specifically to fully guard the *corners* of P. For both the base case and the inductive step, we split into cases based on the horizontal hallways that start and end P. For the base case, we have the following:

Case 1: Suppose that $h = 2$, so that P has a length-2 hallway at one end and a length-3 hallway at the other end. Note that P needs at least one guard to guard it, but that that guard can be placed anywhere on the intersection between the two horizontal hallways to guard P. This gives us a guarding of P with exactly $1 = \frac{c}{4}$ guards, as desired.

Case 2: Suppose P starts and ends with a length-2 hallway, and that $h = 1$. Note again that P needs at least one guard to guard it, but this guard can be placed anywhere in the single length-2 hallway that equals P. This gives us a guarding of P with exactly $1 = \left\lceil \frac{c}{4} \right\rceil$ guards, as desired.

Case 3: Suppose P starts and ends with a length-3 hallway. If $h = 1$, then P is just a single length-3 hallway. Then similarly to case 2, one guard is obviously both necessary and sufficient, as desired.

If $h = 3$, then by corollary 3.1, P needs at least $\left\lceil \frac{c}{5} \right\rceil = 2$ guards to guard it. Then if ℓ_1 and ℓ_2 are the starting and ending length-3 hallways, we can place one guard in ℓ_1, and another guard on the border intersection between ℓ_2 and the length-2 hallway in the middle, giving us a guarding of P with 2 guards, as desired.

Lastly, if $h = 5$, then again by corollary 3.1, at least $\left\lceil \frac{c}{5} \right\rceil = 2$ guards are needed to fully guard the corners of P. The unique guarding of P is shown in figure 3.7. Then we see that we can guard P with exactly $2 = \left\lfloor \frac{c}{4} \right\rfloor$ guards, as desired.

For the inductive step, let P be a polyomino corridor which takes the form described above. We once again split into cases.

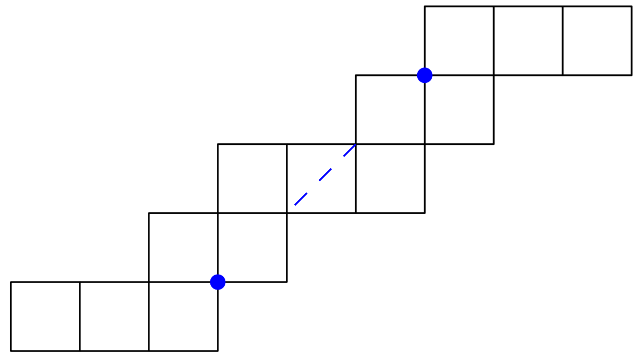

Figure 3.7: The unique guarding of the case when $c = 10$ and P starts and ends with a length-3 hallway

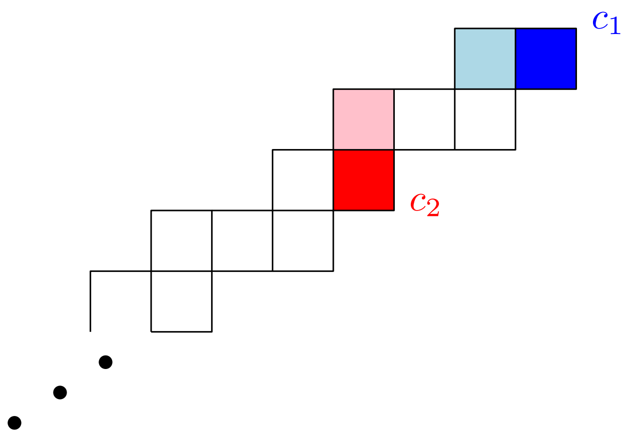

Figure 3.8: The 2-length hallway end of P in case 1 of the inductive step, or one of the ends in case 2

Case 1: Suppose h is even and that $h \geq 4$. See figure 3.8 for reference. Then the extremal corner c_1 in the length-2 hallway at the end has to be guarded by some guard. Note that whatever guard guards c_1, it has to be in the blue-highlighted hallway that contains c_1. But then it can't guard c_2, or any corner to the left of c_2, as it would need to be in one of the hallways that contained c_2 or any of the other corners, which can't possibly intersect with the blue-highlighted hallway.

In other words, since by theorem 3.1 every corner has to be totally seen by some guard, we still need to guard the $c - 4$ corners including c_2 and all corners left of it. These corners are exactly those in a smaller polyomino that has a length-2 hallway at one end and a length-3 hallway at the other, so by our inductive hypothesis, we require $\frac{c-4}{4} = \frac{c}{4} - 1$ guards to guard all of those corners. But then we need the one additional guard that sees c_1, which means that we need at least $\frac{c}{4} - 1 + 1 = \frac{c}{4}$ guards to guard the corners of P. But note also that the entire polyomino including c_2 and everything left of it *can* be guarded with $\frac{c}{4} - 1$ guards by

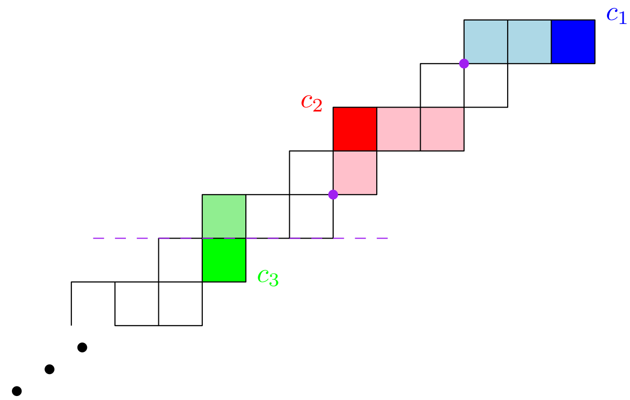

Figure 3.9: One end of P in case 3 of the inductive step

our inductive hypothesis. Then by placing our one additional guard on the lower border of the light blue shaded corner adjacent to c_1, we have a guarding of P with exactly $\frac{c}{4}$ guards, as desired.

Case 2: Suppose h is odd, $h \geq 3$, and that P has a length-2 hallway at each end. Figure 3.8 and the associated argument also apply in this case. The rest of the corners from c_2 and left lie in a smaller polyomino that has a length-2 hallway at each end, so by our inductive hypothesis, $\lceil \frac{c-4}{4} \rceil = \lceil \frac{c}{4} \rceil - 1$ are required to guard those corners to the bottom left. Then in addition to the guard needed for c_1, a total of $\lceil \frac{c}{4} \rceil$ guards are required to guard P. And using a similar scheme as to case 1, we can place the guard for c_1 in the same spot and guard the polyomino starting at c_2 with $\lceil \frac{c}{4} \rceil - 1$ guards, giving us a guarding for P that uses exactly $\lceil \frac{c}{4} \rceil - 1 + 1 = \lceil \frac{c}{4} \rceil$ guards.

Case 3: Suppose P has a length-3 hallway at each end, and that $c > 10$. See figure 3.9 for reference. The extremal corner c_1 has to be seen by some guard g_1. But note then that g_1 has to be in the hallway that contains c_1. In particular, the corner c_2 cannot be seen by g_1, as the hallways containing c_2 do not intersect with the region g_1 must be in. Thus, we need some other guard, say g_2, to see corner c_2. But again, g_2 must be in one of the hallways that contains c_2. So then in particular, the corner c_3 (and every corner to the left of c_3) cannot be seen by either g_1 or g_2, as the hallways that contain those corners cannot possibly intersect the hallways containing c_1 or c_2.

Since by theorem 3.1 every corner has to be totally seen by some guard, we still need to guard the $c - 10$ corners including c_3 and all corners left of it. These corners are exactly those in a smaller polyomino that has a length-3 hallway at one end and a length-2 hallway at the other, so by our inductive hypothesis, we require at least $\frac{c-10}{4} = \frac{c-2}{4} - 2$ guards to guard all of those corners. Note that by the fact that that many guards are necessary and sufficient for that case, this quantity is an integer, so we have that $\frac{c-2}{4}$ is an integer, and so

$\frac{c-2}{4} = \frac{c}{4} - \frac{1}{2} = \left\lfloor \frac{c}{4} \right\rfloor$. So in addition to the two guards needed for c_1 and c_2, we have that we need at least $\left\lfloor \frac{c}{4} \right\rfloor - 2 + 2 = \left\lfloor \frac{c}{4} \right\rfloor$ guards to guard P. But note that by our inductive hypothesis again, those $\left\lfloor \frac{c}{4} \right\rfloor - 2$ guards are sufficient to guard everything left of and including c_3. Note also that using the same guards in figure 3.7, we can guard everything above c_3, as shown in figure 3.9. Thus, we have a guarding of P using exactly $\left\lfloor \frac{c}{4} \right\rfloor$ guards, as desired.

This completes the proof of the claim. QED.

CHAPTER 4: PARTIAL RESULTS AND FUTURE WORK

4.1 ITERATIVE CONJECTURING

As mentioned in section 3.5, the addition of 2-length hallways complicates the problem significantly without the ability to make extra assumptions about how “nice” the polyomino corridor is to deal with (as in section 3.3). When we began thinking about combinatorial bounds in terms of the number of corners, we briefly thought that the problem could be solved thinking only about guarding the corners. However, this is very quickly not true at all. Two polyominoes can require a different number of guards while only differing in a single interior tile.

As demonstrated in figure 4.1, the left polyomino requires only two guards - each of them guards exactly half of the central tile, split along a diagonal. However, if we add one more tile in the center (as in the right polyomino), now two guards are not sufficient. Since there are 10 corners to guard, two guards would have to each guard 5 corners without sharing in order to guard all the corners. The placement of the blue guards is the only possible configuration of guards such that each sees 5 corners. However, this configuration is not enough to guard the entire polyomino, as there is a region in the middle that remains unguarded, and thus requires one more guard to fully guard.

We had also originally considered the possibility that in any guarding of a polyomino corridor, every cell is fully seen by some guard. This also turns out to not be true - the polyomino on the left in figure 4.1 had a special behavior that contradicts this idea: namely, that the cell in the middle is split on its diagonal by the two guards, with each half being seen by a different guard. This is still the only interesting behavior we’ve found that contradicts our assumption by necessity.

Observing that tile split down the middle gave us the idea to potentially expand our previous conjecture to the following: in any guarding of a polyomino corridor, every cell is

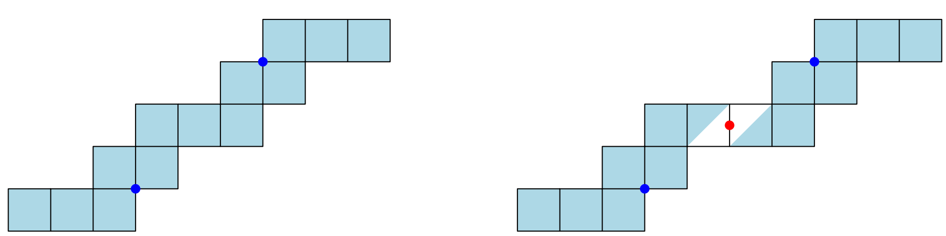

Figure 4.1: A demonstration that guarding the corners of a polyomino is not sufficient to guard the entire polyomino

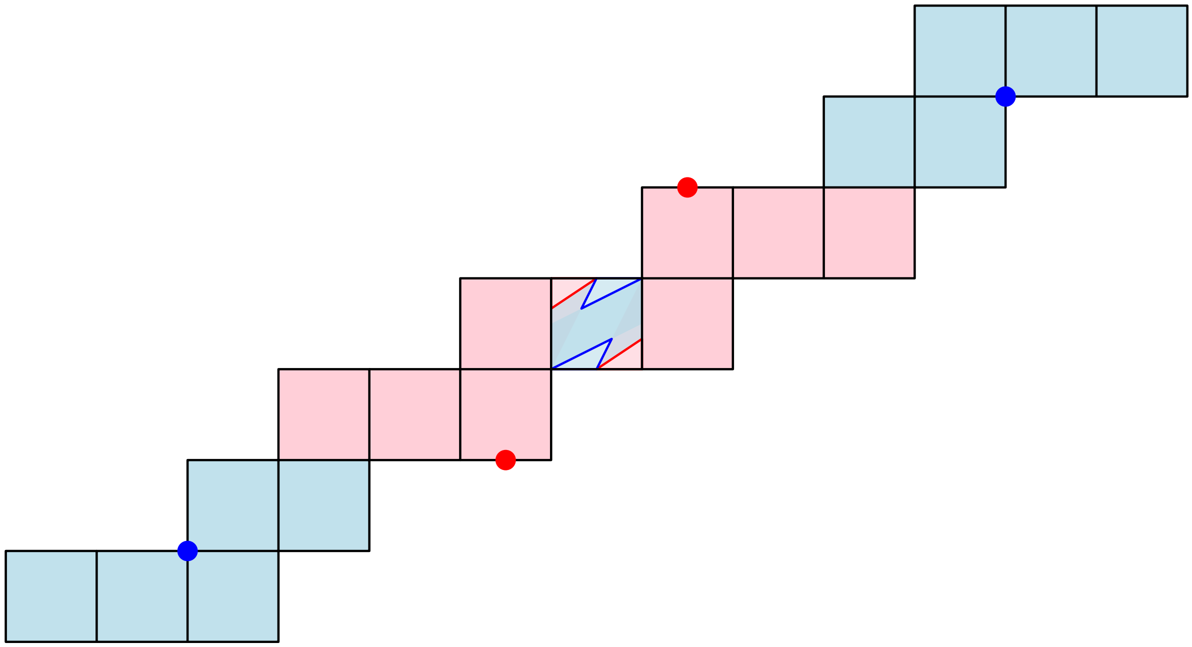

Figure 4.2: An example of an optimal guarding of a polyomino where a single cell requires all 4 guards to fully see it

either fully seen by some guard, or seen by two guards split in half along a diagonal. This conjecture is similar to the ideas discussed in theorem 3.4, theorem 3.5, and lemma 3.3. However, we discovered two things about this. Firstly: it's not true. Figure 10 gives an example of a polyomino that requires 4 guards to be fully guarded. We will not present a full proof of this, but if we were to try to guard this polyomino with 3 guards, two of them would need to see 5 corners, and one would need to see 4. All possible configurations of two guards seeing 5 corners leaves 4 corners remaining that are impossible to see all at once. Thus, 4 guards are required for this polyomino. What is given is a guarding that uses 4 guards, but such that the central tile requires all four guards to fully guard it. The regions of the central tile seen by each guard are shaded, but the highlighted regions are regions that are respectively seen by only one guard, showing that all four are indeed required.

However, this isn't the only optimal guarding of this polyomino. Indeed, another optimal guarding is given in figure 4.3. This guarding is much more "normal", in the sense that it follows our conjecture; every cell is seen entirely by some guard (as per the given shading). As yet, we haven't found a polyomino that *doesn't* admit such a "nice" guarding. This led us to further relax our conjecture to the following: for any polyomino corridor P, *there exists* a guarding of P such that every cell is either fully seen by some guard, or seen by two guards split in half along a diagonal.

However, this leads us to the second problem: this property is probably not strong enough

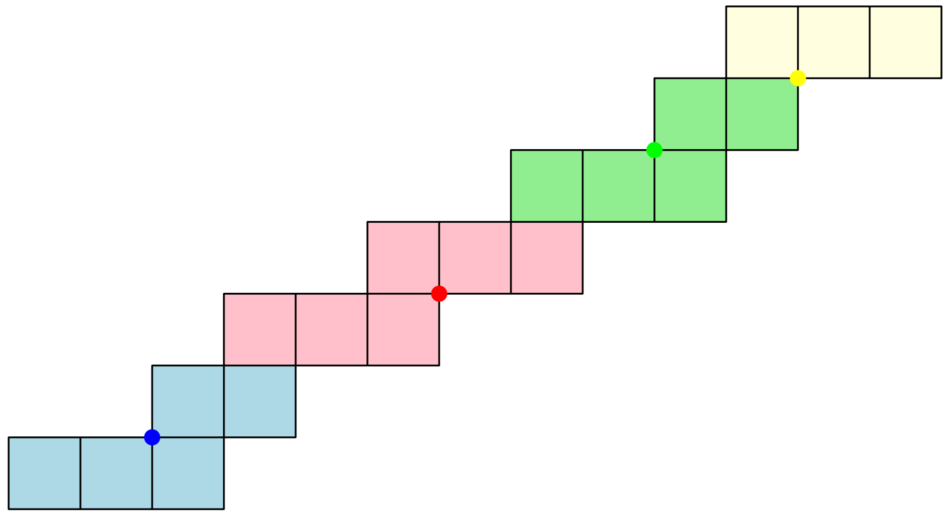

Figure 4.3: Another optimal guarding of the polyomino in figure 4.2, that does not have any strange behavior

to lead to a fast algorithm for the general case. In the case where all hallways are long (discussed in section 3.4), we were able to use the associated property to prove an even stronger property in theorem 3.6. However, we cannot do this in the general case because of the issue with long-distance visibility. In figure 4.4, the placed guard sees half of a tile split on a diagonal from a substantial distance (which can be as arbitrarily long as we'd like), while the region it sees in between its location and that half-tile is not composed of fully-seen corners. The existence of long, non-trivial visibilities such as this prevents us from being able to come up with a nice dynamic programming algorithm if this property is all that we assume.

The numerous counterexamples we found led us to the final version of our conjecture, which is as follows:

Conjecture 4.1. For any polyomino corridor P, there exists an optimal guard set G with $|G| = k$ such that we can partition P into k interior-disjoint connected regions such that:

1. Each region is entirely seen by a single guard
2. The boundaries of the regions are composed of cell boundaries and cell diagonals

We will call such a guarding an *ideal guarding* of P, which induces an *ideal partition* into *ideal regions*. We still have not yet found a counterexample to this conjecture, and in fact we expect it to be true. Unfortunately though, a proof of it (or indeed anything close) has remained elusive. It's possible that some argument about perturbing some existing optimal guard set could yield some traction here, but there's a significant number of details that

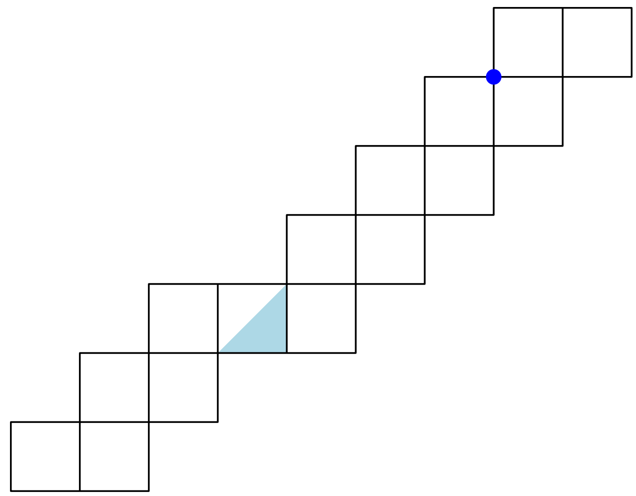

Figure 4.4: An example of a guard seeing a non-trivial portion of a cell from a large distance

need to be attended to for any such argument to work (most of them having to do with long-distance visibility potential).

4.2 A HYPOTHETICAL ALGORITHM

Assuming the truth of conjecture 4.1, we can make the following conjecture:

Conjecture 4.2. There is an $\mathcal{O}(n)$ dynamic programming algorithm to calculate an optimal guard set for polyomino corridors.

The dynamic programming algorithm would look very similar to the algorithm presented in theorem 3.8. To help us out, we have the following lemma:

Lemma 4.1. If a polyomino P has an ideal guarding, then the ideal regions span at most 6 hallways.

Proof. Note that for an ideal region to enter a hallway, it must enter one of the corners in the hallway. But by lemma 3.2, we have that an ideal region can contain at most 5 corners (since the guard that sees the region can see at most 5 corners). But then the ideal region can contain at most the four hallways between the 5 corners, plus possibly one more on each side of the extremal corners. Thus, the ideal region spans at most 6 hallways, as desired. QED.

We describe further lemmas to help us with the design of our algorithm:

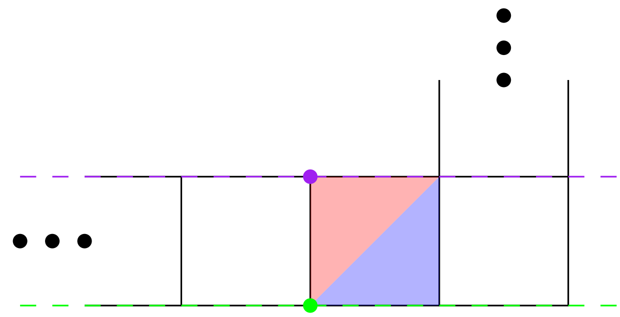

Figure 4.5: An ideal region never needs to have a diagonal border in a cell adjacent to a non-corner

Lemma 4.2. If a polyomino P has an ideal guarding, then it has an ideal guarding where the only diagonal borders of ideal regions are within a non-corner tile adjacent to two corners, and the diagonal borders have both endpoints on reflex vertices.

Proof. We will prove this by adjusting a given ideal guarding. Consider some ideal guarding of P, and consider some region R. Suppose that there is a diagonal border of R that is not in a non-corner tile adjacent to a corner. We now split into cases.

Case 1: If the diagonal border is itself within a corner, then we have eight total cases, based on figure 3.1. If the triangular region of the corner that is in R contains the point p as indicated in figure 3.1, then we have that the guard g_R that guards R sees the corner entirely, and so we can include the entire corner in R and remove the second triangle from the other region it was in. Otherwise, the point p is the single point outside of the triangular region in R. In this case, the other triangular region is part of another ideal region R', and the guard $g_{R'}$ that sees R' sees the point p, and thus sees the entire corner. Then we can assign the entire corner to R' and remove it from R, getting rid of the "bad" diagonal border.

Case 2: If the diagonal border is within a non-corner cell that isn't adjacent to two corners, then the cell is split into two triangles. Say they belong to ideal regions R_1 and R_2. Since the cell c is not adjacent to two corners, it is adjacent to another non-corner on one of its sides. Consult figure 4.5. Without loss of generality, the arrangement of guarded regions will look something like this (possibly with the triangles flipped). Suppose that the red region belongs to R_1. Then the purple point and the green point are in R_1, and are seen by the guard g_{R_1}. But to see both of those points, g_{R_1} needs to be both below the purple dashed line and above the green dashed line. But then g_{R_1} is in the hallway containing this cell, and so it sees the entire hallway, including the cell c, in entirety. Then we can just assign the entire hallway to R_1, removing the problematic diagonal border, and remove whatever portions of the hallway belong to other regions (in particular R_2).

Case 3: Otherwise, the diagonal border is in a non-corner cell that is adjacent to two

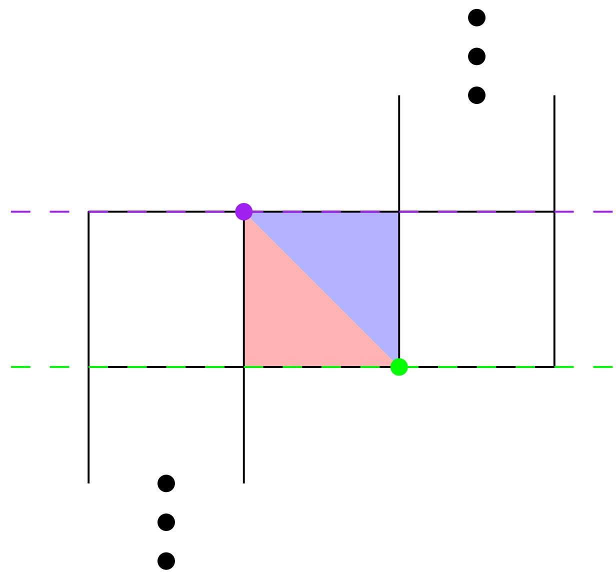

Figure 4.6: An ideal region never requires the diagonal to avoid reflex vertices

corners (the middle cell of a length-3 hallway), but the diagonal is not incident with a reflex vertex. Then we have the case shown in figure 4.6. Suppose the blue region is in some ideal region R. Then g_R sees both the purple and green points, but to do so has to lie between the purple and green dashed lines. But then it sees the whole hallway, and so we can assign the entire hallway to R, removing it from all other regions, and eliminating the problematic border.

Note that by theorem 3.4, if the diagonal is in the middle cell of a length-3 C-hallway, then the C-hallway can be assigned entirely to whichever guard sees it totally, again eliminating the potential for a problematic border. Thus, the only diagonal borders that might *need* to exist in an ideal region are those between reflex vertices in the middle cell of a length-3 hallway. QED.

Our final lemma is the following simple lemma:

Lemma 4.3. If a polyomino P has an ideal guarding, then it has an ideal guarding where there are no ideal region borders between two non-corner cells.

Proof. This lemma follows simply from the idea that if a non-corner cell c is totally seen by some guard g, then g must be in the hallway containing c, and so g sees the entire hallway. Then we can assign the entire hallway (either up to the corner or including the corner) to the ideal region that g sees, removing the problematic internal border. QED.

We can now give the algorithm described by conjecture 4.2, which is also proof of the following theorem:

Theorem 4.1. A minimum ideal guarding of any polyomino corridor P can be computed in $\mathcal{O}(n)$ time.

Proof. We first note that if an ideal region has a diagonal border in it, then the diagonal border is part of an extremal portion of the region, and the region cannot extend beyond it. Now because of lemma 4.1, we know that ideal regions can span at most 6 hallways. Then we can enumerate all possible arrangements of 6 hallways, and all possible ideal regions that could lie within them. We can do this because we know that we can safely restrict our ideal regions so that their only borders are edges of the polyomino itself, borders between corners and non-corners, and reflex-reflex diagonals within the middle cell of a length-3 hallway. Within a 6-hallway region the maximum number of these features is finite, regardless of the length of the hallways involved. Then we can enumerate all possible ideal region shapes. We now partition our polyomino corridor into chunks. The chunks that we partition the polyomino into are as follows:

- Corner cells
- Half cells of the middle cell of a length-3 hallway, split by a reflex-reflex diagonal
- Entire hallway interiors that don't fall into the second category

Note that there are $\mathcal{O}(n)$ of these chunks, where n is the number of cells in the polyomino. We represent each chunk with an array entry $A[i]$, based on its position along the polyomino, and add a sentinel value $A[0] = 0$. All the $A[i]$ for $i > 0$ are initialized to ∞. $A[i]$ will represent the smallest number of ideal regions needed to partition the polyomino up to and including the i^{th} chunk. Then we iterate through the values $A[i]$. At each, we iterate through all finitely many possible ideal region shapes, and place them to cover chunk i in the finitely many ways that are possible. Each possible shape covers the chunks as far back as chunk j for some $j < i$. Then we check the value of $A[j-1]$, And assign $A[i] = \min(A[i], A[j-1]+1)$. We do this for all entries $A[i]$, $i > 0$.

When the algorithm terminates, $A[m]$ (where m is the number of chunks) will contain the minimum number of guards necessary to guard the polyomino. Note that we can also store, in each chunk, some identifier for the ideal region that covers that chunk in the optimal solution. Then we can iterate backward from $A[m]$ to find the set of chunks (and their associated guards), thus producing a guard set with $A[m]$ guards. Assuming the truth of conjecture 4.1, this will produce an optimally-sized ideal guarding, which in turn is an optimal guarding of the original polyomino. Note that we iterate through $m = \mathcal{O}(n)$ chunks, doing a constant amount of work for each, given us a runtime of $\mathcal{O}(n)*\mathcal{O}(1) = \mathcal{O}(n)$. QED.

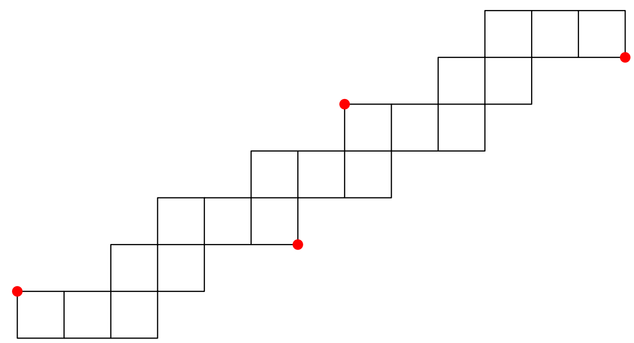

Figure 4.7: A witness/independent set in a polyomino corridor

4.3 SEARCHING FOR A BETTER LOWER BOUND

When trying to solve the Art Gallery Problem practically in other settings, it is useful to iteratively construct what is called a *witness* set - that is, a set of points such that no two of them are visible to the same guard. If you can construct a witness set of size k for a polyomino corridor P, this immediately implies that guarding P requires at least k guards, as any guard can see at most one of the points in the witness set, and all need to be seen. We initially thought that it may be able to construct such a witness set for polyomino corridors from a discretized independent set problem, using as vertices the convex vertices of our polyomino, and connecting with edges those that are mutually visible by some point. For the polyomino corridor in figures 4.2 and 4.3, we can find a witness set from the convex vertices (as shown in figure 4.7) that proves at least 4 guards are necessary to guard the polyomino.

However, it turns out that we can't always find such a witness set from the convex corners, and in fact, we sometimes cannot find such a witness set at all. Consider the polyomino corridor P given in figure 4.8, and suppose that we were trying to find within P a witness set of size 3 (recall from figure 4.1 that 3 guards are necessary to guard P). If we were trying to find a witness set of size 3, we must have one of our witness points in the pink region in the left image. If all 3 points were outside of the pink region, then because the remainder of the polyomino can be guarded by the red guards, by the pigeonhole principle at least two of the witness points would be seen by the same guard, which is impossible. Thus there must be some witness point in the pink region, and a symmetrical argument applies to the other side of the polyomino. So we must have one witness point in each of the pink regions on the right side of the image. But then regardless of where we put the third point, it will either be in one of the green regions, or one of the blue regions. But the blue regions are visible to one of the blue guards (which sees one of the pink witness points), and the green regions are visible to one of the green guards (which also see one of the pink witness points). Thus, even

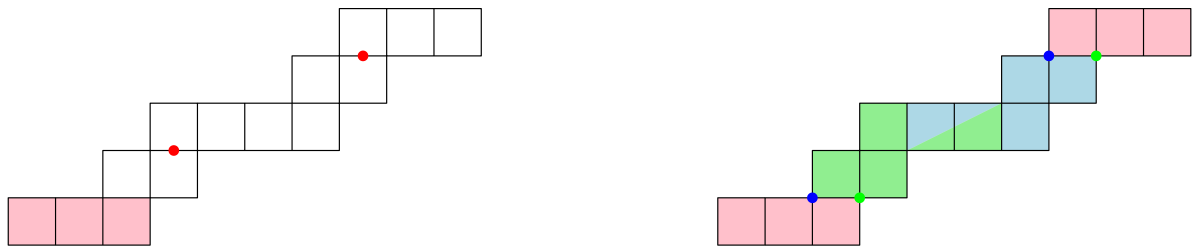

Figure 4.8: A polyomino corridor that does not admit a witness set equal to the size of its optimal guarding

if we allow picking our witness set from *anywhere* within the polyomino, we can't always find a set with the same size as the minimum number of guards needed.

REFERENCES

[1] Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow. The art gallery problem is $\exists\mathbb{R}$-complete. *Journal of the ACM* 69(1):1–70. Association for Computing Machinery (ACM), 2022.

[2] Louigi Addario-Berry, Omid Amini, Jean-Sébastien Sereni, and Stéphan Thomassé. Guarding art galleries: The extra cost for sculptures is linear. *Algorithm Theory – SWAT 2008*, 41–52, 2008. Springer Berlin Heidelberg.

[3] Hannah Alpert and Érika Roldán. Art gallery problem with rook and queen vision. *Graphs and Combinatorics* 37(2):621–642, 2021.

[4] Therese Biedl, Mohammad T. Irfan, Justin Iwerks, Joondong Kim, and Joseph S.B. Mitchell. Guarding polyominoes. *Proceedings of the Twenty-Seventh Annual Symposium on Computational Geometry*, p. 387–396, 2011. SoCG '11, Association for Computing Machinery.

[5] Therese Biedl and Saeed Mehrabi. On orthogonally guarding orthogonal polygons with bounded treewidth. *Algorithmica* 83(2):641–666. Springer Science and Business Media LLC, 2020.

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

[7] Alon Efrat and Sariel Har-Peled. Guarding galleries and terrains. *Information Processing Letters* 100(6):238–245, 2006.

[8] Subir Kumar Ghosh. Approximation algorithms for art gallery problems in polygons. *Discrete Applied Mathematics* 158(6):718–722, 2010.

[9] Simon Hengeveld and Tillmann Miltzow. A practical algorithm with performance guarantees for the art gallery problem, 2020.

[10] Hamid Hoorfar and Alireza Bagheri. Guarding path polygons with orthogonal visibility, 2017.

[11] Chuzo Iwamoto and Toshihiko Kume. Computational complexity of the r-visibility guard set problem for polyominoes. *Discrete and Computational Geometry and Graphs*, 87–95, 2014. Springer International Publishing.

[12] Justin Iwerks and Joseph S.B. Mitchell. The art gallery theorem for simple polygons in terms of the number of reflex and convex vertices. *Information Processing Letters* 112(20):778–782, 2012.

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

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

[15] T. S. Michael and Val Pinciu. The art gallery theorem, revisited. *The American Mathematical Monthly* 123(8):802–807. Mathematical Association of America, 2016. ⟨http://www.jstor.org/stable/10.4169/amer.math.monthly.123.8.802⟩.

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

[17] Dietmar Schuchardt and Hans-Dietrich Hecker. Two np-hard art-gallery problems for ortho-polygons. *Mathematical Logic Quarterly* 41(1):261–267, 1995.

[18] Ana Paula Tomás. Guarding thin orthogonal polygons is hard. *Fundamentals of Computation Theory*, 305–316, 2013. Springer Berlin Heidelberg.

[19] Davi C. Tozoni, Pedro J. De Rezende, and Cid C. De Souza. Algorithm 966: A practical iterative algorithm for the art gallery problem using integer linear programming. *ACM Trans. Math. Softw.* 43(2). Association for Computing Machinery, 2016.

[20] Chris Worman and J. Mark Keil. Polygon decomposition and the orthogonal art gallery problem. *International Journal of Computational Geometry & Applications* 17(02):105–138. World Scientific Pub Co Pte Lt, 2007.