3

MOBILE GUARDS

3.1. INTRODUCTION

In this chapter we explore an interesting variant of the art gallery problem suggested by Toussaint. Rather than modify the shape of the polygons as in the previous chapter, we modify the power of the guard. Specifically, each guard is permitted to "patrol" an interior line segment. Let s be a line segment completely contained in the closed polygonal region P: $s \subseteq P$. Then $x \in P$ is said to be seen by s, or is covered by s, if there is a point $y \in s$ such that the line segment $xy \subseteq P$. Thus x is covered by the guard if x is visible from some point along the guard's patrol path. This is the notion of *weak visibility* from a line segment introduced in Avis and Toussaint (1981b) (*strong visibility* requires x to be seen from every point of s), a concept further explored in Chapter 8.

The main reason that mobile guards are interesting is that they lead to some clean theorems, some difficult theorems, and to interesting open problems. Secondarily they connect to the important notion of edge visibility, to be discussed further in Chapters 7 and 8. A covering by mobile guards induces a partition into edge-visible polygons.

We present two long proofs in this chapter. The first establishes that $\lfloor n/4 \rfloor$ mobile guards are occasionally necessary and always sufficient to cover an n vertex polygon. The second proof, which is quite complex, establishes the equivalent result for orthogonal polygons: $\lfloor (3n+4)/16 \rfloor$ mobile guards are necessary and sufficient. This latter quantity may seem ugly in comparison to the simpler fractions we have encountered so far, but there is a clean logic behind it, as revealed in Table 3.1. Mobile guards are more powerful than stationary guards: only 3/4's as many are needed, in both general and orthogonal polygons—the second column is 3/4 times the first. Moreover, orthogonal polygons are 3/4's easier to cover than general polygons: the second row is 3/4 times the first. Thus orthogonal polygons require about $(3/4)^2 \lfloor n/3 \rfloor$ mobile guards.

The first proof, presented in Section 3.2, is entirely combinatoric, following the outline of Chvátal's proof (Section 1.2.1). The second proof,

Table 3.1

Guard→ Shape	Stationary	Mobile
General	$\lfloor n/3 \rfloor$	$\lfloor n/4 \rfloor$
Orthogonal	$\lfloor n/4 \rfloor$	$\lfloor (3n+4)/16 \rfloor$

presented in Section 3.3, is an instance where no reduction to combinatorics has been discovered, and complex geometric reasoning seems necessary. The chapter closes with a discussion of related results.

3.2. GENERAL POLYGONS[1]

We first define various types of guards, both geometric and combinatorial. Three geometric mobile guards types with different degrees of patrol freedom can be distinguished. An *edge guard* is an edge of P, including the endpoints. A *diagonal guard* is an edge or internal diagonal between vertices of P, again including the endpoints. A *line guard* is any line segment wholly contained in P. (Recall that P is a closed region.) Geometric guards are said to *cover* the region they can see.

The combinatorial counterparts of these guards are obtained by defining a *guard* in a triangulation graph T of a polygon P to be a subset of the nodes of T. Then a *vertex guard* in T is a single node of T, an *edge guard* is a pair of nodes adjacent across an arc corresponding to an edge of P, and a *diagonal guard* is a pair of nodes adjacent across any arc of T. The analog of covering is domination: a collection of guards $C = \{g_1, \ldots, g_k\}$ is said to *dominate* T if every triangular face of T has at least one of its three nodes in some $g_i \in C$.

The goal of this section is to prove that $\lfloor n/4 \rfloor$ combinatorial diagonal guards are sometimes necessary and always sufficient to dominate the triangulation graph of a polygon with $n \geq 4$ vertices. It is clear that if a triangulation graph of a polygon can be dominated by k combinatorial vertex guards, then the polygon can be covered by k geometric vertex guards. The implication is that a proof of the sufficiency of a $\lfloor n/4 \rfloor$ of combinatorial diagonal guards in a triangulation graph establishes the sufficiency of the same number of geometric diagonal and line guards in a polygonal region.

Necessity is established by the generic example due to Toussaint shown in Fig. 3.1: each 4 edge lobe requires its own diagonal guard.

1. An earlier version of this section appeared in O'Rourke (1983a), © 1983 D. Reidel.

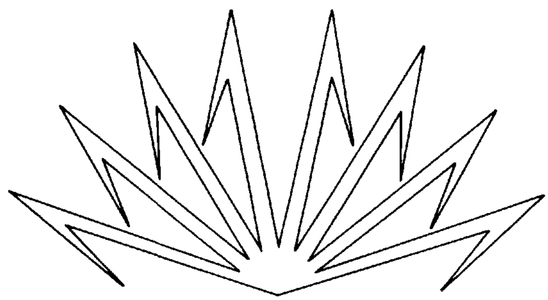

Fig. 3.1. A polygon that requires $\lfloor n/4 \rfloor$ edge, diagonal, or line guards.

3.2.1. Sufficiency Proof

The proof is by induction and follows the main outlines of Chvátal's inductive proof (and Honsberger's exposition (Honsberger 1976)). Before commencing the proof, it will be convenient to establish certain facts that will be used in various cases of the proof. The most important of these concerns "edge contractions." Let P be a polygon and T a triangulation graph for P, and let e be an edge of P, and u and v the two nodes of T corresponding to the endpoints of e. The *contraction* of e is a transformation that alters T by removing nodes u and v and replacing them with a new node x adjacent to every node to which u or v was adjacent.[2] Compare Figs. 3.2a and 3.2d. Note that an edge contraction is a *graph* transformation, not a polygon transformation: the geometric equivalent ("squashing" the polygon edge) could result in self-crossing polygons. Edge contractions are nevertheless useful because of the following lemma.

LEMMA 3.1. Let T be a triangulation graph of a polygon P, and T' the graph resulting from an edge contraction of T. Then T' is a triangulation graph of some polygon P'.

Proof. We construct a figure with curved edges corresponding to T', then straighten the edges to obtain P'.

Let P_t be the planar figure corresponding to the triangulation T, and let e be the edge contracted and u and v its two endpoints in P_t. Let the vertices to which u and v are connected by diagonals and edges be $y_0, \ldots, y_i$ and $z_0, \ldots, z_j$, respectively, with $y_0 = v$ and $z_0 = u$, and the remainder labeled according to their sorted angular order. See Fig. 3.2a. Note that $y_1 = z_1$ is the apex of the triangle supported by e.

Now introduce a new vertex x on the interior of e, and connect the y and z vertices to x by the following procedure. Connect y_1 to x; this can be done without crossing any diagonals because y_1 is the apex of a triangle on whose base x lies. Remove the diagonal (u, y_1). Connect y_2 to x within the region

2. Harary calls this transformation an *elementary contraction* (Harary 1969).

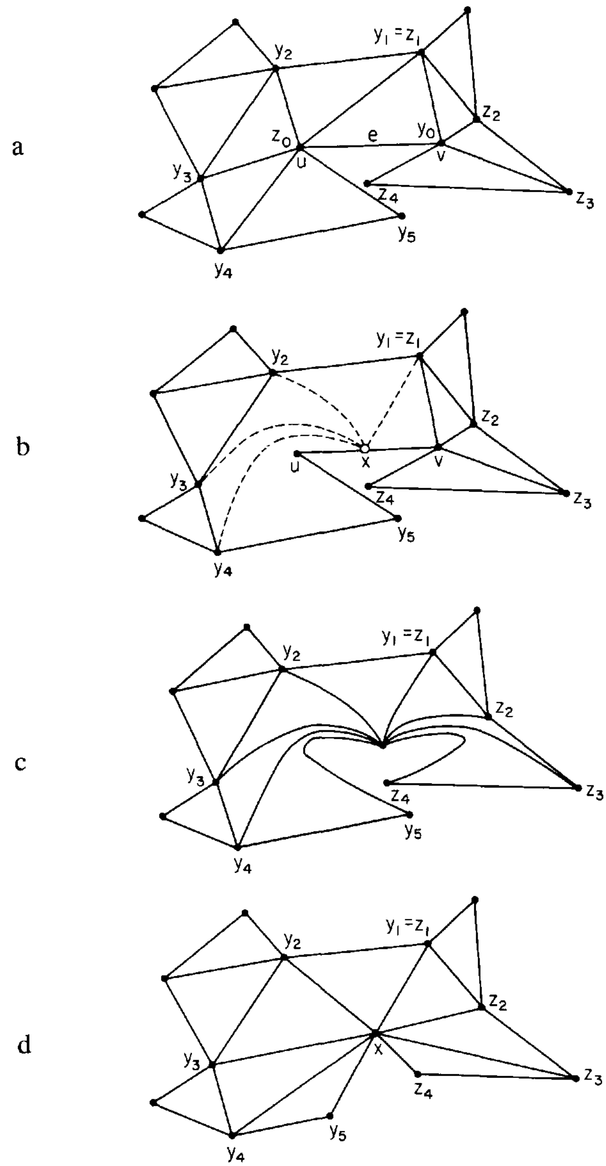

Fig. 3.2. If all the arcs in a triangulation graph (a) incident to u and v are made adjacent to x (b and c), the resulting graph may be deformed into a straight line graph (d).

bounded by (x, y_1, y_2, u); the line may need to be curved but again no crossings are necessary. Remove the diagonal (u, y_2). Continue in this manner (see Fig. 3.2b) until all the y's have been connected to x. Then apply a similar procedure to the z vertices. The result is a planar figure whose connections are the same as those of T'. See Fig. 3.2c.

Finally, apply Fáry's theorem (Giblin 1977): for any planar graph drawn in the plane, perhaps with curved lines, there is a homeomorphism[3] in the plane onto a straight-line graph such that vertices are mapped to vertices and edges to edges. Applying such a homeomorphism to the figure constructed above yields P', a polygon that has T' as one of its triangulations. See Fig. 3.2d. □

The main use of this contraction result is the following.

LEMMA 3.2 Suppose that $f(n)$ combinatorial diagonal guards are always sufficient to dominate any n-node triangulation graph. Then if T is an arbitrary triangulation graph of polygon P with one vertex guard placed at any one of its n nodes, then an additional $f(n-1)$ diagonal guards are sufficient to dominate T.

Proof. Let u be the node at which the one guard is placed, and let v be a node adjacent to u across an arc corresponding to an edge e of P. Edge contract T across e, producing the graph T' of $n-1$ nodes. By Lemma 3.1 T' is a triangulation graph, and so can be dominated by $f(n-1)$ diagonal guards. Let x be the node of T' that replaced u and v. Suppose that no guard is placed at x in the domination of T'. Then the same guard placements will dominate T, since the given guard at u dominates the triangle supported by e, and the remaining triangles of T have dominated counterparts in T'. Again compare Figs. 3.2a and 3.2d. If a guard is used at x in the domination of T', then this guard can be assigned to v in T, with the remaining guards maintaining their position. Again every triangle of T is dominated. □

We note in passing that the same lemma holds for other types of guards, but we will only need to use it with diagonal guards. Intuitively, one can view this lemma as saying that one edge can be "squashed" out for guard coverage calculations if a guard is assigned to either of the edge's endpoints.

The next three lemmas establish special diagonal guard results for small triangulation graphs.

LEMMA 3.3. Every triangulation graph of a pentagon ($n = 5$) can be dominated by a single combinatorial diagonal guard with one endpoint at any selected node.

Proof. Let T be a triangulation graph of a pentagon, and let the selected node be labeled 1. It is easy to show that there are only five distinct triangulations. In each case, a single combinatorial diagonal guard (pair of adjacent nodes), with one end at node 1 can dominate the graph (see Fig. 3.3). □

LEMMA 3.4 Every triangulation graph of a septagon ($n = 7$) can be dominated by a single combinatorial diagonal guard.

3. A *homeomorphism* is a continuous one-one onto mapping whose inverse is also continuous; intuitively it is a deformation without tearing or pasting—that is, it preserves topological properties.

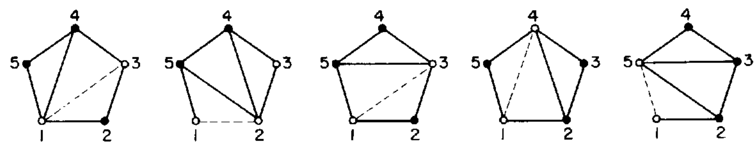

Fig. 3.3. A pentagon can be dominated by a single diagonal guard (shown dashed) with one end at node 1.

Proof. Let T be a triangulation graph of a septagon, and let d be an arbitrary internal diagonal. This diagonal partitions the seven boundary edges of T according to either $2+5=7$ or $3+4=7$; clearly the partition $1+6=7$ is not possible.

Case 1 $(2+5=7)$. Let $d=(1,3)$. Then d supports another triangle T, either $(1,3,4)$, $(1,3,5)$, $(1,3,6)$, or $(1,3,7)$. Only two of these cases are distinct.

Case 1a $(T=(1,3,4))$. Then $(1,4,5,6,7)$ is a pentagon (see Fig. 3.4a). By Lemma 3.3, this pentagon can be covered with a single diagonal guard with one end of node 1. This guard dominates the entire graph.

Case 1b $(T=(1,3,5))$. Choose diagonal $(1,5)$ for the guard (see Fig. 3.4b). Regardless of how the quadrilateral $(1,5,6,7)$ is triangulated, all of T is dominated.

Case 2 $(3+4=7)$. Let $d=(1,4)$. Then both ways of triangulating the quadrilateral $(1,2,3,4)$ lead to situations equivalent to Case 1a above. □

LEMMA 3.5 Every triangulation graph of an enneagon ($n=9$) can be dominated by two combinatorial diagonal guards such that one of their endpoints coincides with any selected node.

Proof. Let T be a triangulation graph of an enneagon, let the selected node be labeled 1, and let d be any internal diagonal with one end at 1. This diagonal partitions the boundary edges of T according to either $2+7=9$, $3+6=9$, or $4+5=9$.

Case 1 $(2+7=9)$. Let $d=(1,3)$. The diagonal d supports another triangle T whose apex is at either 4, 5, 6, 7, 8, or 9. Only three of these cases are distinct.

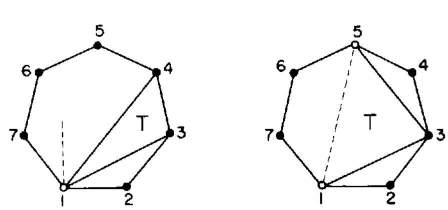

Fig. 3.4. A septagon can be dominated by a single diagonal guard.

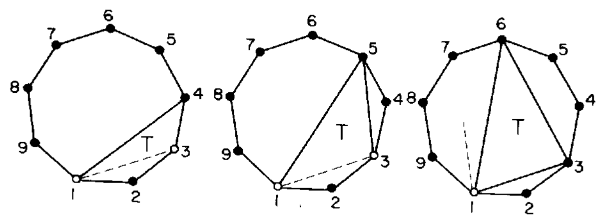

Fig. 3.5. A enneagon can be dominated by two diagonal guards, with one of their ends at node 1.

Case 1a ($T = (1, 3, 4)$). Dominate the septagon $(1, 4, 5, 6, 7, 8, 9)$ with one guard by Lemma 3.4, and use $(1, 3)$ for the second guard (see Fig. 3.5a).

Case 1b ($T = (1, 3, 5)$). Dominate the septagon $(1, 3, 5, 6, 7, 8, 9)$ with one guard by Lemma 3.4, and use $(1, 3)$ for the second guard (see Fig. 3.5b).

Case 1c ($T = (1, 3, 6)$). Dominate the hexagon $(1, 2, 3, 4, 5, 6)$ with one guard by Lemma 3.4, and dominate the pentagon $(1, 6, 7, 8, 9)$ with one guard whose endpoint is at 1 by Lemma 3.3 (see Fig. 3.5c).

Case 2 $(3 + 6 = 9)$. Let $d = (1, 4)$. If diagonal $(1, 3)$ is present, then we have exactly Case 1a above. Otherwise diagonal $(2, 4)$ is present, and one guard along $(1, 2)$ together with a guard for the septagon as in Case 1a suffices.

Case 3 $(4 + 5 = 9)$. Let $d = (1, 6)$. This is equivalent to Case 1c above. □

Finally we must establish the existence of a special diagonal that will allow us to take the induction step, just as Lemma 1.1 did for Chvátal's proof.

LEMMA 3.6. Let P be a polygon of $n \geq 10$ vertices, and T a triangulation graph of P. There exists a diagonal d in T that partitions T into two pieces, one of which contains $k = 5, 6, 7$, or 8 arcs corresponding to edges of P.

Proof. Choose d to be a diagonal of T that separates off a minimum number of polygon edges that is at least 5. Let $k \geq 5$ be this minimum number, and label the vertices $0, 1, \ldots, n-1$ such that d is $(0, k)$. See Fig. 3.6. The diagonal d supports a triangle T whose apex is at t, $0 \leq t \leq k$. Since k is minimal, $t \leq 4$ and $k - t \leq 4$. Adding these two inequalities yields $k \leq 8$. □

With the preceding lemmas available, the induction proof is a nearly straightforward enumeration of cases.

THEOREM 3.1 [O'Rourke 1983]. Every triangulation graph T of a polygon of $n \geq 4$ vertices can be dominated by $\lfloor n/4 \rfloor$ combinatorial diagonal guards.

Proof. Lemmas 3.3, 3.4, and 3.5 establish the truth of the theorem for $5 \leq n \leq 9$, so assume that $n \geq 10$, and that the theorem holds for all $n' < n$. Lemma 3.6 guarantees the existence of a diagonal d that partitions T into

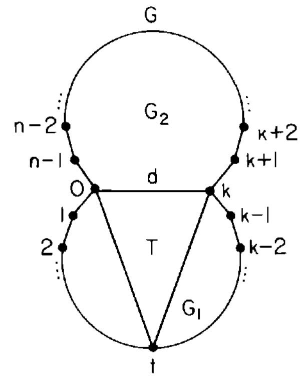

Fig. 3.6. The diagonal d separates G into two pieces, one of which (G_1) shares $5 \le k \le 8$ edges with G.

two graphs T_1 and T_2 where T_1 contains k boundary edges of T with $4 \le k \le 8$. Each value of k will be considered in turn.

Case 1 ($k = 5$ *or* 6). T_1 has $k + 1 \le 7$ boundary edges including d. By Lemma 3.4, T_1 can be dominated with a single diagonal guard. T_2 has $n - k + 1 \le n - 5 + 1 = n - 4$ boundary edges including d, and by the induction hypothesis, it can be dominated with $\lfloor (n-4)/4 \rfloor = \lfloor n/4 \rfloor - 1$ diagonal guards. Thus T_1 and T_2 together can be dominated by $\lfloor n/4 \rfloor$ diagonal guards.

Case 2 ($k = 7$). The presence of any of the diagonals $(0, 6)$, $(0, 5)$, $(1, 7)$, or $(2, 7)$ would violate the minimality of k. Consequently, the triangle T in T_1 that is bounded by d is either $(0, 3, 7)$ or $(0, 4, 7)$; since these are equivalent cases, suppose that T is $(0, 3, 7)$. The quadrilateral $(0, 1, 2, 3)$ has two distinct triangulations. Each will be considered separately.

Case 2a ($(1, 3)$ is included.). Dominate the pentagon $(3, 4, 5, 6, 7)$ with one diagonal guard with one end at node 3. This is possible by Lemma 3.3. This guard dominates all of T_1. Since T_2 has $n - 7 + 1 = n - 6$ boundary edges, it can be dominated by $\lfloor (n-6)/4 \rfloor \le \lfloor n/4 \rfloor - 1$ diagonal guards by the induction hypothesis. This yields a domination of T by $\lfloor n/4 \rfloor$ diagonal guards.

Case 2b ($(0, 2)$ is included.). Form graph T_0 by adjoining the two triangles $T = (0, 3, 7)$ and $T' = (0, 2, 3)$ to T_2 (see Fig. 3.7). T_0 has $n - 7 + 3 = n - 4$ edges, and so can be dominated by $\lfloor (n-4)/4 \rfloor = \lfloor n/4 \rfloor - 1$ diagonal guards by the induction hypothesis. In such a domination, at least one of the vertices of $T' = (0, 2, 3)$ must be a diagonal guard endpoint. There are three possibilities:

(0) If node 0 is a guard end, then T_0 can be extended to include $(0, 1, 2)$ without need of further guards.

(2) If node 2 is a guard end, then T_0 can again be extended to include $(0, 1, 2)$.

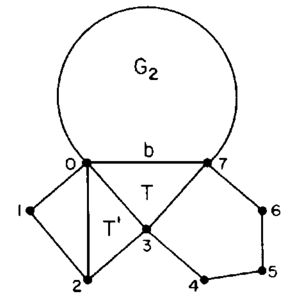

Fig. 3.7. G_0 is formed by adding T and T' to G_2.

(3) If node 3 is a guard end, then there are three possible locations for the other end of the guard. If the other end is at either node 0 or 2, then we fall into the two cases above. If the other end is at node 7, then replace the diagonal guard $(3, 7)$ with $(0, 7)$. Every triangle that was previously dominated is still dominated, and again T_0 can be extended to included $(0, 1, 2)$.

Thus all but the pentagon $(3, 4, 5, 6, 7)$ can be dominated with $\lfloor n/4 \rfloor - 1$ diagonal guards, and the pentagon only requires a single diagonal guard by Lemma 3.4, resulting in a total of $\lfloor n/4 \rfloor$ diagonal guards for all of T.

Case 3 ($k = 8$). T_1 has $k + 1 = 9$ boundary edges, and so by Lemma 3.5, it can be dominated with two diagonal guards, one of whose endpoints is at node 0. Now T_2 has $n - k + 1 = n - 7$ boundary edges. By Lemma 3.2, the one guard at node 0 permits the remainder of T_2 to be dominated by $f(n - 7 - 1) = f(n - 8)$ diagonal guards, where the function $f(n')$ specifies a number of diagonal guards that are always sufficient to dominate a triangulation graph of n' nodes. By the induction hypothesis, $f(n') = \lfloor n'/4 \rfloor$. Therefore, $\lfloor (n - 8)/4 \rfloor = \lfloor n/4 \rfloor - 2$ diagonal guards suffice to dominate T_2. Together with the two allocated to T_1, all of T is dominated by $\lfloor n/4 \rfloor$ diagonal guards. □

COROLLARY. Any polygon P of $n \geq 4$ edges can be covered by $\lfloor n/4 \rfloor$ geometric diagonal or line guards.

Proof. The diagonal guard result follows immediately from the theorem. Since diagonal guards are special cases of line guards, the same number of these more powerful guards clearly suffice. □

3.2.2. Edge Guards

The above proof depends on the fortunate identity between the number of combinatorial and geometric diagonal guards necessary and sufficient to dominate and cover triangulation graphs and polygons, respectively. This identity is not known to hold for edge guards, however. No polygons are known to need more than $\lfloor (n + 1)/4 \rfloor$ *geometric* edge guards (see Fig. 3.8),

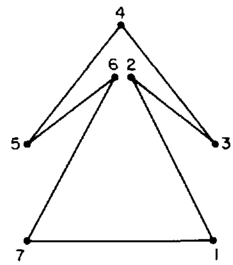

Fig. 3.8. A polygon of seven edges that requires two edge guards.

but triangulation graphs exist that require $\lfloor 2n/7 \rfloor = \lfloor n/3.5 \rfloor$ *combinatorial* edge guards (see Fig. 3.9). Thus it appears that a different proof technique is required in this case.

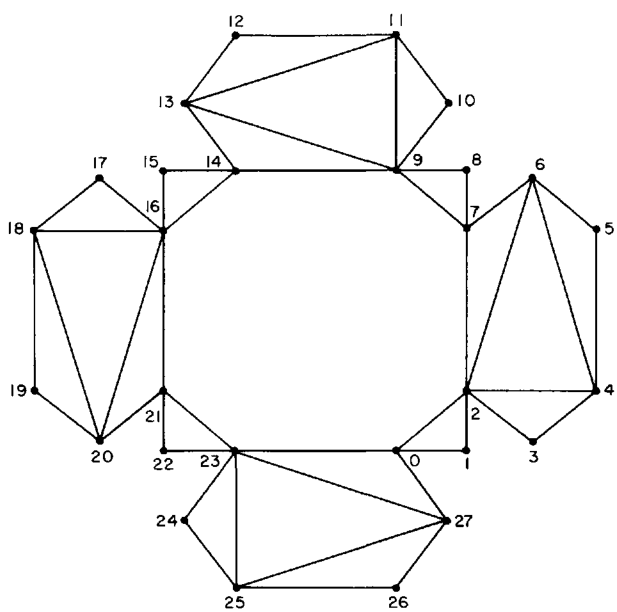

Fig. 3.9. A triangulation graph that requires two edge guards per seven edges. The central octagon may be triangulated arbitrarily.

3.3. ORTHOGONAL POLYGONS

In this section we present Aggarwal's proof that $\lfloor (3n+4)/16 \rfloor$ mobile guards are sufficient for covering an n vertex simple orthogonal polygon (Aggarwal 1984). The occasional necessity of this number of mobile guards is established by a connected series of swastika-like polygons, as shown in Fig. 3.10. The single swasktika shown in Fig. 3.10a with $n = 20$ requires four guards, one per arm; note that $(3 \cdot 20 + 4)/16 = 64/16 = 4$. Merging two

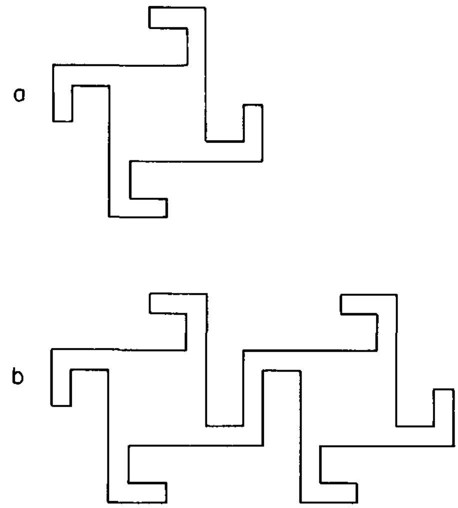

Fig. 3.10. Polygons that require $\lfloor (3n+4)/16 \rfloor$ mobile orthogonal guards: (a) $n = 20$ and $g = 4$; (b) $n = 36$ and $g = 7$.

20-vertex swastika's together removes four vertices at the join, yielding $n = 36$, as in Fig. 3.10b. This polygon requires seven guards, one for each of the six isolated arms, and one at the join; note that $(3 \cdot 36 + 4)/16 = 112/16 = 7$. Joining k swastikas results in an $n = 16k + 4$ vertex polygons that requires $3k + 1$ guards; and note that $\lfloor (3n+4)/16 \rfloor = 3k + 1$. The necessity for other values of n is established by attaching a spiral of the appropriate number of edges to one arm of a swastika. Figure 3.11 shows that a spiral addition of 6 edges requires one guard more than the swastika; a spiral addition of 12 edges requires two guards more. These are the critical additions; spirals with a different number of edges do not require a different number of guards.

This establishes the necessity of $\lfloor (3n+4)/16 \rfloor$ guards. We now turn to sufficiency. Aggarwal's proof is at least superficially similar in structure to the proof for general polygons in the preceding section. The proof is by induction. A small number of quadrilaterals are cut off from the given polygon, these small number covered separately, and the remainder of the polygon handled recursively. The difficulties arise at the interface between

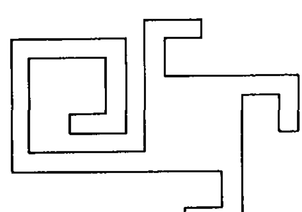

Fig. 3.11. Addition of a spiral establishes necessity for other values of n.

the quadrilaterals cut off and the remainder. In the previous section, interfacing required choosing the diagonal guard with one end at the interface, and applying the "edge-squashing" lemma (3.2) to reduce the number required in the remainder; in effect a guard is shared across the interface. In Aggarwal's proof, the delicacy of the interface requires a complex strategy to complete the induction proof.

Besides the increased complexity, the proof differs in two additional aspects from that of Theorem 3.1. First, it uses geometric constructions throughout, as opposed to reducing the geometric problem to a purely combinatorial one. It is unclear if this is essential; this point will be revisited in Section 3.4. Second, the remainder of the polygon is often modified and needs to be requadrilateralized. The proof of Theorem 3.1 maintained the same triangulation throughout. The combination of these differences result in a unique and complicated proof. It remains to be seen if a simpler approach can establish the same result.

Before commencing with the details, it may be helpful to sketch the main outline of the proof. It will be shown below in Lemma 3.8 that there is always a diagonal d in any quadrilateralization of an orthogonal polygon that cuts off 2, 3, or 4 quadrilaterals. If four quadrilaterals are cut off by d, then properties of quadrilateralizations of orthogonal polygons permit only two essentially different cases, and the induction carries through with a bit of sharing in the vicinity of d. If three quadrilaterals are cut off by d, then there are five distinct cases to handle, only one of which requires extensive sharing at the interface. Finally, if two quadrilaterals are cut off by d, then there are seven cases, most of which require sharing, some rather complicated. All the sharing is accomplished through one complex lemma (3.21). In all cases it will be shown that applying the induction hypothesis to the remainder of the polygon, taking into account any interface sharing, results in $\lfloor (3n+4)/16 \rfloor$ guards. We assume throughout that the polygon is in "general position" in that no two vertices can be connected by a vertical or horizontal line that does not intersect the boundary of the polygon.

We will first discuss structural properties of orthogonal polygons that will be used throughout the remainder of the section. Then we will establish the lemmas used to share at the interface, and finally prove the theorem.

3.3.1. Properties of Orthogonal Polygons

We will conduct the argument in terms of the number of quadrilaterals q in a quadrilateralization of the polygon rather than in terms of the number of vertices n. Our first two lemmas relate these quantities.

LEMMA 3.7. For any quadrilateralization of an orthogonal polygon of n vertices into q quadrilaterals, $n = 2q + 2$.

Proof. The sum of the interior angles of an orthogonal polygon of n vertices is $180(n-2)$ degrees. But since there are q quadrilaterals, each of 360 degrees, $360q = 180(n-2)$, or $n = 2q + 2$. □

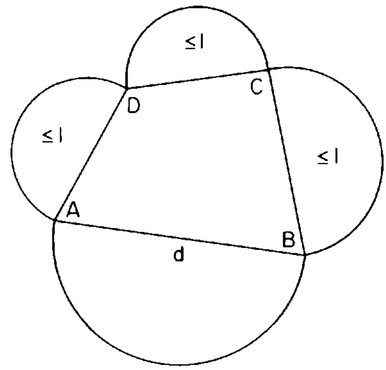

Fig. 3.12. Diagonal d cuts off a minimum number of quadrilaterals that is at least 2.

The same lemma holds for any quadrilateralizable polygon, even those that are not orthogonal.

Since q is fixed for any polygon, we will sometimes say "the number of quadrilaterals in P" rather than "the number of quadrilaterals in any quadrilateralization of P."

Applying this lemma to the sufficiency bound of $\lfloor(3n+4)/16\rfloor$ shows that it is equivalent to $\lfloor(3q+5)/8\rfloor$. It is in this form that the bound will appear throughout the proof.

The following lemma is the equivalent of Lemma 3.6.

LEMMA 3.8. Let P be an orthogonal polygon and Q a quadrilateralization of P. There exists a diagonal d in Q that partitions P into two pieces, one of which contains $q = 2$, 3, or 4 quadrilaterals of Q.

Proof. Choose d to be a diagonal of Q that separates off a minimum number of quadrilaterals that is at least 2. Let $q \geq 2$ be this minimum. Let $ABCD$ be the quadrilateral supported by $d = AB$ towards the piece with q quadrilaterals; see Fig. 3.12. The number of quadrilaterals in the BC, CD, and DA regions illustrated in the figure is each less than 2—that is, less than or equal to 1—otherwise q would not be minimal. Therefore, $q \leq 4$. □

It will often be useful to use the dual of a quadrilateralization. Let every quadrilateral of a quadrilateralization Q be a node of a graph $\bar{Q}$, where two nodes are adjacent in $\bar{Q}$ iff their corresponding quadrilaterals share a diagonal.[4] The following is immediate (compare Lemma 1.3).

LEMMA 3.9. For any quadrilateralization Q of an orthogonal polygon, the dual $\bar{Q}$ is a tree with each node of degree no more than 4.

As an application of this observation, we can obtain an alternate proof of Lemma 3.8. Choose any root r for $\bar{Q}$, and let x be a leaf at maximum

4. As mentioned in Chapter 1, this is the graph theoretic "weak dual," weak because no node is assigned to the exterior face.

distance from r, and let y be the parent of x. Then all of the nodes adjacent to y not on the ry path must be leaf nodes; otherwise there would be a path longer than rx. Thus the diagonal of y that crosses the ry path cuts of 2, 3, or 4 quadrilaterals, depending on whether y is of degree 2, 3, or 4, respectively.

One of the main tools used throughout the proof is a cut, a tool previously used in Section 2.5. A *cut* L in an orthogonal polygon P is a maximal interior line segment in P (maximal in the sense that any line segment properly containing L contains a point exterior to P) that contains an edge and a reflex vertex of P. L partitions P into two or three pieces, depending on whether it contains one or two reflex vertices respectively; see Fig. 3.13. In either case, the following holds.

LEMMA 3.10. The sum of the number of quadrilaterals in the pieces defined by a cut L of P is equal to the number of quadrilaterals in P.

Proof. Suppose L partitions P into two pieces P_1 and P_2 as in Fig. 3.13a. Let P, P_1, and P_2 have n, n_1, and n_2 vertices and q, q_1, and q_2 quadrilaterals, respectively. Then $n_1 + n_2 = n + 2$, as L introduces one new vertex, counted in each of P_1 and P_2. Lemma 3.7 shows that

$$q_1 + q_2 = \frac{n_1 - 2}{2} + \frac{n_2 - 2}{2} = \frac{n - 2}{2} = q.$$

If L partitions P into three pieces as in Fig. 3.13b, then L can be considered as a combination of two "half" cuts, each resolving just one reflex vertex. The first partitions P into two pieces, and the second partitions one of the pieces into two, resulting in three pieces. Applying the result just established for two pieces yields the lemma for three pieces. ☐

We now present a series of lemmas detailing the relationship between a diagonal of a quadrilateralization and the local structure of the polygon; henceforth "diagonal" means diagonal of a quadrilateralization. Recall that the *orientation* of an edge is horizontal or vertical. We will say that edges a and b are *to the same side* of d if they are in the same piece of P partitioned off by d; note that a and b may be in opposite half-planes defined by d but still to the same side.

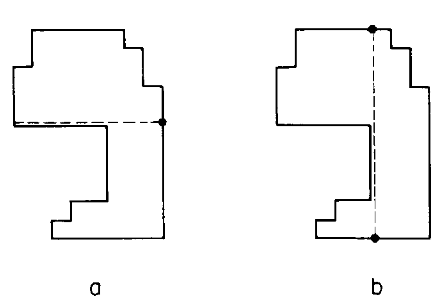

Fig. 3.13. A cut partitions a polygon into two (a) or three (b) pieces.

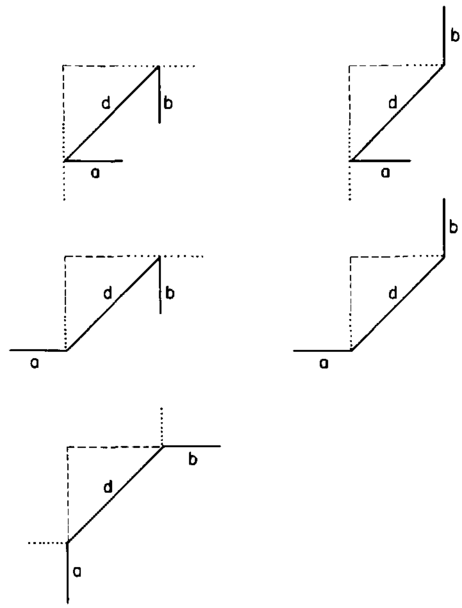

Fig. 3.14. The five possible arrangements when a and b have opposite orientations. The dotted lines represent possible orientations of the edges to the other side of d; the dashed lines indicate an added right angle that forms a subpolygon.

LEMMA 3.11. Let a and b be edges of P adjacent and to the same side of a diagonal d. Then a and b have the same orientation.

Proof. Without loss of generality orient d with positive slope with the polygon P_1 containing a and b below. Assume for contradiction that a and b have different orientations. Then there are five distinct possible combinations of a and b: hanging up, down, left, or right from the endpoints of d, as shown in Fig. 3.14. The other three possible combinations force a and b to not be to the same side of d. We can derive a contradiction in two ways. First note that in all five cases, addition of a right angle above d produces a new orthogonal polygon P'; perhaps it will be necessary to put this new angle on a different "level" as defined in Section 2.2, but this will not affect the angle sums. Let q_1 be the number of quadrilaterals in P_1. Then the sum of the internal angles of P' is $360q_1 + 180$. But this implies that P' is not quadrilaterizable, in contradiction to Theorem 2.1.

For a second proof, recall Lubiw's scheme of assigning "types" to each vertex of an orthogonal polygon such that they alternate type 1 and type 2 in a traversal of the boundary (Section 2.4.2, especially Fig. 2.36). In all five cases of Fig. 3.14, d connects two vertices of the same type. Thus the strict alternation is destroyed, and P_1 has an odd number of vertices. But this contradicts the assumption that P_1 is quadrilateralizable, since any polygon partitioned into quadrilaterals must have an even number of vertices. ☐

Because no internal angle of a quadrilateral can be greater than 270°

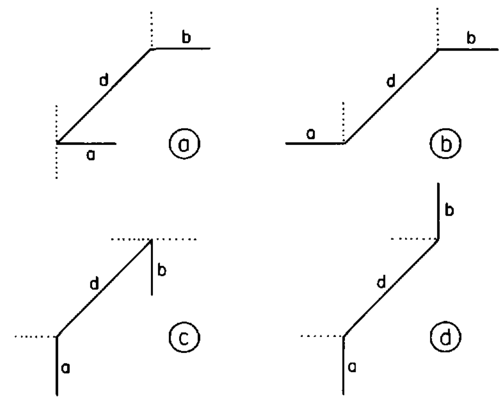

Fig. 3.15. The four possible arrangements of a and b when d is in its standard orientation. The dotted lines indicate the possibilities for the edges adjacent to d and to a and b.

(since all are subangles of either 90° or 270°), only four configurations are possible for d, a, and b, as illustrated in Fig. 3.15. We will raise this observation to a lemma for later reference.

LEMMA 3.12. The only configurations possible for a diagonal d and its two adjacent edges a and b to one side (perhaps after rotation and reflection to orient d with positive slope) are those shown in Fig. 3.15.

Although not needed for the proof of the art gallery theorem, we now turn our attention to characterizing the quadrilateral trees of orthogonal polygons. This is accomplished by showing that Lemma 3.12 restricts the configuration possible for a quadrilateral of a specific degree to a finite set of possibilities, and that only certain configurations can "mate" with one another as adjacent quadrilaterals. We start by showing that a degree 4 quadrilateral can have only one configuration. A *configuration* is defined by the orientations of the edges of the polygon adjacent to each vertex of a quadrilateral, and the type (convex/reflex) of the vertices.

Let a reflex vertex whose exterior angle is in the first quadrant (between the positive x and positive y axes) be called type 1, in the second quadrant (between the positive y and the negative x axes), type 2, and similarly for type 3 and 4.

LEMMA 3.13. Let $ABCD$ be a quadrilateral of degree 4 in $\bar{Q}$ for any orthogonal polygon P. Then A, B, C, and D are each reflex vertices of P, of types $1, 2, 3$, and 4 in counterclockwise order.

Proof. Assume to the contrary that at least A is convex. Without loss of generality let A be a lower left corner as illustrated in Fig. 3.16a. Then Fig. 3.15a shows that edge a' forces b and b' to have the orientations shown at B; b' forces c and c' as shown at C; and c' forces d and d' at D. But now d' lies inside $ABCD$, contradicting the assumption that $ABCD$ is an internal quadrilateral.

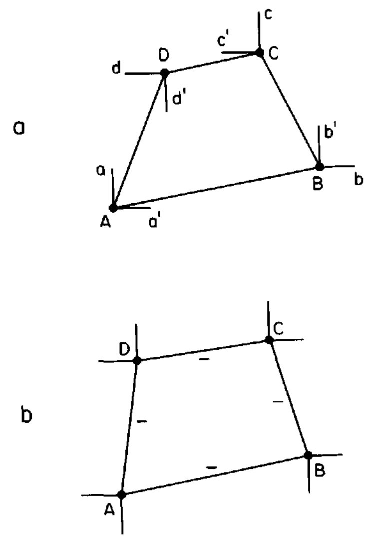

Fig. 3.16. If A is convex (a), d' is forced to be internal to $ABCD$; if A is reflex (b), the degree 4 quadrilateral has a unique configuration.

Now let A be a type 3 reflex vertex. Following the same logic as above forces the configuration shown in Fig. 3.16b, establishing the lemma. ☐

The possible configurations proliferate for quadrilaterals of smaller degree, but the proofs proceed the same way, repeatedly applying the constraints imposed by Lemma 3.12, and will only be sketched.

LEMMA 3.14. A quadrilateral of degree three can have just one of the four configurations shown in Fig. 3.17.

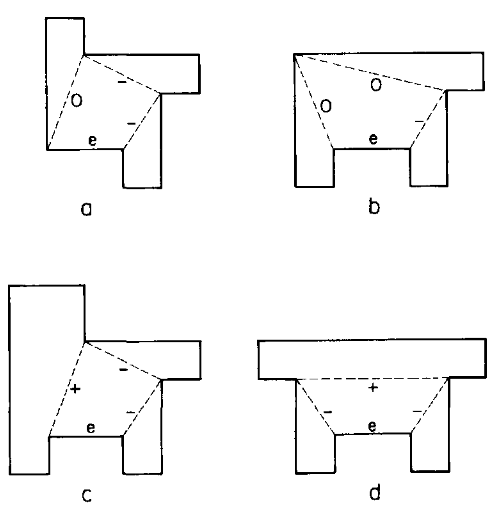

Fig. 3.17. The four configurations possible for a degree 3 quadrilateral.

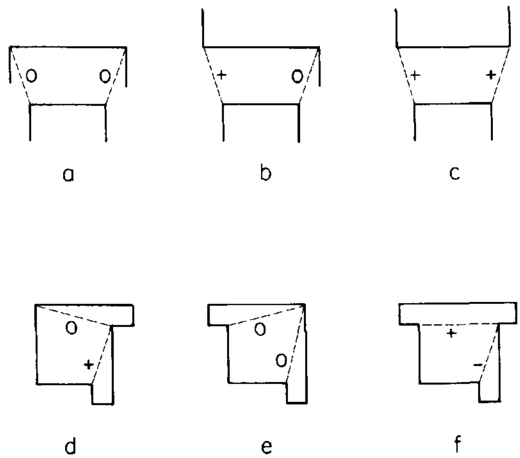

Fig. 3.18. The six configurations possible for a degree 2 quadrilateral.

Proof. Let e be the edge of the quadrilateral shared with the polygon. It is easily shown using Lemma 3.12 that both endpoints of e cannot be convex. If one endpoint is convex and the other reflex, Fig. 3.17a is forced. If both are reflex, three configurations are possible, shown in Figs. 3.17b–3.17d. □

The +, −, and 0 markings in Figs. 3.16 and 3.17 (and in the figures to follow) will be explained later.

LEMMA 3.15. A quadrilateral of degree 2 can have just one of the six configurations shown in Fig. 3.18.

Proof. If the two edges shared with the polygon are non-adjacent, then the three configurations shown in Figs. 3.18a–3.18c are possible. If the shared edges are adjacent, then the three configurations shown in Figs. 3.18d–3.18f are possible. □

LEMMA 3.16. A quadrilateral of degree 1 can have just one of the two configurations shown in Fig. 3.19.

This completes the classification of the possible configurations of the quadrilaterals in an orthogonal polygon. In order to study which configurations can mate with one another, we introduce the concept of "charge" on a diagonal. Let a and b be edges to the same side and adjacent to a diagonal d

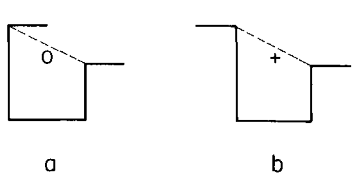

Fig. 3.19. The two configurations possible for a degree 1 quadrilateral.

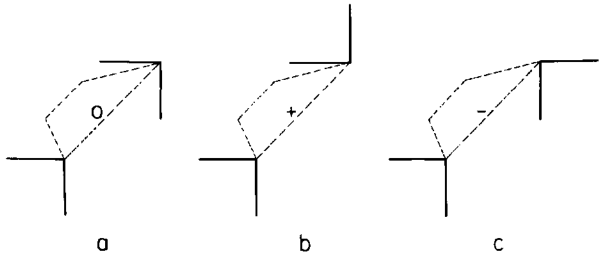

Fig. 3.20. Definitions of the three diagonal charges.

of a quadrilateral. If a and b lie in the same half-plane determined by d, then we will say they have the same *parity*; otherwise they have opposite parity. Thus in Figs. 3.15a and 3.15c, a and b have the same parity, and in Figs. 3.15b and 3.15d they have opposite parity. The *charge* on a diagonal d of a quadrilateral q, with respect to q, is 0 if the adjacent edges to both sides of d have the same parity (Fig. 3.20a), + if the adjacent edges to the q side have the same parity, and the adjacent edges to the opposite side of d have opposite parity (Fig. 3.20b), and − if the adjacent edges to the q side have opposite parity, and those to the opposite side of d have the same parity (Fig. 3.20c). Note that charge is defined with respect to a quadrilateral, so that each diagonal has a charge defined on either side.

LEMMA 3.17. The net charge on any diagonal in a quadrilateralization of an orthogonal polygon must be zero: the charges must be 0/0, +/−, or −/+.

Proof. This is immediate from the definition of charge: a 0 charge on one side is a 0 from the other side, and a + charge on one side is a − from the viewpoint of the other side. □

For the purpose of determining which configurations of quadrilaterals can mate with one another, each configuration can be reduced to a square symbol labeled with charges. The symbols corresponding to the configurations established in Lemmas 3.13–3.16 are displayed in Fig. 3.21 in the same order in which they appear in Figs. 3.16–3.19. We will refer to these symbols as, for example, [3b], meaning the b symbol for a degree 3 quadrilateral as displayed in Fig. 3.21. All quadrilateral trees of orthogonal polygons can be constructed by gluing these symbols together such that each diagonal is uncharged.

We may finally state and prove the characterization theorem.

THEOREM 3.2 [O'Rourke 1985]. A tree is a quadrilateral tree for a simple orthogonal polygon iff no node has degree greater than 4, and the tree contains no path connecting two degree 4 nodes by a sequence of zero or more degree 3 nodes—that is, the path degree sequence (4 3* 4) does not occur.

Proof. It is immediate that two degree 4 nodes cannot be adjacent, since

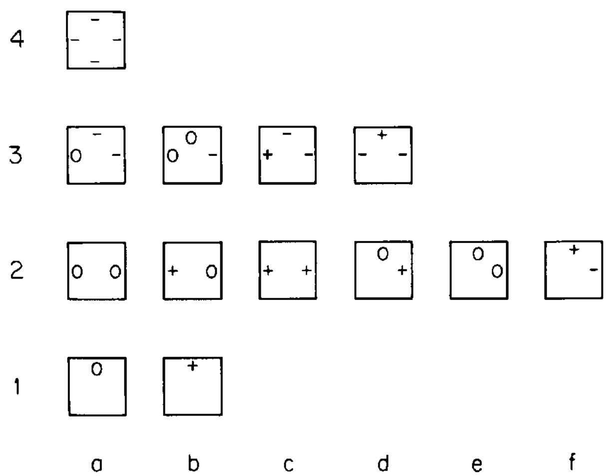

Fig. 3.21. Symbols for all possible quadrilateral configurations. The numbers to the left indicate the degree of the quadrilaterals, the letters below distinguish different configurations.

the symbol [4] has a negative charge on every diagonal. The degree sequence (4 3 3 3. . .) can be achieved by mating [4] with the + charge of either [3c] or [3d], and then mating $-/+$ again with either [3c] or [3d], and so on. But it is clear that the last degree 3 quadrilateral in such a sequence has two $-$ diagonals free, neither of which can mate with [4]. Thus the degree sequence (4 3 . . . 3 4) cannot occur in the quadrilateral tree of any orthogonal polygon.

Now we show that any tree that does not contain a (4 3* 4) path can be realized as the quadrilateral tree of an orthogonal polygon, by assigning square symbols to each node such that all diagonals are uncharged. Assign to each degree 4 node the only choice, [4]. For each connected subtree S composed of degree 3 nodes, distinguish two cases. If S is adjacent to a [4], assign [3c] to each node in S, aligning the charges to balance. This will leave only $-$ charges on the unmatched diagonals of S. If S is not adjacent to a [4], assign one of the leaves [3a], and all the other nodes of S [3c] as in the first case. Now there is one unmatched 0 diagonal, and the remaining unmatched diagonals are negatively charged. The important point is that no unmatched diagonal has a + charge. Next assign each degree 1 node adjacent to a negative diagonal [1b], and all others [1a]. Note that a degree 1 node will not be adjacent to a + charge by construction. Finally assign the degree 2 nodes one of the symbols to cancel the charges appropriately. Since the only charge configuration not available with degree 2 nodes is one with two negative diagonals, this will always be possible as long as a degree 2 node does not have to mate with two positive diagonals. But by construction, all free diagonals are either 0 or $-$. This completes the construction and the proof. □

The construction procedure is illustrated in Fig. 3.22. Figure 3.22a shows a tree that does not contain the forbidden degree sequence, and Fig. 3.22b

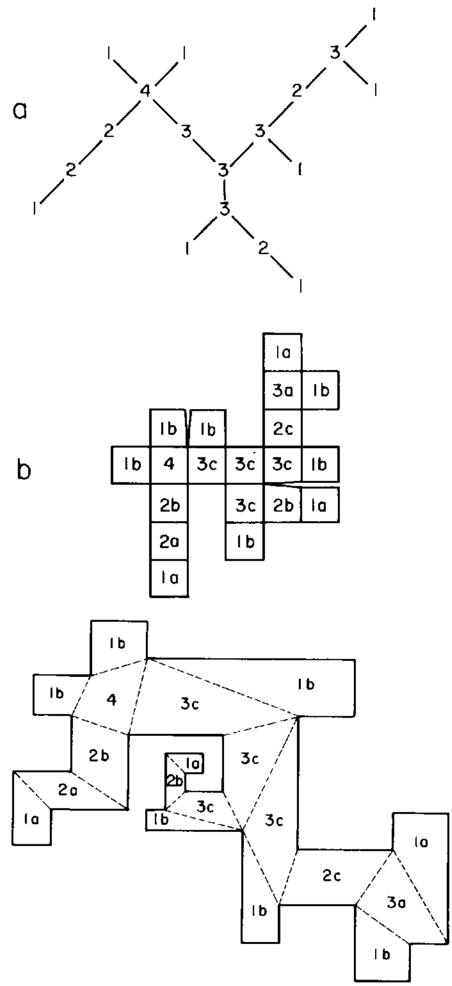

Fig. 3.22. A non-forbidden tree (a), a selection of symbols matching the tree degrees (b), and an orthogonal polygon realizing the symbols (c).

shows the symbols assigned by the construction, glued together appropriately to cancel charges. Finally Fig. 3.22c shows an orthogonal polygon that results by replacing the symbols by their corresponding configurations. It is clear that there are many options in the transition from the symbols to the actual polygon, but the transition is always possible by adjusting the lengths of the edges to avoid overlap, in a manner similar to the local scale changes used in Culberson and Rawlins (1985).

3.3.2. Sharing Lemmas

In this section we develop three "sharing lemmas" similar in spirit to Lemma 3.2 in the proof for general polygons in Section 3.2. They all have

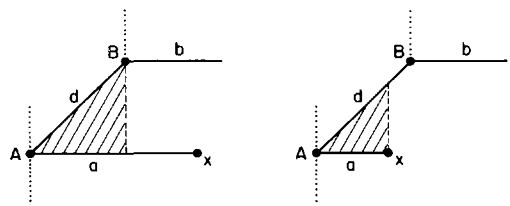

Fig. 3.23. The partial shadow of a diagonal.

the following flavor: "Suppose the induction hypothesis holds, and we are given a polygon with one (or more) guards placed in particular locations 'free.' Then an additional X guards suffice for total coverage." Here X will always be just the right amount to establish the induction hypothesis. I am calling these "sharing" lemmas because in effect they are sharing "fractional" guards across the induction dividing diagonal.

The induction hypothesis that is the premise of these lemmas is:

Induction Hypothesis (IH). Any orthogonal polygon with $q' < q$ quadrilaterals may be covered with $\lfloor (3q' + 5)/8 \rfloor$ mobile orthogonal guards.

First we present a specialized geometric lemma that will be needed in the proofs of the sharing lemmas. Let a and b be the two edges adjacent to and to the same side of a diagonal d, with the same parity. Thus we have either Fig. 3.15a or 3.15c. These situations are clearly identical after rotation and reflection, and we will henceforth consider just Fig. 3.15a. In this situation, define the *partial shadow* of d to be the closed triangular region defined by d, a, and a vertical line through either x, the right endpoint of a, or through the vertex incident to d and b, whichever is leftmost. See Fig. 3.23.

LEMMA 3.18. The partial shadow of a diagonal in a quadrilateralization of an orthogonal polygon is empty.

Proof. The partial shadow is only defined in the situation illustrated in Fig. 3.23. Let A be the vertex incident to d and a as shown. Assume the shadow is not empty, and let e be the leftmost vertical edge in the shadow. Then A and e must be part of a quadrilateral Q. But there is no vertex that can serve as the fourth for Q: it cannot lie to the right of e, for then Q would be non-convex; it cannot lie collinear with e, for then our general position assumption is violated; nor can it lie to the left of e, since e is leftmost. □

The following lemma is almost the direct analog of Lemma 3.2.

LEMMA 3.19. If P is a polygon of q quadrilaterals with one guard placed along a convex edge e (one whose endpoints A and B are both convex vertices), then assuming IH, P can be covered with an additional $\lfloor [3(q-1)+5]/8 \rfloor$ guards.

Proof. The proof is by induction on q. The lemma is clearly true when $q = 1$. Assume it is true for $q' < q$. Let $Q = ABCD$ be the quadrilateral

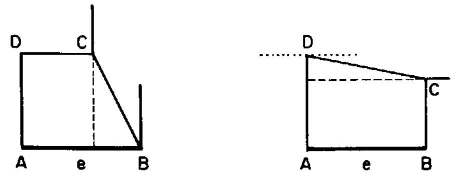

Fig. 3.24. If $ABCD$ has degree 1 and e is guarded, one quadrilateral may be removed.

containing $e = AB$. The proof proceeds by cases depending on the degree of Q. If $deg(Q) = 1$, it follows easily; $deg(Q) = 2$ requires more work; and $deg(Q) = 3$ is not possible.

Case ($deg(Q) = 1$). Either BC (or symmetrically DA) or CD is the sole internal diagonal of Q. In either situation, illustrated in Fig. 3.24, a cut through C partitions P into a covered rectangle and a polygon of $q - 1$ quadrilaterals by Lemma 3.10. Applying IH establishes the lemma.

Case ($deg(Q) = 2$). The internal diagonals of Q are either adjacent or not.

Case 2.1 (Non-adjacent Diagonals) The only situation possible is shown in Fig. 3.25a, corresponding to Fig. 3.18a. The two cuts illustrated partition P into 3 pieces, a rectangle bound by the cuts, and two orthogonal polygons P_1 and P_2 of, say, q_1 and q_2 quadrilaterals. By Lemma 3.10, $q_1 + q_2 + 1 = q$. Now note that the guard along e is a guard between two convex vertices in

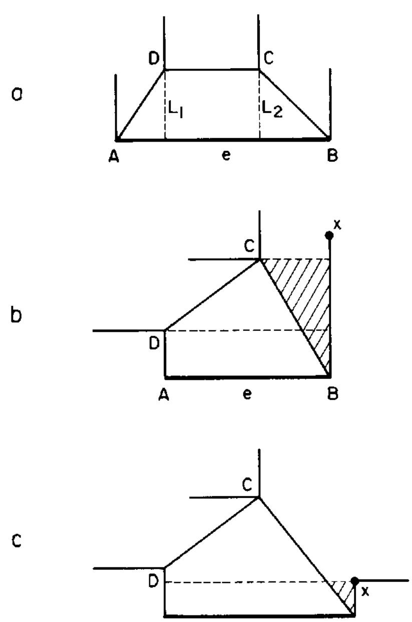

Fig. 3.25. If $ABCD$ has degree 2 and e is guarded, one quadrilateral may be removed, either by induction (a), or by removal of a rectangle (b and c).

each of P_1 and P_2. Since $q_1 < q$ and $q_2 < q$, the induction hypothesis for this lemma applies. Therefore, P can be covered with

$$\lfloor [3(q_1 - 1) + 5]/8 \rfloor + \lfloor [3(q_2 - 1) + 5]/8 \rfloor$$

additional guards. Tedious analysis shows that this is less than or equal to $\lfloor [3(q-1)+5]/8 \rfloor$, establishing the lemma.

Case 2.2 (Adjacent Diagonals). Only Fig. 3.18d is possible, which we will further partition into the two cases shown in Figs. 3.25b and 3.25c. Let BC and CD be the diagonals of Q; C must be above D. The two figures are distinguished by whether x, the upper endpoint of the vertical edge incident to B, is higher or lower than D. In the former case (Fig. 3.25b), a cut through D, and in the latter case (Fig. 3.25c), a cut through x, is guaranteed by the emptiness of the partial shadow of BC (Lemma 3.18) to partition P into a covered rectangle and a polygon of $q-1$ quadrilaterals. Applying IH establishes the lemma.

That the case $deg(Q) = 3$ is not possible is immediate from the possible configurations shown in Fig. 3.17: e is not a convex edge in any of the possible configurations. □

The next sharing lemma in effect "squashes out" two quadrilaterals.

LEMMA 3.20. If P is a polygon of q quadrilaterals with two guards placed on consecutive convex edges AB and BC, then assuming IH, P can be covered with an additional $\lfloor [3(q-2)+5]/8 \rfloor$ guards.

Proof. The proof is similar to the preceding one. The structural possibilities are clearly the same as in that proof, but with BC here playing the role of AB there. Let Z, A, B, C, D, and E be consecutive vertices on the boundary of P. Let Q be the quadrilateral including BC; Q is not necessarily $ABCD$.

Case 1 ($deg(Q) = 1$). Either $Q = ABCD$ with AD the internal edge (Fig. 3.26a), or $Q = BCDE$ with BE the internal edge (Fig. 3.26b). In the first instance D is reflex, and a horizontal cut through it leaves a covered rectangle and a polygon of $q-1$ quadrilaterals that satisfies Lemma 3.19. Applying that lemma establishes the result. In the second instance, E is reflex, and a vertical cut leaves a covered rectangle and a polygon of $q-1$

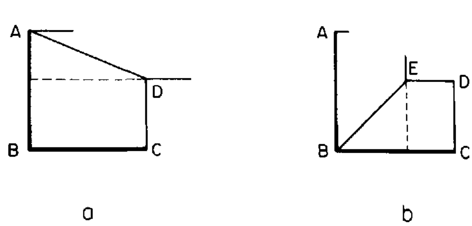

Fig. 3.26. If $ABCD$ has degree 1 and AB and BC are guarded, one quadrilateral may be removed.

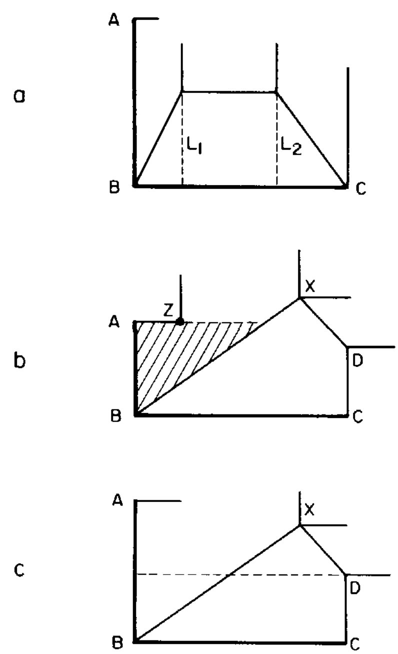

Fig. 3.27. If $ABCD$ has degree 2 and AB and BC are guarded, either induction applies (a), the situation is impossible (b), or a quadrilateral may be removed (c).

quadrilaterals that satisfies the induction hypothesis (and Lemma 3.19). In all cases, then, the result holds.

Case 2 ($deg(Q) = 2$). The same two cases apply as in Lemma 3.19.

Case 2.1 (Non-adjacent Diagonals.) The situations must be as in Fig. 3.27a. The two cuts L_1 and L_2 partition P into a covered rectangle, a polygon P_1 with q_1 quadrilaterals that satisfies the induction hypothesis, and a polygon P_2 of q_2 quadrilaterals that satisfies Lemma 3.19, where $q_1 + q_2 + 1 = q$. Applying both results yields coverage with

$$\lfloor [3(q_1 - 2) + 5]/8 \rfloor + \lfloor [3(q_2 - 1) + 5]/8 \rfloor$$

additional guards. A tedious analysis reveals this to be no larger than $\lfloor [3(q-2)+5]/8 \rfloor$ for $q \geq 2$.

Case 2.2 (Adjacent Diagonals) A cannot be a vertex of Q: Figs. 3.18d, 3.18e, and 3.18f do not permit three consecutive convex vertices. Therefore $Q = BCDX$, corresponding to Fig. 3.18d, with D reflex. Now if X is above A, as in Fig. 3.27b, Z is in the partial shadow of BX, a contradiction. So we are left with the situation shown in Fig. 3.27c. A cut through D establishes the result as in Case 1 (compare Fig. 3.26a). □

The last sharing lemma is the most complex. It takes the form: if certain sharing conditions hold, then the remainder of the polygon needs one full guard less than $\lfloor (3q+5)/8 \rfloor$. The sharing conditions are rather complicated, but essentially the idea is to place two guards crossing each other orthogonally such that the previous two lemmas apply to the pieces of the resulting partition.

LEMMA 3.21. If P is a polygon of q quadrilaterals with a guard placed along a maximal segment L_1 that contains a polygon edge that is situated in P such that

(a) there are two cuts orthogonal to L_1 that partition off rectangles touching (and thereby covered by) L_1, and
(b) L_1 cuts the remainder (P with the two rectangles removed) into one or two (i.e., not three) pieces.

then assuming IH, P can be covered with an additional $\lfloor (3q+5)/8 \rfloor - 1$ guards.

Proof. Note that L_1 is not necessarily a cut, but could be a convex edge. Let P_1 and P_2 be the two pieces separated by L_1 after removal of the two rectangles; P_2 may be empty. The premise of the lemma is a bit ungainly, but is composed to have two geometric consequences:

(1) If P_2 is not empty, there is at least one reflex vertex on L_1 in $P_1 \cup P_2$.
(2) L_1 lies on a convex edge of both P_1 and P_2.

We first support these claims. If L_1 is a cut, it partitions P into two or three pieces. If L_1 cuts P into three pieces, the premise can only be satisfied if the third piece is composed of one or both of the rectangles cut off, disallowing Fig. 3.28a for example. If L_1 cuts P into two pieces, then if the second piece is composed of one or both of the cut off rectangles, then P_2 is empty, as in

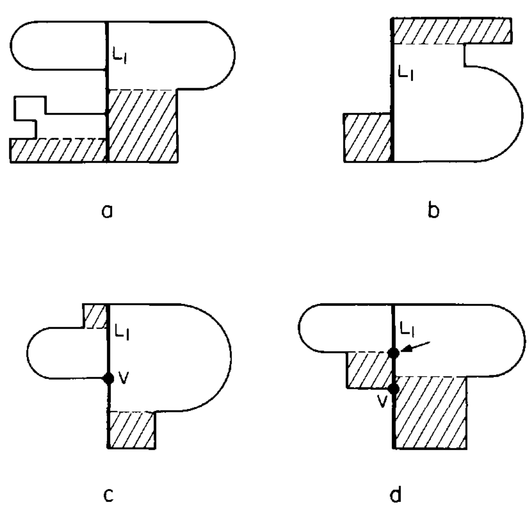

Fig. 3.28. The cut in (a) does not satisfy the conditions of the lemma; in (b), P_2 is empty; in (c) and (d), L_1 contains a reflex vertex.

Fig. 3.28b, for example. Let V be the reflex vertex on L_1 (the one in $P_1 \cup P_2$ in case there are two reflex vertices on L_1). Then either V is in $P_1 \cup P_2$, as in Fig. 3.28c, or there is at least one other reflex vertex on L_1 introduced by the orthogonal cuts, as in Fig. 3.28d. Finally, since L_1 is a supporting line for both P_1 and P_2, it constitutes a convex edge in each.

Let P_1 and P_2 have q_1 and q_2 quadrilaterals. Then $q_1 + q_2 = q - 2$ by Lemma 3.10. We now apply the previous two sharing lemmas to establish the claim.

By property (2), Lemma 3.19 applies to both P_1 and P_2, resulting in complete coverage with

$$\lfloor [3(q_1 - 1) + 5]/8 \rfloor + \lfloor [3(q_2 - 1) + 5]/8 \rfloor = \lfloor (3q_1 + 2)/8 \rfloor + \lfloor (3q_2 + 2)/8 \rfloor$$

additional guards. A tedious case analysis shows that this quantity is no greater than $\lfloor (3q + 5)/8 \rfloor - 1$ for all possible mod 8 residues of q_1 and q_2 *except* in the single case when both $q_1 \equiv 2 \pmod 8$ and $q_2 \equiv 2 \pmod 8$. We now concentrate on this "hard" case. Note that P_2 cannot be empty in this case.

Introduce a cut L_2 orthogonal to L_1 through the reflex vertex V guaranteed by property (1), and place a guard along L_2. L_2 must partition one of P_1 or P_2, say P_2, into two pieces P_2' and P_2'', with q_2' and q_2'' quadrilaterals; it may or may not partition P_1, as illustrated in Fig. 3.29. If L_2 partitions P_1, call the pieces P_1' and P_1'' with q_1' and q_1'' quadrilaterals. Now although $q_1' + q_1'' = q_1$ by Lemma 3.10, $q_2' + q_2'' = q_2 + 1$, since L_1 already resolved the reflex vertex V.

Note that the conditions for the application of Lemma 3.20 hold for both P_2' and P_2''; L_1 and L_2 both lie on convex edges in each. Therefore, all of P_2

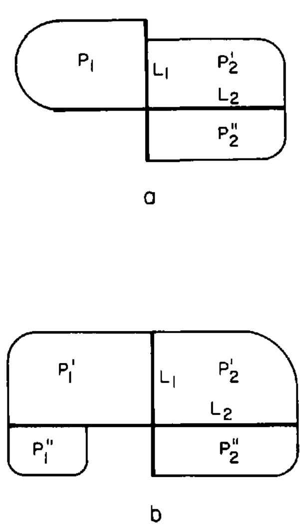

Fig. 3.29. The cut L_2 may (a) or may not (b) partition P_1.

can be covered with

$$g_2 = \lfloor [3(q_2' - 2) + 5]/8 \rfloor + \lfloor [3(q_2'' - 2) + 5]/8 \rfloor \tag{1}$$

guards. If L_2 does not partition P_1, then we apply Lemma 3.20 to P_1 to cover it with

$$g_1 = \lfloor [3(q_1 - 2) + 5]/8 \rfloor \tag{2}$$

guards. If L_2 does partition P_1, then Lemma 3.20 can be applied to P_1' above L_2, and Lemma 3.19 to P_1'' below L_2, resulting in

$$g_1' = \lfloor [3(q_1' - 2) + 5]/8 \rfloor + \lfloor [3(q_1'' - 1) + 5]/8 \rfloor \tag{3}$$

guards. Using the special case assumption that $q_1 = 8k_1 + 2$, (2) yields $g_1 = 3k_1$, and a case analysis and $q_1' + q_1'' = q_1$ shows that (3) implies $g_1' \leq 3k_1$. Therefore, $3k_1$ guards suffice for P_1 in either case. The assumption $q_2 = 8k_2 + 2$ and $q_2' + q_2'' = q_2 + 1$ leads to (1) to $g_2 \leq 3k_2$. Thus a total of $3(k_1 + k_2)$ guards suffice. Finally, $q = q_1 + q_2 = 8(k_1 + k_2) + 6$ implies that $\lfloor (3q + 5)/8 \rfloor - 2 = 3(k_1 + k_2)$, which together with the 1 guard placed on L_2, establishes the lemma. □

3.3.3. Proof of Orthogonal Polygon Theorem

We have finally assembled enough lemmas to prove the main theorem.

THEOREM 3.3 [Aggarwal 1984]. $\lfloor (3q + 5)/8 \rfloor = \lfloor (3n + 4)/16 \rfloor$ mobile guards are sufficient to cover any orthogonal polygon P of q quadrilaterals and n vertices.

Proof. The proof is by induction on q. If $q \leq 2$, then 1 guard clearly suffices. Assume now the induction hypothesis IH. Fix an arbitrary quadrilateralization of P. Lemma 3.8 established that there is a diagonal d that cuts off a minimal number k of 2, 3, or 4 quadrilaterals. These constitute the three cases of the proof, which we consider in reverse order.

Case $k = 4$. Recall from the proof of Lemma 3.8 (see Fig. 3.12) that d must be a diagonal of a degree 4 quadrilateral Q, say $Q = ABCD$ with $d = DA$. Lemma 3.13 shows that A, B, C, and D must all be reflex vertices. Let A be left of and lower than D, which can be achieved without loss of generality by rotation and reflection. We can distinguish three cases, only two of which are real possibilities, depending on the horizontal sorting of B, C, and D. We will use the notation $X < Y$ to mean that point X is strictly left of point Y.

Subcase ($C < D$ (Fig. 3.30a)). This case violates Lemma 3.11, as completion of the polygon between B and C as illustrated demonstrates.

Subcase ($C > D$ and $B < D$ (Fig. 3.30b)). Place a guard on the vertical cut L_1 through D as illustrated. Then this cut satisfies the conditions of Lemma

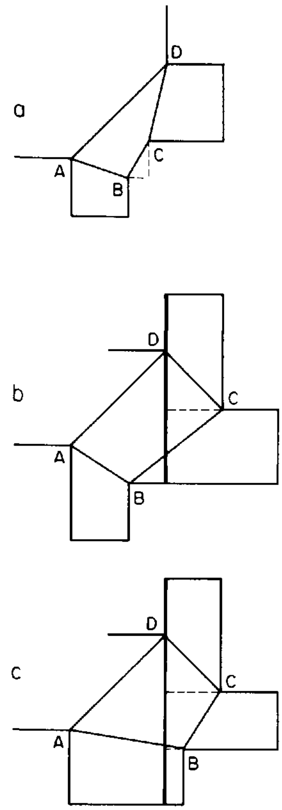

Fig. 3.30. If $ABCD$ has degree 4, then either the situation is impossible (a), or Lemma 3.21 applies (b and c).

3.21, with P_2 empty. Applying that lemma yields coverage of P_1 with $\lfloor(3q+5)/8\rfloor - 1$ guards, which, together with the guard along L_1, establishes the theorem.

Subcase ($C > D$ and $B > D$ (Fig. 3.30c)). Again Lemma 3.21 applies with L_1 the vertical cut through D.

Case $k = 3$. The proof of Lemma 3.8 shows that d is a diagonal of a degree 3 quadrilateral Q. Let $Q = ABCD$ with $d = DA$. Orient d as in Fig. 3.31a, and assume without loss of generality that the edges of P' adjacent to d are horizontal, with A reflex. (Figure 3.17 shows that at most one of A, B, C, D is convex, so one end of d is always reflex. If the other end is convex, d angles away from the reflex vertex as in Fig. 3.31a; if the other end is reflex, then either d angles away as in Fig. 3.31a, or it will after reflection in the x axis.) We distinguish five cases, depending mainly on which edge of Q is a polygon edge. In each case, Lemma 3.21 is invoked.

Subcase (BC is a polygon edge.) BC must be horizontal and below A,

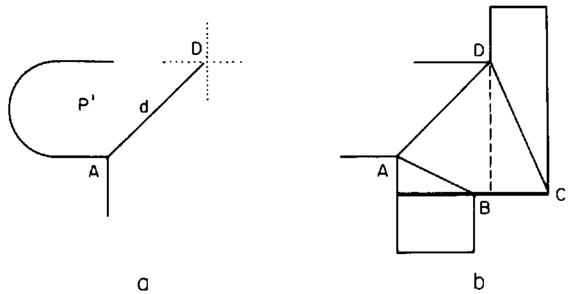

Fig. 3.31. The configuration when $k = 3$: the dotted edges in (a) represent the two possible orientations of the polygon edge at D. If BC is a polygon edge, Lemma 3.21 applies.

otherwise either Q is non-convex or Lemma 3.11 is violated. Consequently the parity of the horizontal edges adjacent to d in P' is the same, and the situation is as illustrated in Fig. 3.31b. A horizontal cut through B satisfies the conditions of Lemma 3.21 (with P_2 empty), and the theorem follows by placing a guard along the cut and applying Lemma 3.21.

Subcase (CD is a polygon edge). CD must be vertical to satisfy Lemma 3.11, and B must be left and below C since Q is convex. Regardless of the vertical placements of A, B, and C, a vertical cut through B satisfies Lemma 3.21. Figure 3.32 shows that in each of the three possible vertical sortings (A, C, B), (C, A, B), and (C, B, A), in a, b, c respectively, the theorem follows by placing a guard along the cut and applying Lemma 3.21.

Subcase (AB is a polygon edge). Distinguish further subcases, depending on the location of C with respect to B and D.

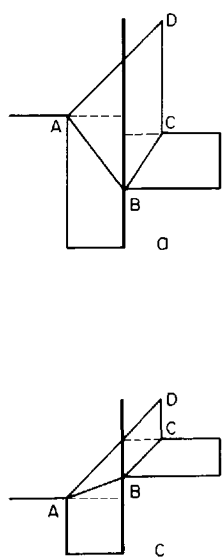

Fig. 3.32. If CD is a polygon edge, Lemma 3.21 applies.

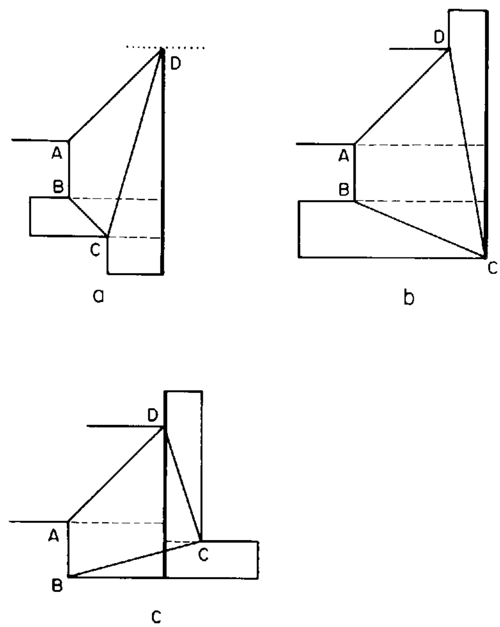

Fig. 3.33. If AB is a polygon edge, Lemma 3.21 applies in all cases.

Subsubcase (C is below B and left of D (Fig. 3.33a).). Let L_1 be the maximal vertical segment containing D. This satisfies Lemma 3.21, regardless of whether or not D is reflex.

Subsubcase (C is below B and right of D (Fig. 3.33b).). Let L_1 be the vertical edge containing C. This satisfies Lemma 3.21.

Subsubase (C is above B and left of D). This case violates Lemma 3.11 and so is not possible.

Subsubcase (C is above B and right of D (Fig. 3.33c).). Let L_1 be the maximal vertical segment containing D. This satisfies Lemma 3.21.

Case $k = 2$. Although this case is simplest in some sense, it requires the most extensive sharing, since so little is cut off by d. Fortunately, all the sharing is concentrated into Lemma 3.21. We partition the problem into two subcases, depending on whether the edges adjacent to d in P' have the same or opposite parity. Let $d = DA$ as usual, and let A be below D so that A is always reflex.

Subcase (Same Parity). Let d be oriented as in Fig. 3.34a. Place a guard along the vertical edge through D if D is convex (Figs. 3.34b and 3.34c), or along the first vertical edge hit by a horizontal cut through A (Fig. 3.34d). In all cases, Lemma 3.21 applies, with P_2 empty.

Subcase (Opposite Parity). Orient d as in Fig. 3.35a. That both A and D

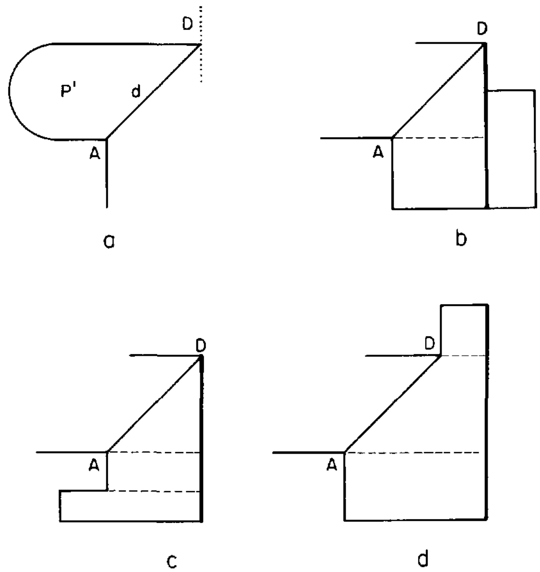

Fig. 3.34. When the edges adjacent to d have the same parity, Lemma 3.21 applies in all cases.

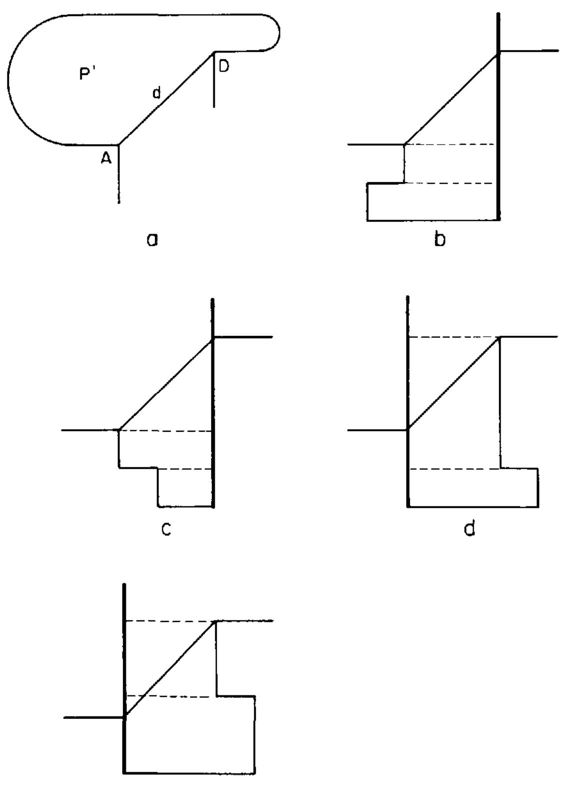

Fig. 3.35. When the edges adjacent to d have opposite parity, Lemma 3.21 applies in all cases.

are reflex with their adjacent edges oriented as shown can be seen by examination of Fig. 3.18. Of the four vertices in the chain counterclockwise between A and D, exactly one is reflex. If the first or second (counterclockwise from A) is reflex (Figs. 3.35b and 3.35c), a vertical cut L_1 through D satisfies the conditions of Lemma 3.21. If the third or fourth vertex from A is reflex (Figs. 3.35d and 3.35e), a vertical cut through A satisfies Lemma 3.21. In all cases, placing a guard along L_1 and applying Lemma 3.21 yields coverage by $\lfloor(3q+5)/8\rfloor - 1 + 1$ guards, establishing the theorem.

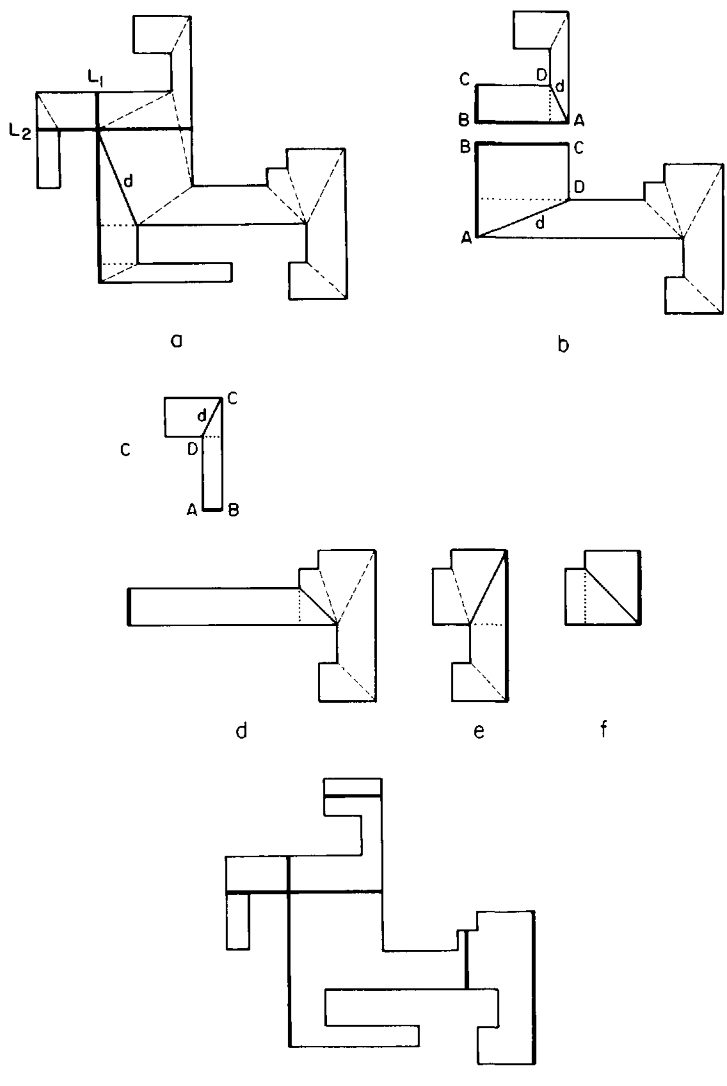

Fig. 3.36. An "execution" of the proof of Theorem 3.3 and the lemmas it invokes. The final guard placement is shown in (g).

We have exhausted all possibilities, and therefore the theorem is established. □

The proof just presented is constructive, and therefore can be converted to an algorithm. The algorithm is highly inefficient, however, since requadrilateralization is implicitly required at almost every step. It will help understanding the proof if we step through a small example, tracking the proof through the various lemmas and "executing" them as procedures.

Consider the polygon shown in Fig. 3.36a. It has $n = 26$ vertices and $q = 12$ quadrilaterals. The theorem then says that six guards suffice; actually four suffice in this case. Using the quadrilateralization in Fig. 3.36a, d is a diagonal that cuts off a minimum number k of quadrilaterals; in this case, $k = 2$. Following the theorem, the $k = 2$ case (opposite parity: Fig. 3.35b reflected) invokes Lemma 3.21. In our particular case, the cut L_1 and the abutting rectangles are shown. L_1 partitions P into pieces with $q_1 = 2$ and $q_2 = 10$ quadrilaterals. This is the hard case of the lemma, and requires a second cut L_2 shown. Two non-trivial pieces remain, and for both Lemma 3.21 invokes Lemma 3.20, because there are guards on two consecutive convex edges (Fig. 3.36b). Both pieces fall under the same case of Lemma 3.20 ($deg(Q) = 1$: Fig. 3.26a reflected), both introducing a cut and invoking Lemma 3.19 for a guard along a single convex edge. Figure 3.36c shows the smaller piece. Lemma 3.19 makes a cut (following Fig. 3.24b) and applies the IH, which in this case is trivial since the remaining piece is a rectangle, which is assigned its own guard. Figure 3.36d shows the larger piece. Again Lemma 3.19 cuts and applies IH to the polygon shown in Fig. 3.36e. We are now back at the "top level" in the main theorem. In the quadrilateralization shown, d cuts off a minimum $k = 2$ quadrilaterals. The case here is the same parity one (Fig. 3.34c), and introduces a guard along the vertical edge shown. The top remainder is handled by Lemma 3.19, because the guard forms a convex edge. Lemma 3.19 then (unnecessarily in this case) invokes IH again, this time at the basis, and a guard is assigned to the rectangle in Fig. 3.36f. The resulting five guards assigned are shown in Fig. 3.36g.

3.4. Discussion

The guards used for the general polygon theorem (Theorem 3.1) are combinatorial: visibility is needed only at the two vertices at the endpoints of diagonals. The guards used for the orthogonal polygon theorem (Theorem 3.3) are geometric: visibility is required throughout their length. The guards used in the two theorems differ in other respects. Several features of the orthogonal guards are:

(1) Visibility is required throughout the length of the guard.
(2) Guards are oriented horizontal or vertical only.
(3) Each guard can be chosen to include an edge of the polygon.
(4) Visibility is only required orthogonal to the guard.
(5) The patrols of two guards may pass through one another.

These were not conditions imposed on the problem, but rather those that "fell out" of Aggarwal's proof. It would be interesting to disallow the fifth condition above: do not permit the lines of two guards to cross. But the most interesting question concerning these qualifications on guard "power" is whether (1) is necessary: can the same result be achieved with combinatorial guards, as in the general polygon case?

Aggarwal has proven several other results on mobile guards (Aggarwal 1984). The most important is that for quadrilaterizable polygons—that is, those that can be partitioned into convex quadrilaterals, $\lfloor n/5 \rfloor$ guards are necessary and sufficient. Since $\lfloor n/5 \rfloor > \lfloor (3n+4)/16 \rfloor$ for all $n > 20$, this result does not contradict Theorem 3.3. Despite Theorem 3.2, which characterizes the quadrilateral trees of orthogonal polygons, it remains an open problem to characterize those polygons that are quadrilateralizable. Aggarwal's proof of the $\lfloor n/5 \rfloor$ result differs in two ways from the proof of Theorem 3.3: first, it is entirely combinatorial, and second, it is much longer: at one point 93 separate cases are considered! Several other of his mobile guard results for specialized polygons will be discussed in the next chapter.