Chapter 2

Guarding Polygons

2.1 Chvátal's Art Gallery Theorem

We start with the proof of Chvátal's Art Gallery Theorem. The proof presented here is due to S. Fisk [56].

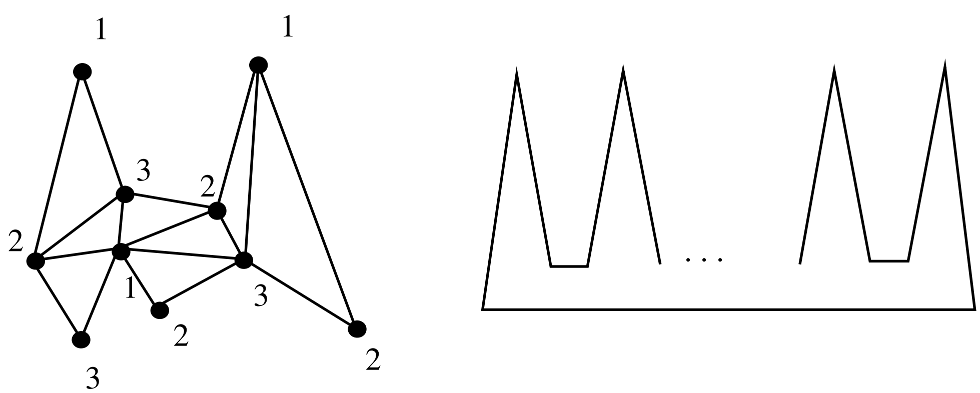

Figure 2.1: Illustration of proof of Chvátal's Art Gallery Theorem.

Theorem 2.1.1 $\lfloor \frac{n}{3} \rfloor$ *stationary guards are always sufficient and occasionally necessary to illuminate a polygonal art gallery with n vertices.*

Proof: Consider an arbitrary simple polygon P with n vertices. Obtain a triangulation T of P by adding $n-1$ interior diagonals. (See Figure 2.1.)

By Theorem 1.2.3 we can color the vertices of P using three colors $\{1, 2, 3\}$ such that any two vertices joined by an edge of P or a diagonal of T receive different colors. This partitions the vertex set of P into three chromatic classes C_1, C_2 and C_3. Since the vertices of each triangle of T receive different colors, each chromatic class guards P. Place a guard at each vertex of the smallest chromatic class and our result follows.

To see that $\lfloor \frac{n}{3} \rfloor$ guards are sometimes needed, consider the comb polygon $Comb_m$ with $n = 3m$ vertices presented in Figure 2.1. It is easy to see that to guard P_m we need at least m guards. ∎

2.1.1 A linear time implementation of Chvátal's result

Using Fisk's proof, together with Chazelle's linear time triangulation algorithm, we can develop a linear time algorithm to find $\lfloor \frac{n}{3} \rfloor$ guards to protect a polygon as follows:

Algorithm 2.1.1

1. *Triangulate p.*
2. *Three color the vertices of $GT(P)$.*
3. *Place the guards at the smallest chromatic class.*

To show that Algorithm 2.1.1 runs in linear time, all we have to show is that Step 2 of Algorithm 2.1.1 can be implemented in linear time. To show this, we notice that by Excercise 3 of Section 1.3, $GT(P)$ always has a vertex v of degree 2. We now notice that any 3-coloring of $GT(P) - v$ can be extended to a 3-coloring of $GT(P)$. It is now easy to see that this procedure leads to a linear time algorithm.

Chvátal's proof of Theorem 2.1.1 was slightly more complicated than Fisk's. We now outline Chvátal's proof. The details will be left as excercises for the reader. The main idea behind Chvátal's proof is the following result:

Lemma 2.1.1 *Any triangulation of a polygon with n vertices has a diagonal d that splits P into two polygons P_1 and P_2, such that P_1 has 5, 6, or 7 vertices, $n \geq 6$.*

Using an inductive argument and a case analysis, we can now finish Chvátal's proof as follows:

If P has 3, 4 or 5 vertices, P can always be guarded with a single vertex. Suppose then that P has at least six vertices.

Case 1: If P_1 has five vertices, then P_2 has $n-3$ vertices. Guard P_1 with a single guard and P_2 with $\lfloor \frac{n-3}{3} \rfloor$ guards. This produces a total of $\lfloor \frac{n}{3} \rfloor$ guards that guard P.

The cases when P_1 has 6 and 7 vertices can be solved in a similar way, and are left as excercises.

2.2 Input sensitive algorithms

Although Theorem 2.1.1 provides a general upper bound on the number of guards required to guard any polygon with n vertices, most polygons can be guarded with fewer than $\lfloor \frac{n}{3} \rfloor$ guards. It is natural then to ask for the existence of an efficient algorithm to find the minimun number of vertex guards to protect a polygon.

Unfortunately, the existence of such an algorithm is highly unlikely, due to the following result of D. T. Lee and A.K.Lin [86]:

Theorem 2.2.1 *The minimum vertex guard problem for polygons is NP-complete.*

The proof of this result will be studied later. This result gives rise to the study of algorithms that somehow take advantage of the structure of a polygon to obtain a better bound on the number of vertices required to guard them. For example, convex polygons can be guarded with one guard.

A vertex of a polygon P is called *reflex* if the internal angle of P at v is greater than π. We now prove:

Theorem 2.2.2 *Le P be a polygon with r reflex vertices, $r \geq 1$. Then r guards are always sufficient and occasionally necessary to guard P.*

The following result will be useful to prove Theorem 2.2.2.

Lemma 2.2.1 *Any polygon with r reflex vertices can be partitioned into at most $r+1$ convex polygons with disjoint interiors.*

Proof: Let $\{r_1, \ldots, r_k\}$ be the reflex vertices of P. For $i = 1, \ldots, k$ extend a ray starting at r_i that bisects the internal angle of P at r_i until it hits the boundary of P or a previously drawn ray; see Figure 2.2. An easy inductive argument can now be used to show that this set of rays divides P into $r + 1$ convex polygons as desired. ∎

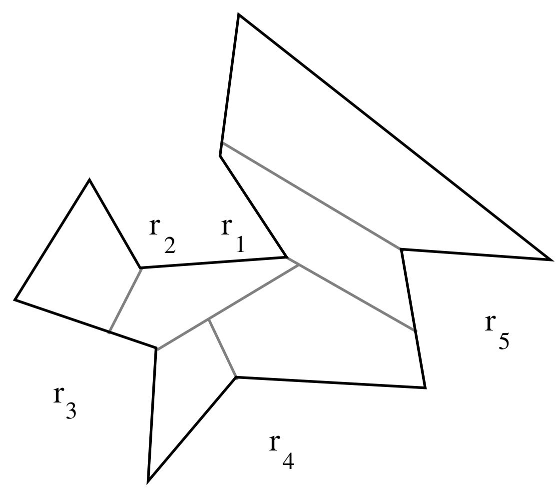

Figure 2.2: Partitioning a polygon into $r + 1$ convex polygons.

We can now prove our result:

Proof of Theorem 2.2.2: We first show that r guards always suffice. Partition P as in the previous lemma. Notice that each of the convex polygons obtained is such that at least one of its vertices is a reflex vertex of P. Our result now follows.

A family of polygons for which r guards are necessary is shown in Figure 2.3. ∎

2.3 Orthogonal Polygons

Of particular interest is the study of guarding problems for orthogonal polygons, that is; polygons whose edges are all parallel to the x- or y-axis. Perhaps one of the motivations for the study of these polygons is that most "real life"

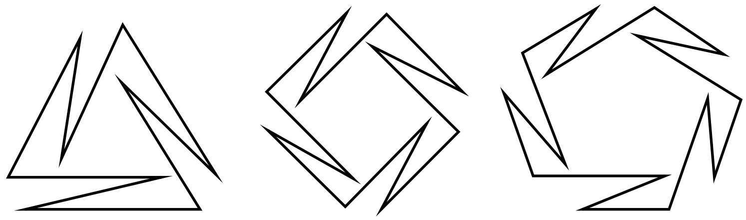

Figure 2.3: A family of polygons that require r guards.

buildings are, after all, "orthogonal". From a mathematical point of view, their inherent structure allows us to obtain very interesting and aesthetic results, and I believe this is an even stronger motivation to study them.

The first major result here is due to Kahn, Klawa and Kleiman [82]. They proved:

Theorem 2.3.1 *Any orthogonal polygon with n vertices can be guarded with at most $\lfloor \frac{n}{4} \rfloor$ guards.*

We will present three proofs for Theorem 2.3.1 The first one is new, and short. We then give two more proofs of the same theorem that although longer, are interesting in their own right, and demonstrate some interesting properties of partitioning of polygons into convex quadrilaterals and L-shaped polygons.

We start by proving the following result due to O'Rourke:

Lemma 2.3.1 *Let P be an orthogonal polygon with n vertices, r of which are reflex. Then $r = \frac{n-4}{2}$.*

Proof: Since P has n vertices, the sum of the internal angles of P is exactly $(n-2)\pi$. Notice that all the internal angles of P are of size $\frac{\pi}{2}$ or $\frac{3\pi}{2}$, depending on whether they are generated by a convex or reflex vertex respectively.

Thus if we have r reflex and $n-r$ convex vertices, we have:

$$(n-r)(\frac{\pi}{2}) + r(\frac{3\pi}{2}) = (n-2)\pi$$

Solving for r gives the desired result. ∎

A horizontal or vertical *cut* of P is the extension of a horizontal or vertical edge of P at a reflex vertex towards the interior of P until it hits the boundary of P. A horizontal or vertical cut is called an *odd cut* if it splits P into two non-empty polygons P_1 and P_2 with n_1 and n_2 vertices respectively such that at least one of n_1 or n_2 equals $4k+2$ for some k.

O'Rourke noted in his proof of the orthogonal art gallery theorem that if we can prove that any orthogonal polygon has an odd cut, then by an inductive argument, Theorem 2.3.1 would follow. Indeed, we can easily verify that if $n_1 + n_2 \leq n+2$, and one of n_1 or n_2 equals $n_1 = 4k+2$ then $\lfloor \frac{n_1}{4} \rfloor + \lfloor \frac{n_2}{4} \rfloor \leq \lfloor \frac{n}{4} \rfloor$.

Given an orthogonal polygon P, we label its edges as *top*, *right*, *bottom*, and *left* in the natural way, and then call a reflex vertex of P a *top-left* reflex vertex if the edges of P incident to it are a top and a left edge. *Top-right*, *bottom-right*, and *bottom-left* edges are defined in a similar way; see Figure 2.4.

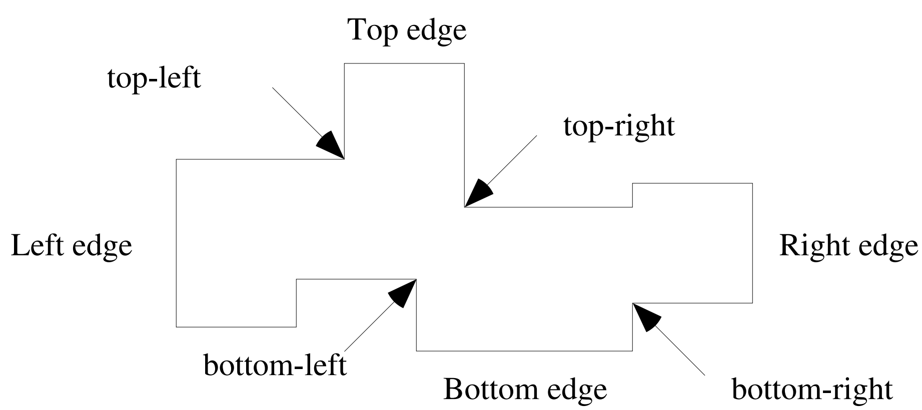

Figure 2.4: Classifying the edges and vertices of an orthogonal polygon.

We are ready to prove Theorem 2.3.1.

Proof: Split the set of reflex vertices of P into two sets, S_1 containing all the top-right and bottom-left reflex vertices of P, and S_2 containing all the top-left and bottom-right vertices. Since P has $\frac{n-4}{2}$ reflex vertices, one of S_1 or S_2 has at most $\lfloor \frac{n}{4} \rfloor$ vertices. Suppose it is S_1. If placing a light at every vertex of S_1 illuminates all of P, we are done. Suppose then that there is a point p in P not illuminated by S_1. Consider the longest horizontal line segment c containing p, and contained in P. Let e and f be the edges of P containing the endpoints of c; see Figure 2.5.

Consider the largest rectangle containing c and contained in P. Let e' and f' be edges of P that intersect the top and bottom edges of R respectively. Since p is not visible from any point in S_1, it follows that e and e' meet at the top-left corner point of R. Similarly f and f' meet at the right-bottom corner point of R.

Let q and q' be the top-right and bottom-left vertices of R. If they are vertices of P, it follows that P is R and there is nothing to prove.

Two cases arise:

1. Neither of q and q' are vertices of P.

2. Exactly one of them, say q, is a vertex of P.

If neither of q and q' are vertices of P, it is easy to see that we can generate two horizontal cuts that generate a rectangle contained in R and containing c as shown in Figure 2.5. It now follows that one of these two horizontal cuts is an odd cut of P.

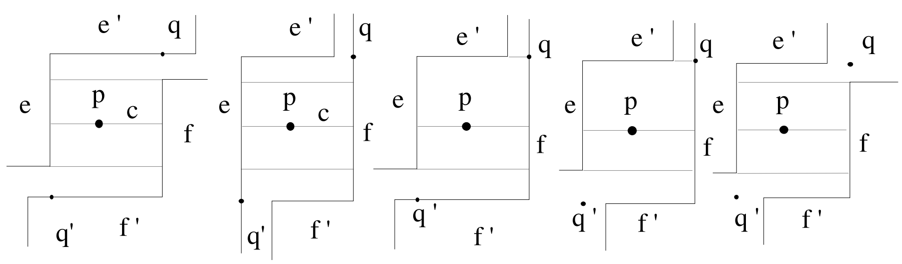

Figure 2.5: Up to symmetry, these are all the cases arising when none of q and q' are vertices of P.

Suppose then that only q is a vertex of P. If e and f' are properly contained in the left and bottom edge of R, then by extending the horizontal edge of P, incident to the bottom vertex of e, and the vertical edge of P incident to the left endpoint of f', we obtain a polygon with $n-4$ vertices that by induction can be guarded with $\lfloor \frac{n-4}{4} \rfloor$ guards. Since R can be guarded with a single guard, our result follows; see Figure 2.6.

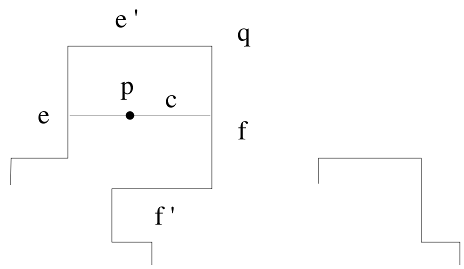

Figure 2.6: The case when e and f' are properly contained on the left and bottom edges of R.

Suppose then without loss of generality that f' properly contains the bottom edge of R and let w be the bottom vertex of e. Consider the horizontal line segment l joining w to a point in the base of R. Slide this segment until it hits a vertical edge of P or it reaches the left endpoint of f' or the leftmost endvertex of horizontal edge g incident to the bottom vertex of e. In the second case, if we reach the left endpoint of f', this point generates a vertical odd cut of P that leaves a polygon with six edges to its right, and another with $n-4$ on the left. In this case our result follows again by induction. The case when we reach the leftmost point of g follows in a similar way.

Suppose then that we hit a vertical edge of P. Let x be the highest vertex of P contained in l. Then we can generate two cuts of P at x, a horizontal cut h and a vertical cut h'. Let P' be the orthogonal subpolygon of P to the left of h' obtained when we cut P along h', and P'', the subpolygon on top of h obtained by cutting P along h. If P' has m vertices, P'' contains $m+2$ vertices, and thus either h or h' is an odd cut; see Figure 2.7. Our result follows. ∎

2.4 Partitioning orthogonal polygons into convex pieces

A convex quadrilaterization of an orthogonal polygon P is a partition of P into a set of convex quadrilaterals with disjoint interiors such that the edges of these quadrilaterals are either edges of P or diagonals joining pairs of

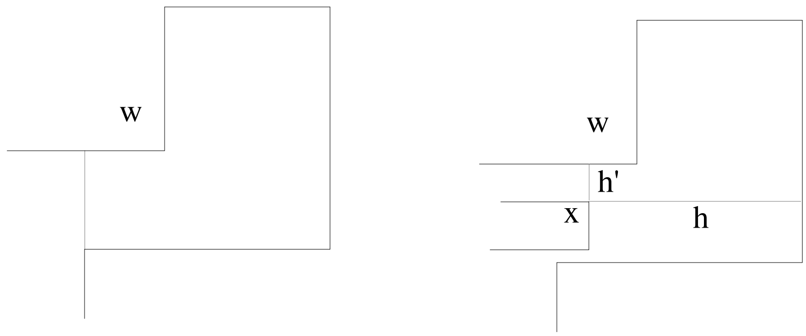

Figure 2.7: The case when f' properly contains the bottom edge of R.

vertices of P; see Figure 2.8.

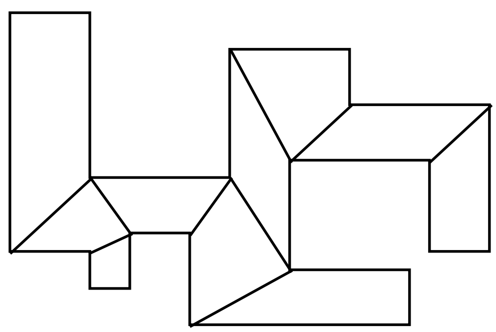

Figure 2.8: A quadrilaterization of an orthogonal polygon.

The proof of Theorem 2.3.1 is based on the following result:

Theorem 2.4.1 *Any orthogonal polygon is convex quadrilaterizable.*

We will return to the proof of this result in the following section; we now prove Theorem 2.3.1.

Proof of Theorem 2.3.1: In Figure 2.9, we show a family of orthogonal polygons that require $\lfloor \frac{n}{4} \rfloor$ guards.

To show that $\lfloor \frac{n}{4} \rfloor$ guards are always sufficient, we first obtain a quadrilaterization Q of P. We next obtain a graph H by adding two diagonals connecting opposite vertices of every quadrilateral of Q, as shown in Figure 2.10.

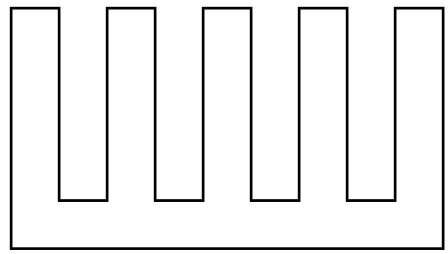

Figure 2.9: An orthogonal polygon that requires $\lfloor \frac{n}{4} \rfloor$ guards.

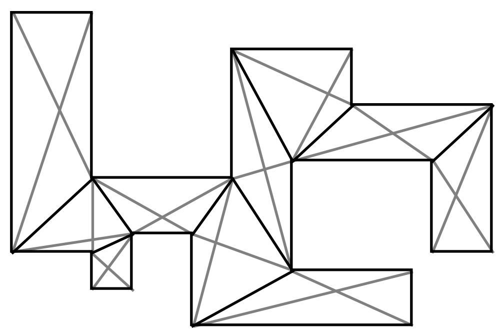

Figure 2.10: A 4-chromatic graph obtained from a quadrilaterization.

We now show that this graph is 4-colorable. To this end, consider the dual graph Q^* of Q. The vertices of Q^* are the quadrilaterals of Q, two of which are adjacent if they share an edge of Q. Clearly Q^* is a tree. Remove from Q a quadrilateral corresponding to a leaf of Q^*. This produces a subgraph H' of H which, by an inductive argument, may be assumed to be 4-colorable. It is easy to see that this coloring can be extended to a 4-coloring of H. Notice that in the resulting 4-coloring, all the vertices of any quadrilateral of Q receive different colors, and thus the vertices of each chromatic class guard P. Place a guard at each vertex of the smallest chromatic class, and our result follows. ∎

We now show that every orthogonal polygon is convex quadrilaterizable. The proof we give here is due to Lubiw [91]. To this end, we define a 1-*orthogonal polygon* P to be a polygon that satisfies the following conditions:

1. All the edges of P, with the possible exception of one distinguished edge denoted by e and called the *slanted edge*, are parallel to the x- or y axis.

2. With the possible exception of e, the edges alternate between horizontal and vertical.

3. All internal angles are less than or equal to $\frac{3\pi}{4}$.

4. The interior of the nose of the slanted edge contains no vertex of P.

5. P has an even number of edges.

Notice that orthogonal polygons are 1-orthogonal. In our definition, we do not explicitly forbid e from being parallel to the x- or y-axis. If P is orthogonal, then e may be chosen to be any of the edges of P. Thus if we show that any 1-orthogonal polygon is convexly quadrilaterizable, Theorem 2.3.1 will be proved.

The *nose* of the slanted edge e of P is the right triangle T towards the interior of P such that:

1. The hypotenuse of T is e.

2. The other two edges of T are parallel to the x- and y-axis.

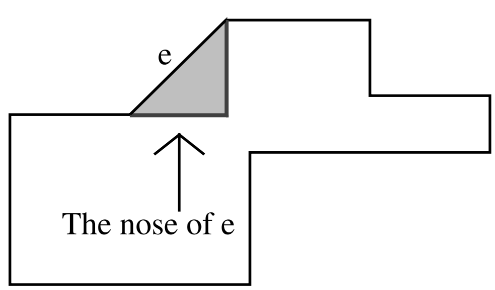

Figure 2.11: A 1-orthogonal polygon.

We now prove:

Theorem 2.4.2 *Any* 1*-orthogonal polygon is convexly quadrilateralizable.*

Proof: If P has four edges, it must be convex, and our result follows. We now show that if P has more than four edges, there is always a quadrilateral Q' whose removal splits P into smaller 1-orthogonal polygons. Using induction our result follows.

Assume without loss of generality that the slope of e lies between 0 and $\frac{\pi}{2}$, and that the vertices of e are labelled u and v such that u is the vertex of e with smallest y-coordinate. Let w be the third vertex of the nose T of e.

Since P has an even number of edges, and all internal angles are of size at most $\frac{3\pi}{2}$, both edges adjacent to e are horizontal. Let e' be the horizontal edge of P that precedes e in the counterclockwise order along the boundary of P. We now proceed to find the four vertices of Q'. The first two vertices of Q' are u and v. Let e" be the semi-open line segment joining v to w, open at w. Slide e" to the right until it hits a vertex of P. If e'' hits a whole edge of P, then the endpoints of this edge, together with u and v are the vertices of Q'. Suppose then that e'' hits a single vertex x of P; x is the third vertex of Q'.

Consider now the semi-open horizontal line segment h joining u to the vertical line through x; h does not contain u, but it contains its right endpoint. Slide it down until it reaches a vertex of P, or it hits a horizontal edge f of P in the middle. In the first case, let y be the rightmost point that h hits. The fourth vertex of Q' is y. This case is illustrated in Figure 2.12.

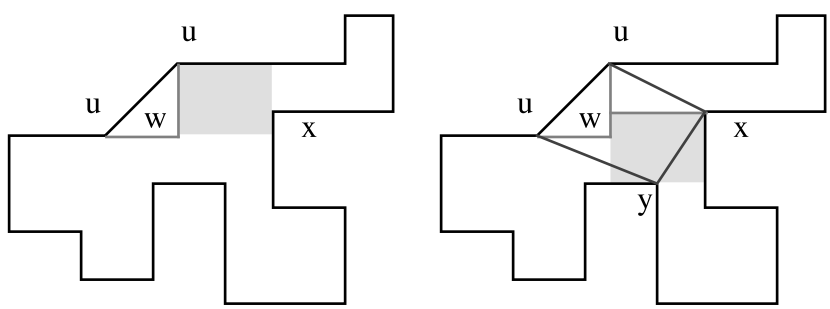

Figure 2.12: Finding Q', case 1.

We now deal with the case when h hits an edge f of P. This can only happen if x is the second end-vertex of e'. This case is illustrated in Figure 2.13. In this case, consider the semi-open vertical line segment h' joining x to f, closed at the bottom and open at the top. Slide h' to the right until it hits a vertex of P. If it hits more than one vertex, choose y, the lowest of these points, to be the last vertex of Q'.

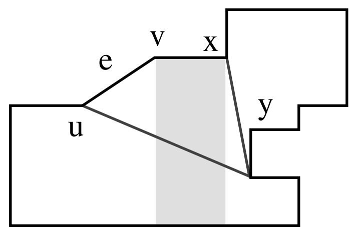

Figure 2.13: Finding Q', case 2.

It is now easy to verify that we can split $P - Q'$ into two or three 1-orthogonal polygons. ■

2.5 Cutting orthogonal polygons into L-shaped pieces

A different proof of Theorem 2.3.1 was obtained by J. O'Rourke [102]. An L-shaped polygon is an orthogonal polygon with six vertices. Notice that any L-shaped polygon can be guarded by a single guard. O'Rourke's main idea is to divide an orthogonal polygon into L-shaped pieces, each of which can be guarded with one guard.

In view of Lemma 2.3.1 O'Rourke rephrases his result in terms of r, the number of reflex vertices of orthogonal polygons. He proves:

Theorem 2.5.1 $\lfloor \frac{r}{2} \rfloor + 1$ *guards are always sufficient and occasionally necessary to guard an orthogonal polygon with r reflex vertices.*

Proof: The proof of Theorem 2.5.1 is by induction on r. If $r = 0$ our result is clearly true. Two cases arise:

1. There is a horizontal or vertical line segment l connecting two reflex vertices of P such that the interior of l is totally contained in the interior of P.

2. No such horizontal or vertical line segment exists. In this case, we say that P is in general position.

Let l be a horizontal line segment as in (1). Then l splits P into two orthogonal polygons Q and R with s and t reflex vertices respectively, such that $s+t=r-2$. By induction, we can guard them with $\lfloor \frac{s}{2} \rfloor + 1$ and $\lfloor \frac{t}{2} \rfloor + 1$ guards respectively. But $(\lfloor \frac{s}{2} \rfloor + 1) + (\lfloor \frac{t}{2} \rfloor + 1) \leq \lfloor \frac{r}{2} \rfloor + 1$ and our result follows.

Some terminology and results will be neede to prove (2). A *cut* of an orthogonal polygon is an extension of one of the edges incident to a reflex vertex of P towards the interior of P until it hits the boundary of P. Notice that a cut through a reflex vertex u "resolves" or eliminates u since it is no longer a reflex vertex in either of the two polygons, say Q and R, into which P is partitioned by the cut. A cut is called an *odd cut* if at least one of Q and R has an odd number of reflex vertices.

To finish the proof of our result, we will use the following result, which will be proved after we finish our proof: *Every orthogonal polygon in general position has an odd-cut.*

Let P be an orthogonal polygon in general position, and consider an odd-cut that splits it into two subpolygons Q and R. Assume without loss of generality that Q has an odd number of reflex vertices. Let s and t be the number of reflex vertices of Q and R respectively. Then by induction, we can guard Q and R with $(\lfloor \frac{s}{2} \rfloor + 1)$ and $(\lfloor \frac{t}{2} \rfloor + 1)$ vertex guards respectively. Notice that $s+t=r-1$, and since s is odd it follows by a simple case analysis that $(\lfloor \frac{s}{2} \rfloor + 1) + (\lfloor \frac{t}{2} \rfloor + 1) \leq \lfloor \frac{r}{2} \rfloor + 1$. Our result now follows. ∎

To finish the proof of Theorem 2.5.1 we now prove that every orthogonal polygon in general position has an odd-cut.

We notice first that the general position condition is essential. The polygon in Figure 2.14(a) has no odd-cuts! Moreover a horizontal odd-cut does not always exist. The polygon in Figure 2.14(b) has no horizontal odd-cut. This complicates things a bit. We divide cuts into horizontal and vertical cuts, which we denote as *H-cuts* and *V-cuts* respectively.

Notice first that if the number of reflex vertices of P is even, then the number of reflex vertices of Q plus those in R equals $r-1$, which is odd, and thus one of Q or R must have an odd number of reflex vertices. Suppose then that P has an odd number of reflex vertices.

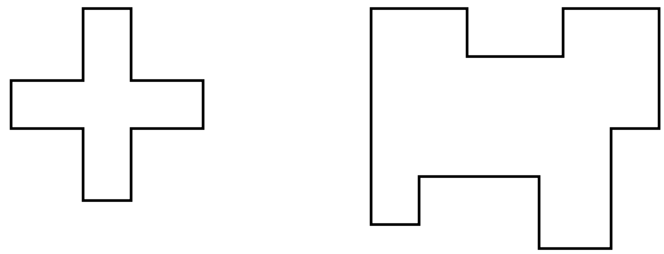

(a) (b)

Figure 2.14: An orthogonal polygon with no odd-cuts, and a polygon with no horizontal cuts.

Two reflex vertices of P are called a *horizontal pair* if they are the end vertices of a horizontal edge of P. A reflex vertex v of P is called *H-isolated* if the second end vertex of the horizontal edge incident at v is convex.

Obtain a partitioning of P by cutting P along an H-cut at each reflex vertex of P that belongs to a horizontal pair; see Figure 2.15. We call this the *H-partitioning* π of P. We now define a H-graph of this partitioning as the oriented graph whose vertex set is the set of regions of this partitioning. Two adjacent regions, X and Y are connected by an edge oriented from X to Y if the H-cut that separates them is an extension of an edge of P on the boundary of X; see Figure 2.15.

We can now classify the vertices of the H-graph of P as follows:

Leaf nodes Vertices with in-degree 1.

Pit nodes Vertices with in-degree 2, and out-degree 0.

Source nodes Vertices with out-degree 2 or 4.

Branch nodes Vertices with out-degree 2 and in-degree 1.

Lemma 2.5.1 *If the H-graph of P contains a pit, then it has an odd-cut.*

Proof: Let K be a pit of P, and let h and h' be the top and lower H-cuts separating it from its neighbours. If K contains no H-isolated vertices, one of h or h' must be an odd-cut. If K has H-isolated vertices and h is an even cut, the highest H-isolated vertex u of K generates an odd-cut, and our result follows; see Figure 2.16. ∎

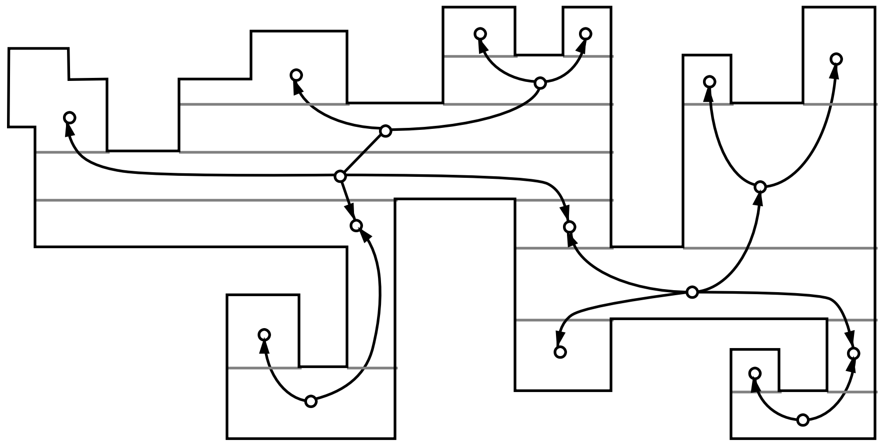

Figure 2.15: The H-partitioning of P and its H-graph.

We now prove:

Lemma 2.5.2 *If P contains no odd-cuts, its H-graph contains a single source, P has exactly one H-isolated vertex, and it is located at the source of its H-graph.*

Proof: Observe that if the H-graph of P contains two sources, then the path joining them must necessarily contain a pit. However since P admits no odd-cuts, its H-graph has no pits. It follows the H-graph of P contains a unique source. Since the H-graph of P is a tree with a unique source, each vertex of it, except for its source, has in-degree 1. Let K be one of these regions of the H-partitioning of P. Then since P admits no odd-cuts, the cut corresponding to the incoming edge of K is an even-cut. Suppose without loss of generality that this edge is a bottom edge of K. Then if K contains H-isolated vertices, the lowest of them generates an odd-cut, which is a contradiction.

However since P contains an odd number of reflex vertices (or any H-cut would be an odd-cut) it must contain at least one H-isolated vertex, and it must be a vertex of the source region of the H-partitioning of P. Moreover, notice that if the source region contains more than one H-isolated vertex, one of them would generate an odd-cut, which is a contradiction. Our result follows. ∎

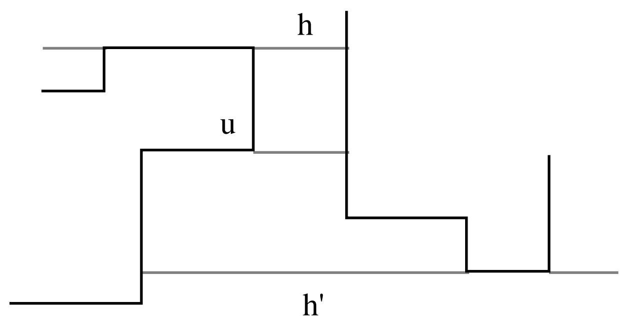

Figure 2.16: Finding an odd-cut in a pit.

We notice that all the results proved for H-cuts hold for V-cuts, and in particular we have that if P admits no vertical odd-cuts, it has exactly one V-isolated vertex located at the source face of the V-graph of P.

We now prove:

Lemma 2.5.3 *Any orthogonal polygon P in general position with an odd number r of vertices admits an odd-cut,* $3 \leq r$.

Proof: Since P admits no odd-cuts, it has exactly one H-isolated vertex u and one V-isolated vertex v. Moreover these vertices are in the sources of the H-graph and the V-graph of P. Any other reflex vertex of P belongs to an H- and a V-pair. Therefore all the reflex vertices of P are located in a chain of reflex vertices that starts, say, in u and ends in v; see Figure 2.17. This implies that both u and v are located in leaf regions of the H- and V-partitions of P, which is a contradiction. ∎

This concludes our proof of Theorem 2.5.1. ∎

2.6 Algorithms for L-shaped partitions

In this section we develop an $O(n \ln n)$ algorithm to partition an orthogonal polygon into L-shaped pieces. Using Chazelle's linear time triangulation algorithm, a linear time implementation of our algorithm can be obtained, however this is outside the scope of our book.

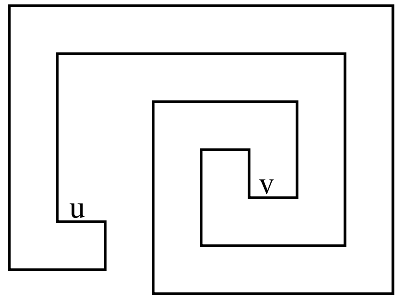

Figure 2.17: A spiral orthogonal polygon.

An orthogonal polygon Q is called a vertical *histogram* if it has a horizontal edge, called the *base* of Q, such that every point of Q is visible from some point in its base; see Figure 2.18.

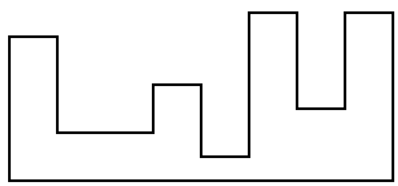

Figure 2.18: A histogram.

Lemma 2.6.1 *If an orthogonal polygon has no horizontal odd-cuts, it can be partitioned into* $\lfloor \frac{r}{2} \rfloor + 1$ *L-shaped pieces.*

Proof : Observe that by Lemma 2.5.2 the H-graph of any such polygon has exactly one source, and one H-isolated vertex located in this face. It is easy to see that this cut partitions P into two histograms. Our result now follows by partitioning each of these histograms by cutting them along a vertical line through every second reflex vertex. ■

We now show how to obtain a partitioning of P into L-shaped pieces. To start, assume that P has no horizontal cuts that join two reflex vertices. This case can be easily dealt with. Furthermore we will assume that P contains an odd number of reflex vertices. The case when P has an even number of reflex vertices will be dealt with at the end of our proof for the odd case. Our approach is justified by the following observation: an odd-cut of a polygon with an odd number of reflex vertices splits it into two subpolygons, each with an odd number of reflex vertices. Moreover we observe that if we cut P along an odd-cut h, and h' is an odd-cut of P different from h, then h' is also an odd-cut of the subpolygon of P that contains it.

We now present our algorithm to find an L-partitioning of a polygon P with an odd number of reflex vertices:

Algorithm 2.6.1

1. *Find all the horizontal cuts of P.*
2. *Find all the horizontal odd-cuts.*
3. *Cut P at each horizontal odd-cut.*
4. *Each resulting piece has no horizontal odd-cut, i.e it is a histogram. Partition it by cutting it along a vertical through every second reflex vertex.*

We notice now that each subpolygon P_i of P obtained in Step 3 of our algorithm has no odd-cuts, and thus by Lemma 2.6.1 can be partitioned into at most $\lfloor \frac{r_i}{2} \rfloor - 1$ L-shaped pieces, where r_i is the number of reflex vertices of P_i. It now follows that our algorithm indeed produces a partitioning of P as desired.

Finding the odd-cuts of P can be done in $O(n \ln n)$ time by a horizontal line sweep. See Appendix.

Determining which cuts are odd-cuts can be accomplished as follows: Pick any reflex vertex of P and label it 1. Set a reflex counter of vertices of P to 1. We now *walk* along the boundary of P. Each time we encounter a reflex vertex of P we increase our counter by one, and label the vertex with the new value of our counter. Each time we encounter the end point of a horizontal cut that is not a vertex of P we simply label it with the value of our counter. A horizontal cut is even if the absolute value of the difference of the labels of its end points is odd. This can clearly be done in linear time.

Step 3 can also be done in linear time. It is obvious that this partition yields at most $\lfloor \frac{r}{2} \rfloor + 1$ L-shaped polygons. Thus we have proved:

Theorem 2.6.1 *An orthogonal polygon with an odd number of reflex vertices can be partitioned into at most* $\lfloor \frac{r}{2} \rfloor + 1$ *L-shaped polygons in* $O(n \ln n)$ *time.*

If P has an even number r of reflex vertices, we first create an extra reflex vertex of P as follows: let e be the lowest horizontal edge of P and u its leftmost vertex. Since we are assuming that P is in general position, all its vertices have different x- and y-coordinares. Let ϵ be such that the horizontal and vertical distance between any two vertices of P is at least ϵ. Cut away from P a square with sides parallel to the coordinate axes, such that the length of its edges is ϵ and u is its lowest leftmost vertex. Let P' be the resulting orthogonal polygon. P' has exactly $r+1$ reflex vertices. The following lemma, left as an exercise to the reader, can now be proved:

Lemma 2.6.2 *Any L-shaped partitioning of* P' *into at most* $\lfloor \frac{r+1}{2} \rfloor + 1$ *pieces can be modified to obtain a similar partitioning of* P *into* $\lfloor \frac{r}{2} \rfloor + 1$ *pieces.*

However since r is even, $\lfloor \frac{r+1}{2} \rfloor + 1 = \lfloor \frac{r}{2} \rfloor + 1$ and we have proved:

Theorem 2.6.2 *An orthogonal polygon with* r *reflex vertices can be partitioned into at most* $\lfloor \frac{r}{2} \rfloor + 1$ *L-shaped polygons in* $O(n \ln n)$ *time.*

2.7 Exercises

Ex. 2.1 Provide all the details of a linear time implementation of Algorithm 2.1.1.

Ex. 2.2 Prove Lemma 2.1.1. Hint: use the fact that the dual graph of a triangulation of a polygon is a tree in which the maximum degree of each vertex is 3.

Ex. 2.3 Finish Chvátal's proof of Theorem 2.1.1.

Ex. 2.4 Verify that $P - Q'$ can indeed be split into two or three 1-orthogonal polygons; i.e. verify that the resulting polygons have an even number of edges and their noses are empty.

Ex. 2.5 Prove Lemma 2.6.2.