ILLUMINATING THE FREE SPACE BETWEEN QUADRILATERALS WITH POINT LIGHT SOURCES*

Gregoria Blanco[1]
Hazel Everett[2]
Jesus Garcia Lopez[1]
Godfried Toussaint[3]

ABSTRACT

Let *S* be a set of *n* pairwise-disjoint, possibly non-convex, quadrilaterals in the plane. We consider the problem of illuminating the "free" space, i.e., the region of the plane excluding the interiors of the quadrilaterals (which act as barriers), by using point light sources. We show that if the placement of the lights is restricted to the vertices of the quadrilaterals (*vertex* lights) then $2n$ lights are always sufficient and sometimes necessary to do the job. Furthermore, all the locations for the $2n$ vertex lights can be found in $O(n^2)$ time. If the quadrilaterals are isothetic rectangles then $2n$ vertex lights can be placed in $O(n)$ time. If the isothetic quadrilaterals have equal width then $n+2$ point lights are always sufficient and their location can be found in $O(n \log n)$ time. At the cost of only one additional light this result reduces the complexity of the algorithm in [HK] for placing $n+1$ point lights in $O(n^{1.5} \log^2 (n) \log \log n)$ time. Furthermore, for isothetic squares it is shown that n point lights are sufficient and can be found in $O(n \log n)$ time. Finally, for a Manhattan configuration of n isothetic rectangles it is shown that 4 point lights are sufficient and a minimally close placement for them can be determined in O(n) time.

1. Introduction

In 1973 Victor Klee asked the question: given an arbitrarily shaped gallery of n walls (polygon of n edges) what is the minimum number of guards (lights) required to guard (illuminate) its interior [Ho]. Vasek Chvatal soon established what has become known as "Chvatal's Art Gallery Theorem": $n/3$ lights are always sufficient and sometimes necessary [Ch]. In 1981 Avis and Toussaint exhibited an efficient algorithm for actually finding the locations where these lights should be installed [AT]. We should add that it is tempting to place a light on every third vertex of the polygon. The reader can easily design a gallery where this placement strategy will fail no matter on which vertex one starts placing the lights.

---

* Research of the last author was supported by NSERC Grant no. OGP0009293 and FCAR Grant no. 93-ER-0291
1. Department of Applied Mathematics, Escuela Universitaria de Informatica, Madrid, Spain.
2. Department of Mathematics, University of Quebec in Montreal, Canada.
3. School of Computer Science, McGill University, Montreal, Canada.

This type of problem for different types of lights (vertex lights [HK93], point or line-segment lights) and different types of environments (rectilinear polygons, polygons with holes, etc.) falls into an area of research which has come to be known as *art gallery theorems and algorithms*. O'Rourke has written an excellent monograph devoted entirely to this subject [O'R] and Shermer has compiled a comprehensive tutorial survey of results obtained since the publication of O'Rourke's book [Sh]. For the latest results and open problems in this area the reader is referred to [O'R92].

Some work has also been done concerning the illumination of collections of different types of pair-wise non-intersecting objects in the plane such as convex sets [UZ], [Fe], rectangles and triangles [CRU], circles [CRUZ] and line segments [CRU], [JL].

In [CRU] several problems are considered concerning illuminating a set S of n isothetic rectangles and triangles in the plane by point light sources. In particular it is shown that $[(4n+4)/3]$ lights are always sufficient. Furthermore, for the case of equal-width rectangles they show that $n+2$ lights suffice. They ignore the question of the computational complexity of actually finding such placements.

Some work has also been directed at illuminating special sets of objects in the sense that the collection of objects is not arbitrarily placed but has the property that there exists a transversal, i.e., a line that intersects them all. For example, Everett et al. [ELRS] show that if the n squares admit a transversal then $[2n/3]+2$ point lights are sometimes necessary and $[2n/3]+7$ are always sufficient. Furthermore they give an $O(n \log n)$ time algorithm for placing these lights.

In this paper we consider the problem of illuminating, with point light sources, the "free" space, i.e., the region of the plane excluding the interiors of a set S of n given, pairwise non-intersecting, quadrilaterals (which act as barriers). We show that if the placement of the lights is restricted to the vertices of the quadrilaterals (*vertex* lights) then $2n$ lights are always sufficient and sometimes necessary to do the job. Furthermore, all the locations for the $2n$ vertex lights can be found in $O(n^2)$ time. If the quadrilaterals are isothetic rectangles (i.e., with sides parallel to the axes) then $2n$ vertex lights can be placed in $O(n)$ time. If the isothetic quadrilaterals have equal width then $n+2$ point lights are always sufficient and their location can be found in $O(n \log n)$ time. At the cost of only one additional light this result reduces the complexity of the algorithm in [HK] for placing $n+1$ point lights in $O(n^{1.5} \log^2 (n) \log \log n)$ time. Furthermore, for isothetic squares it is shown that n point lights are sufficient and can be found in $O(n \log n)$ time. For a Manhattan configuration of n isothetic rectangles it is shown that 4 point lights are sufficient and a minimally close placement for them can be determined in O(n) time.

2. Manhattan Rectangles

Consider the problem of illuminating an ideal Manhattan-style city with the minimum number of ideal lights. By illuminating the city we mean illuminating all the sidewalks and every portion of every street, in short, all the space not including the interior of the city-blocks (quadrilaterals). Five problems immediately arise: (1) how many lights are necessary and sufficient, (2) where should these lights be placed, (3) what is the computational complexity of determining such a placement of lights, (4) is it any easier to illuminate only the sidewalks (after all people don't usually carry flashlights whereas cars are usually equipped with working headlights), and finally (5)

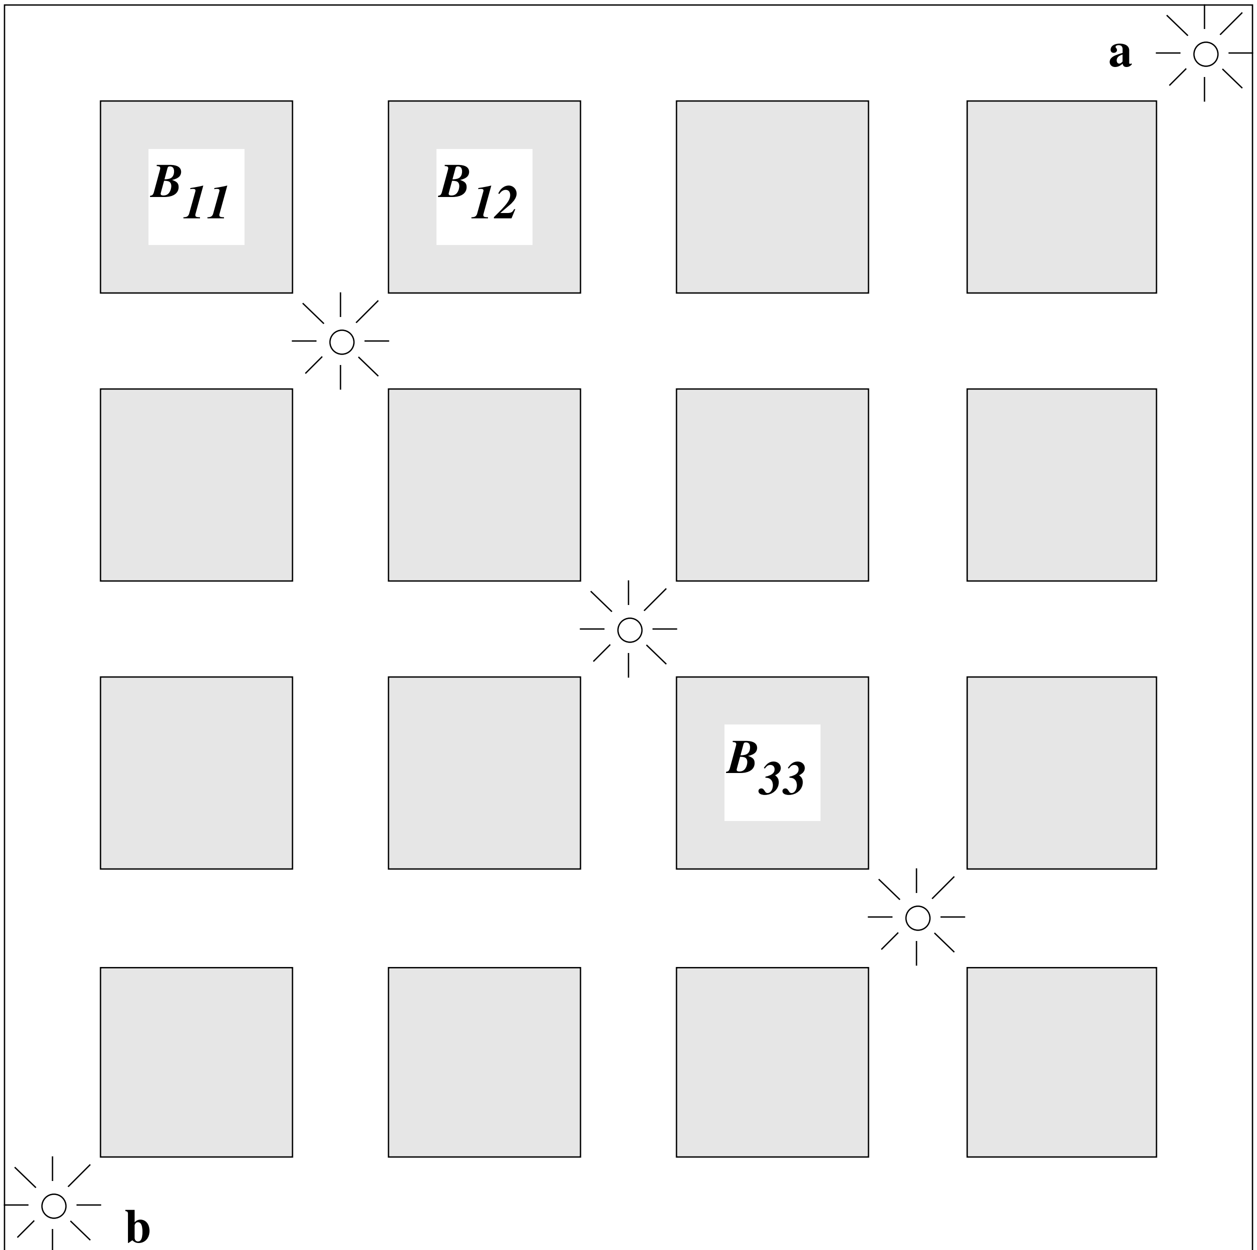

Fig. 1: Illustrating a Manhattan configuration consisting of 4 X 4 = 16 blocks. The five lights so placed will illuminate all the sidewalks and all the streets. However, it is possible to accomplish this task with only four lights.

all of the above problems in the case where we are allowed to place lights only on street corners, i.e., vertices of the squares representing the city blocks (such lights are called *vertex* lights as distinguished from *point* lights). By an ideal Manhattan-style city we mean a city consisting of an array of $k \times k = n$ blocks in which the streets (running horizontally) and the avenues (running vertically) are each of constant width although different streets and avenues may have different widths. By an ideal light we mean that its rays travel an infinite distance in all directions. Such a city is illustrated in Fig. 1 and shall be denoted by M. For the moment we assume that the blocks

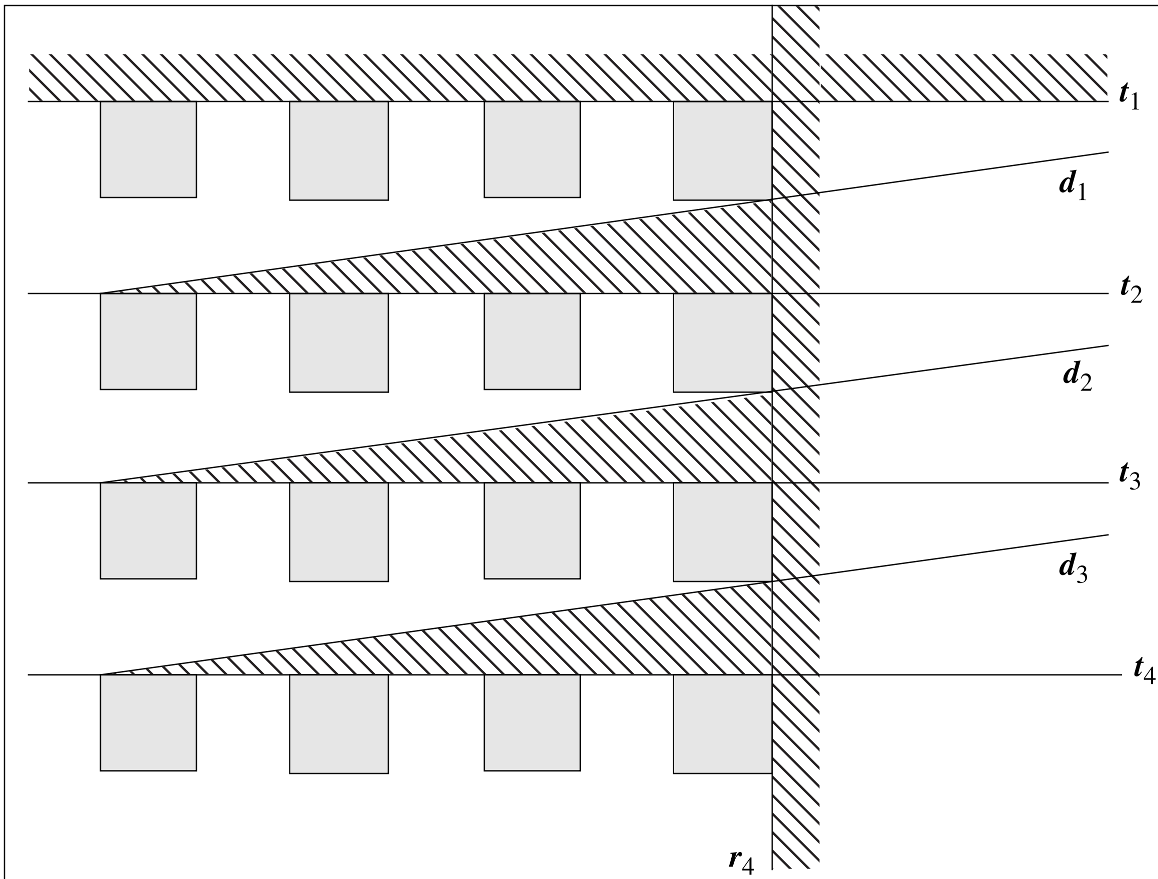

Fig. 2: One light above t_1 and below d_3 illuminates all the space above t_1 and to the right of r_4 as well as the lower right diagonal half of every street.

are squares and that the streets and avenues are all of equal width. Presented with problems (1) and (2) with k=4, the typical *person-in-the-street* comes up, sometimes after considerable thought, with five lights placed as in Fig. 1. Lights a and b light up the "outside" of the city and the remainder each take care of one whole street and avenue to light up the "inner" city. This may appear to be a clever placement of lights even though it is obvious that such a strategy requires placing as many as $\sqrt{n}+1$ lights for cities of arbitrary size. However, it turns out that only four point lights are sufficient. Furthermore, four are sufficient no matter what the value of k is; any size city can be illuminated with no more than four lights as we now demonstrate.

Let $FS(M)$ denote the *free space* to be illuminated, i.e., the region of the plane not including the interior of the blocks. Let B denote the set of blocks B_{ij}, *i,j= 1,2,...,k.* We assume that B is given as a two dimensional array in which each block B_{ij} is specified in terms of the cartesian coordinates of its corners (see Fig. 1). Let $L = \{l_1, l_2, l_3, l_4\}$ denote the set of lights where the l_i i=1,...,4 are points in $FS(M)$ denoting the position of the lights. We say that $FS(M)$ is illuminated by L if for every point x in $FS(M)$ there exists an integer i such that the line segment $[x, l_i]$ does not intersect the interior of any block.

We consider the regions inside and outside the convex hull of $FS(M)$ separately. First consider the streets inside the convex hull of $FS(M)$. Let t_i, i=1,2,...,k denote the horizontal line colinear with the upper sides of the blocks B_{ij} (the south sides of the streets), and let $UH(t_i)$ denote the upper half-planes determined by t_i, i=1,2,...,k. Let d_i, i=1,2,...,k-1 denote the diagonal lines which are critical lines of support (in a counter-clockwise manner) separating the B_{ij} from the $B_{i+1,j}$, and let $LH(d_i)$ denote the lower half-planes determined by d_i, i=1,2,...,k-1. These lines are illustrated in Fig.2. It follows that if a light is placed in the cone $C(d_i,t_{i+1}) = LH(d_i) \cap UH(t_{i+1})$ it illuminates the lower right diagonal half of the street between the B_{ij} and the $B_{i+1,j}$. Therefore, to illuminate the lower right diagonal half of every street it suffices to place a light in the intersection of all such cones. Furthermore, the intersection of all such cones lies to the right of the line colinear with the east sides of the right-most blocks (r_4 in Fig. 2) and therefore such a light will illuminate the region outside the convex hull of $FS(M)$ and to the right of r_4. If we place this same light so that it also lies above $UH(t_1)$ then such a light will illuminate the region outside the convex hull of $FS(M)$ and above t_1. Therefore in general we can illuminate half of $FS(M)$ outside its convex hull and half of the streets inside the convex hull of $FS(M)$ by placing a light in

$$UH(t_1) \cap \bigcap_{1}^{k-1} C(d_i, t_{i+1})$$

The existence of such an intersection region follows from the fact that all the t_i are horizontal and all the d_i have positive slope. In fact it suffices to place the light at the intersection point of t_1 and d_{k-1}. In a similar way by placing a second light symmetrically reflected about the origin (assuming the origin lies at the center of $FS(M)$) we may illuminate the rest of the region outside the convex hull of $FS(M)$ as well as the upper left sides of the streets. It remains to illuminate the avenues and these can be taken care of the same way with two additional lights, one above and the other below $FS(M)$. Given the n blocks, we can determine in O(n) time which ones play a role in determining the half-planes where the four lights must be placed. Finally, their location can be determined in O(1) time. For the unrestricted case of Manhattan rectangles (variable width streets and avenues) the solution region for each light is not determined by two half-planes but may have a description of size O(n). Hence the last step can be done in O(n) time. We therefore have the following theorem:

Theorem 1: A set of n Manhattan rectangles can be illuminated with four point lights. Furthermore, the locations of these four lights can be determined in $O(n)$ time.

3. Equal-Width Rectangles and Isothetic Squares

3.1 Introduction

In [CRU] it was shown that for n equal-width isothetic rectangles (i.e., rectangles with sides parallel to the axes) n+2 point lights are always sufficient to illuminate the free space. Since then Hoffmann and Kaufmann [HK] have shown that for an isothetic polygon with holes and a total of k vertices, [k/4] point lights are always sufficient and sometimes necessary to illuminate the interior of the polygon. Furthermore, they can find a placement for these [k/4] point lights in O($k^{1.5}\log^2(k)\log\log k$) time. This result can be applied to the problems considered in this section for point lights. If we surround the n rectangles or squares with one large enough square we obtain a polygon with holes and $k = 4n + 4$ vertices. Therefore their result shows that n+1 point lights suffice and can be located in $O(n^{1.5} \log^2 (n) \log\log n)$ time. Of course the necessity result no longer holds because our situation is a special case of a polygon with a special type of hole. For the restricted case of n *equal-width* isothetic rectangles it is shown in [CRU] that n+2 point lights suffice. However, they do not address the computational problem of actually placing these n+2 lights. We give in this section a simple algorithm that can locate the n+2 lights in $O(n \log n)$ time but first we discuss one of the sub-problems we will encounter.

3.2 Illuminating Manhattan skylines with two lights

Before considering equal-width isothetic rectangles we show that a so-called Manhattan skyline consisting of n line segments can be illuminated with two point lights in $O(n)$ time. A Manhattan skyline is a one dimensional isothetic (rectilinear) polygonal chain that is horizontally monotonic. In other words the silhouette of isothetic buildings erected on a flat ground. Such a skyline is illustrated in Fig.3. We should remark here that this result has been already established in [ELRS]. However, due to space limitations they only sketched a proof. Here we provide an alternate algorithm in detail that has some merit over their algorithm.

Let B denote the ground line, T the horizontal line through the top of the highest building, R the vertical line through the right-most vertical segment and L the vertical line through the leftmost vertical segment. We say that a Manhattan skyline is illuminated by a set of lights if every point on and above the upper envelope of the union of B and the polygonal chains representing the buildings is visible from at least one light. Denote this set by S (for sky) and denote the boundary of this set by $B(S)$. The following construction will do the job (refer to Fig. 3). An edge of $B(S)$ is called a *lower tab* if it is a local minimum, i.e., at both of its endpoints the polygonal chain turns upwards. Let $m_1, m_2,..., m_k$ denote the midpoints of the lower tabs of $B(S)$. From each such midpoint erect an infinite half line in the positive y direction and rotate this line clockwise about its midpoint until it touches the first point of $B(S)$ other than the midpoint. Denote these diagonal half-lines by d_1, $d_2,..., d_k$. If we locate a light in the intersection of the right half-plane determined by R, the upper half-plane determined by T, and $UH(d_i)$, i=1,2,...,k, we will illuminate S except for the regions above $B(S)$ and below the d_i for i=1,2,...,k, (shaded in Fig. 3) and possibly a portion to the left of the first left staircase of $B(S)$. A left staircase is a polygonal chain monotonically increasing in both the positive x and y directions. It is straightforward to show that all these regions are illuminated with one more light located to the left of L, above T and above a set of half-lines analogous to the d_i but which are rotated with respect to the midpoints $m_1, m_2,..., m_k$ in a counterclockwise direction. This establishes the correctness of the algorithm.

We turn now to the complexity. There are two parts to the algorithm. The first part computes $d_1, d_2,..., d_k$, all the clockwise and counter-clockwise support tangents of $B(S)$ emanating from the midpoints $m_1, m_2,..., m_k$ of the lower tabs. The second computes a point in the intersection of $O(n)$ half-planes. To solve the first part we note that the first edge of the upper convex hull of the portion of $B(S)$ from m_i to r, (the right-most vertex of $B(S)$) determines the supporting tangent d_i. Therefore to compute all the d_i it suffices to make the m_i dummy vertices of $B(S)$ and compute the shortest path tree from r to all vertices of $B(S)$ via S. Because $B(S)$ is monotonic the shortest path tree contains the union of all partial convex hulls of partial chains of $B(S)$ that start at r. Furthermore, the shortest path tree of a monotonic chain can be computed simply in $O(n)$ time using the Graham

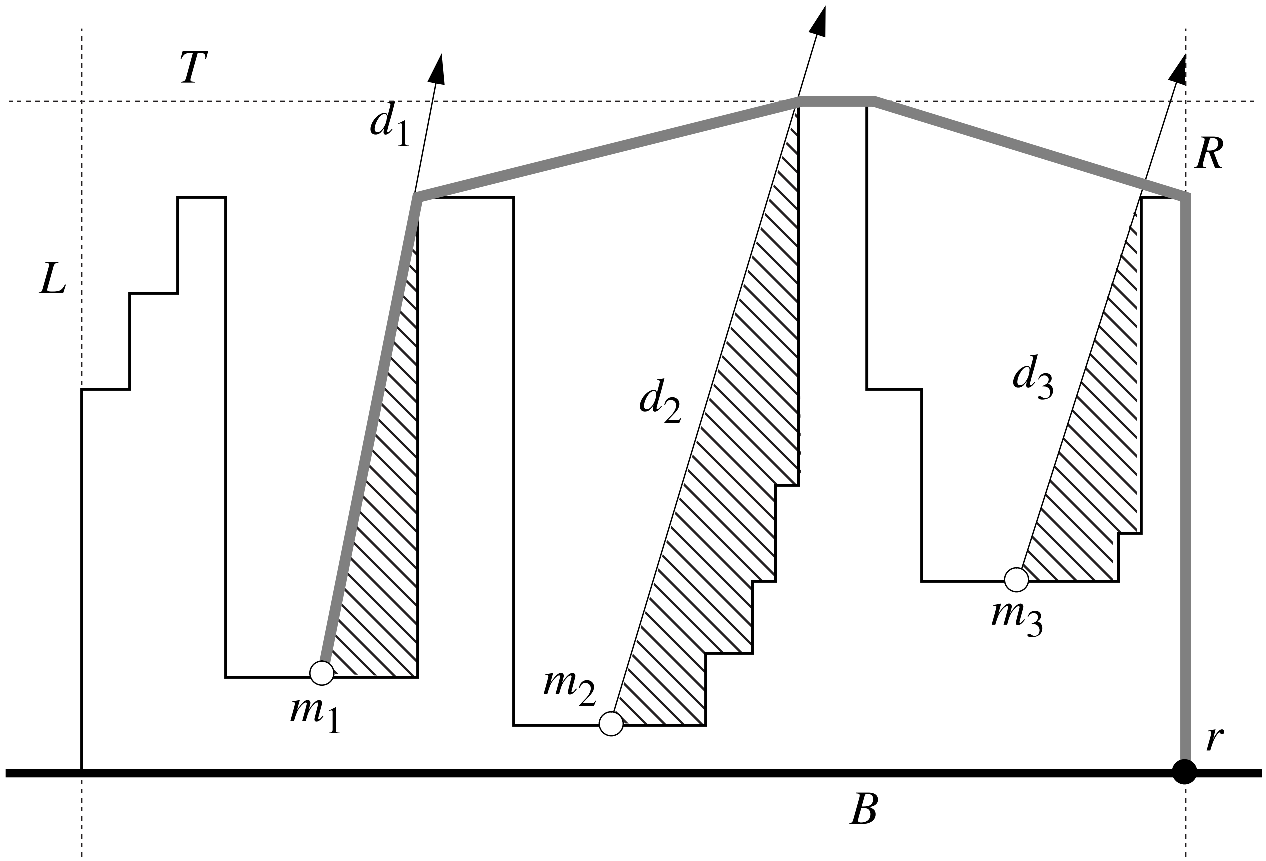

Fig. 3 Illustrating how two point lights (one to the right of R, the other to the left of L, both sufficiently high) are sufficient to illuminate a Manhattan skyline.

scan [ETb]. The second part of the algorithm can be solved in O(n) time by simply finding the highest intersection points of L and R with their corresponding diagonal half lines. We therefore have the following lemma:

Lemma 1: A Manhattan skyline can be illuminated with two point lights in O(n) time.

From the above discussion it follows that if we do not care about the elevation of these two lights then the Graham scan is not needed in the algorithm described above. Indeed in the simpler algorithm of [ELRS] it suffices to place the light on R, for example, at an elevation above the highest intersection point of R with any diagonal half-line. In [ELRS] this is accomplished by essentially computing the intersection point of R with a line that passes through the intersection point of L and T and has positive slope $m = (T\text{-}B)/(h/2)$ where h is the length of the shortest horizontal edge of $B(S)$. While our algorithm is slightly more complicated, the use of the Graham scan will find two lights with possibly much smaller elevations.

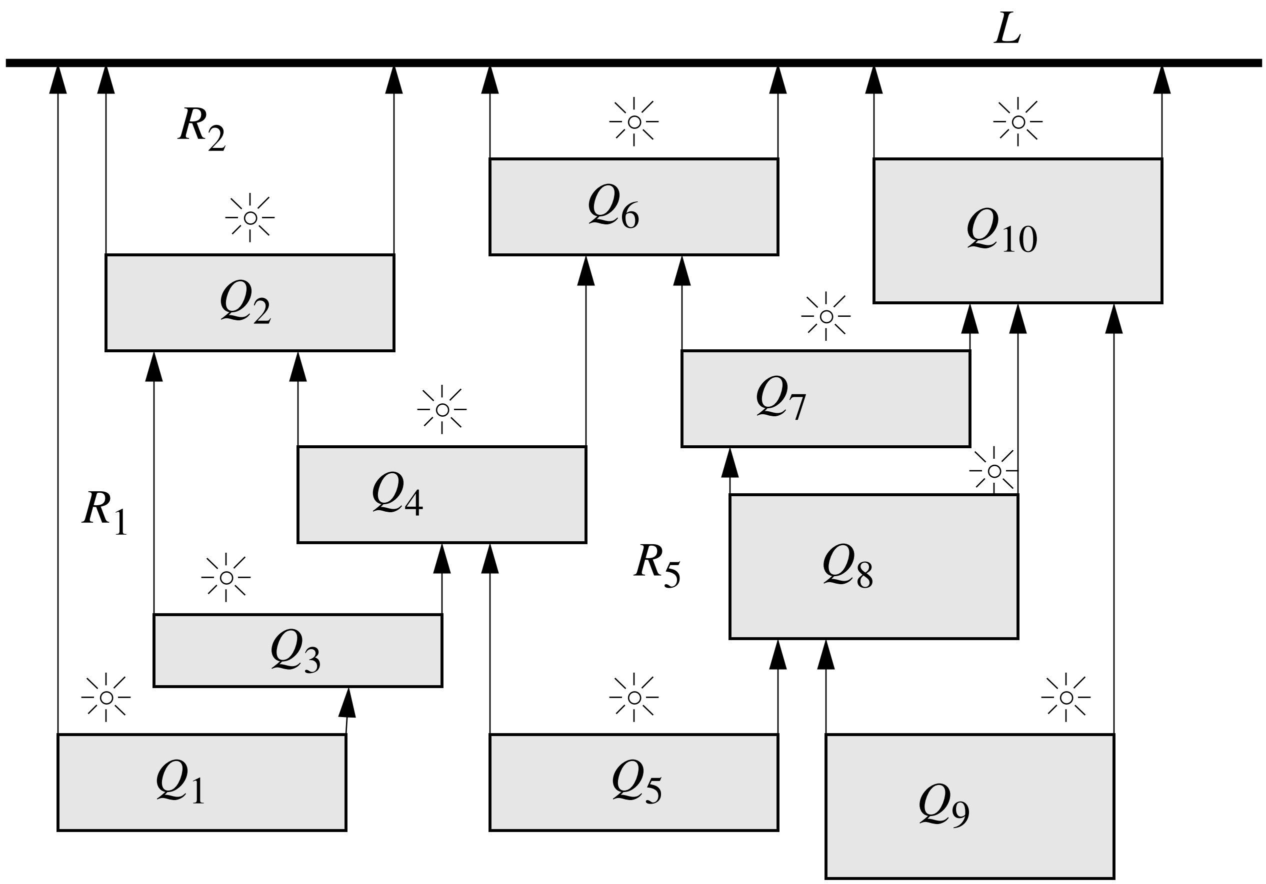

Fig. 4.1Illustrating the construction of the *next-element-subdivision* in the [CRU] proof.

3.3 Illuminating equal-width rectangles

We show that the constructive proof in [CRU] can be implemented in O(n log n) time. First we draw a line L above all the rectangles in O(n) time (see Fig. 4). Then we construct the *next-element-subdivision* from the bottom in an upwards direction using the algorithm of Edelsbrunner, Overmars, and Seidel [EOS] in O(n log n) time. In each pyramidal region so created we can easily find the highest edge in a total O(n) time to place n lights, one above each rectangle. Finally, to place the two lower lights we have one instance of a Manhattan skyline which we can do in O(n) time by lemma 1. We therefore have:

Theorem 2: A configuration of n equal-width isothetic rectangles can be illuminated with n+2 lights in O(n log n) time.

We should note that once the next element subdivision has been computed we can not only find a single placement for the lights in linear time but within the same time bound we can compute a description of a maximal set of light placements. To do this we compute the respective *entire kernels* of the above placement. It suffices to realize that we are dealing with polygonal domains and the kernel of such a domain can be computed in linear time with the algorithm of Lee & Preparata [LP]. Furthermore, for the case of the pyramids an even simpler O(n) time algorithm can be used [TE].

3.4 Illuminating squares

In the illumination scheme of [CRU] there may be more than one rectangle that can be completely projected onto L by the next-element-subdivision. Call these rectangles *exposed.* If that is the case some of these lights are redundant and can be discarded. It suffices to realize that if we perform a next-element subdivision in the opposite direction we obtain a Manhattan skyline which can be illuminated with two lights. These two lights will illuminate all the tops of the exposed rectangles. If there are k exposed rectangles we can thus save k-1 lights with this trick. For squares of equal size we can exploit this property due to the following lemma.

Lemma 2: For $n > 5$ at least one of the four axial directions contains at least three *exposed* squares.

Proof: Let S be a set of $n > 5$ squares. Suppose that all axial directions contain at most 2 exposed squares. Let n_1 and n_2 be the two squares whose upper sides are exposed. Without loss of generality assume that n_1 lies above and to the left of n_2. Let p be the uppermost square in S - $\{n_1,n_2\}$. Since two squares already have their upper sides exposed p cannot be so exposed. Furthermore, p must lie below the line passing through the lower side of n_1 since otherwise the square which prevents the upper side of p from being exposed would be higher up than p contradicting the assumption. Moreover, if p lies to the right of n_2 then it must lie below the line passing through the lower side of n_2, otherwise it would be the third square also illuminated from the top contradicting the assumption. Since p is the uppermost square in S - $\{n_1,n_2\}$ this must be true for all squares in S - $\{n_1, n_2\}$. This implies that the left side of $n_1$and the right side of n_2 are exposed.

Let w be the leftmost square in S - $\{n_1,n_2\}$. The left side of w must be exposed. By an argument similar to above, all squares in S -$\{n_1,n_2, w\}$ must lie to the right of w and therefore the lower side of w is exposed. Let p be the leftmost square in S - $\{n_1, n_2, w\}$ which blocks the upper side of w. Then the left side of p must be exposed because otherwise it would imply there exists a square to the left of p, a contradiction. But the left sides of n_1 and w are exposed

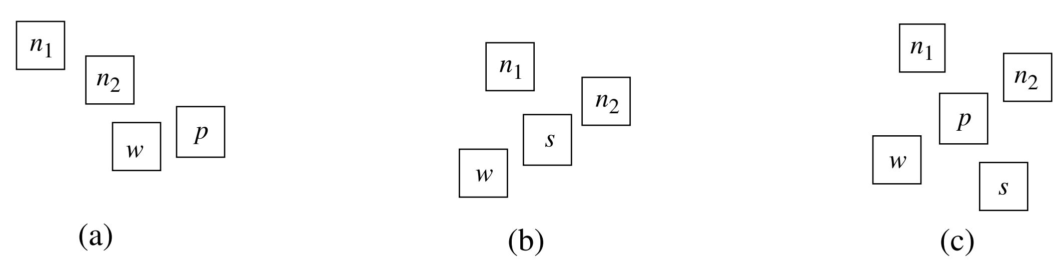

Fig. 5 Illustrating the blocking relations in the proof of lemma 2.

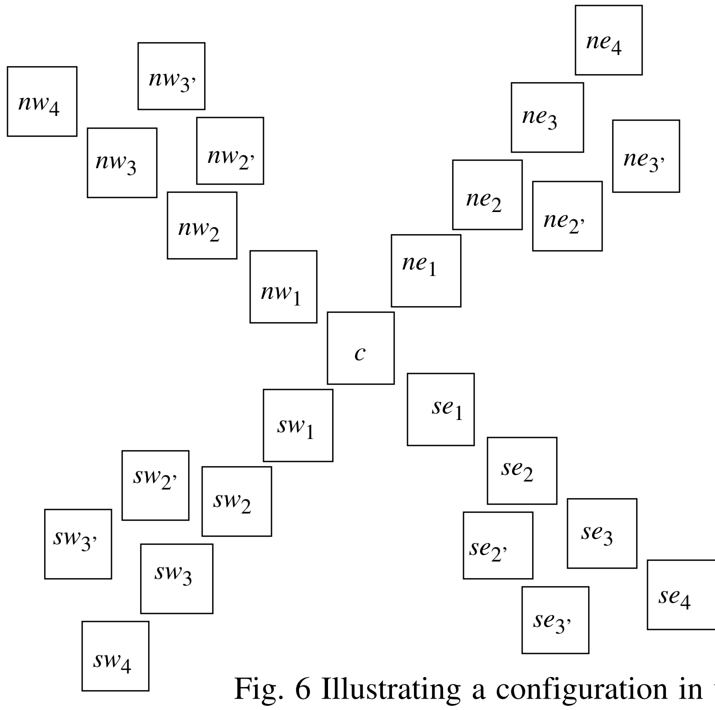

Fig. 6 Illustrating a configuration in which no axial direction contains more than three exposed squares.

so there can be no such p. Thus the upper side of w is blocked either by n_1 or n_2. We will show that it is blocked by n_1. Therefore suppose w is blocked by n_2 and not n_1 and refer to Fig. 5a. Then the lower side of n_1 is exposed. Let p be the lowest square in S - $\{n_1, n_2, w\}$. Since p lies to the right of w and below n_2, the lower side of p must also be exposed. But we already have two squares, namely w and n_1, with lower sides exposed. Therefore the upper side of w is blocked by n_1.

Let s be the lowest square in S - $\{n_1, n_2, w\}$. The lower side of s must be exposed. By an argument similar to above, all squares in S - $\{n_1, n_2, w\}$ must lie above s. If the right side of s is not exposed then it must be blocked by n_2. See Fig. 5b. But then the lower side of n_2 is exposed. So the right side of s must be exposed. Let p be the lowest square in S - $\{n_1, n_2, w, s\}$ which blocks the left side of s. Then the lower side of p is exposed. So the left side of s must be blocked by w. Notice that either s or n_2 must be the right-most square in S and in either case, no other square may have its upper side blocked by s or n_2.

All remaining squares have no exposed sides and they must lie below n_1, to the right of w, above s and to the left of either s or n_2. Let p be the uppermost square in S - $\{n_1, n_2, w, s\}$. The upper edge of p must be blocked by n_1. Thus the left side of w must be to the left of n_1 and since the upper side of w is blocked by n_1, the left side of p must also be blocked by w. Since the left side of s is blocked by w, the lower side of p must be blocked by s and, consequently, s cannot be the right most square. Let q be the uppermost square in S - $\{n_1, n_2, w, s, p\}$ which blocks the right side of p. Then the upper side of q is exposed. Thus the right edge of p is blocked by n_2. There is only one choice for p which is that shown in Fig. 5c. Q.E.D.

It is not the case that for $n > 5$ at least one of the four axial directions contains more than three exposed squares. Given k we show how to construct an arrangement of $n = 8k$-7 squares such that only three squares are exposed in each direction. Figure 6 shows an example for $k = 4$. The arrangement consists of a center square c and four arms. We will show how to construct the south-east arm; the construction for the north-east, north-west and south-west arms is similar. Let $c = se_0 = ne_0 = nw_0 = sw_0$. For $1 \le i \le k$, place square se_i such that the north-west corner of se_i is distance ε above and to the right of the south-east corner of se_{i-1} where $\varepsilon = 1/10k$. Now place square $se_{2'}$ such that the north-west corner of $se_{2'}$ is distance δ below the lower side of se_2 and to the left of the right side of se_1 where $\delta = 1/20k$. For $3 \le i \le k$-1, place $se_{i'}$ such that the north-west corner of $se_{i'}$ is distance δ below and to the left of the south-east corner of $se_{i-1'}$.

From the above construction it follows that no sides of c are exposed. Let SE be the set of squares in the south-east arm. We will show that the only exposed sides among the squares in SE are the right side of se_k and the lower sides of $se_{k-1'}$ and se_k. The situation in the remaining arms is similar. To see that no left side of a square in SE is exposed note that the distance from the lower side of c to the lower side of sw_k is $k(1 + \varepsilon) = k + 1/10$ and the distance from the lower side of c to the lower side of $se_{k-1'}$ is $k + (k - 2)\delta - 2\varepsilon = k + 1/20 - 6/20k$, i.e. the lower side of sw_k is lower than the lower side of $se_{k-1'}$. To see that no upper side of a square in SE is exposed note that the distance from the right side of c to the left side of se_k is $k(1 + \varepsilon) - 1 = k - 9/10$ and the distance from the right side of c to the right side of $ne_{k-1'}$ is $k + (k - 2)\delta - 2\varepsilon = k + 1/20 - 6/20k$, i.e. the right side of $e_{k-1'}$ is further right than the left side of se_k. To see that the only square in SE whose right side is exposed is se_k note that the distance from the lower side of c to the lower side of se_k is $k(1 - \varepsilon) = k - 1/10$ and the distance from the lower side of c to the upper side of $se_{k-1'}$ is $k + (k\text{-}2)\delta - 2\varepsilon - 1 = k - 19/20 - 6/20k$, i.e. the lower side of se_k is lower than the upper side of $se_{k-1'}$. It is clear that the only squares in SE whose lower sides are exposed are $se_{k-1'}$ and se_k. The argument for the squares in the remaining arms is similar. We therefore have the following result:

Theorem 3: Given a collection of n isothetic squares the free space can be illuminated with at most n point lights in $O(n \log n)$ time.

Proof: By lemma 2 it follows that it suffices to apply the previous algorithm in all four isothetic directions and chose the direction which minimizes the number of lights.

4. Isothetic Quadrilaterals

In this section we will show that given a set $S = \{Q_1, Q_2, ..., Q_n\}$ of n pairwise disjoint isothetic but otherwise arbitrary quadrilaterals in the plane we can illuminate the free space with $2n$ *vertex* lights that can be found trivially in O(n) time. Consider the isothetic quadrilaterals in Fig. 7. We place a lights on every north-west and south-east corner of every quadrilateral. Trivially, we have 2n vertex lights which can be located in $O(n)$ time and obviously they illuminate all the sidewalks, i.e., the entire boundary of the free space. We now show that this placement also illuminates the *entire* free space.

The following lemma is useful to this end.

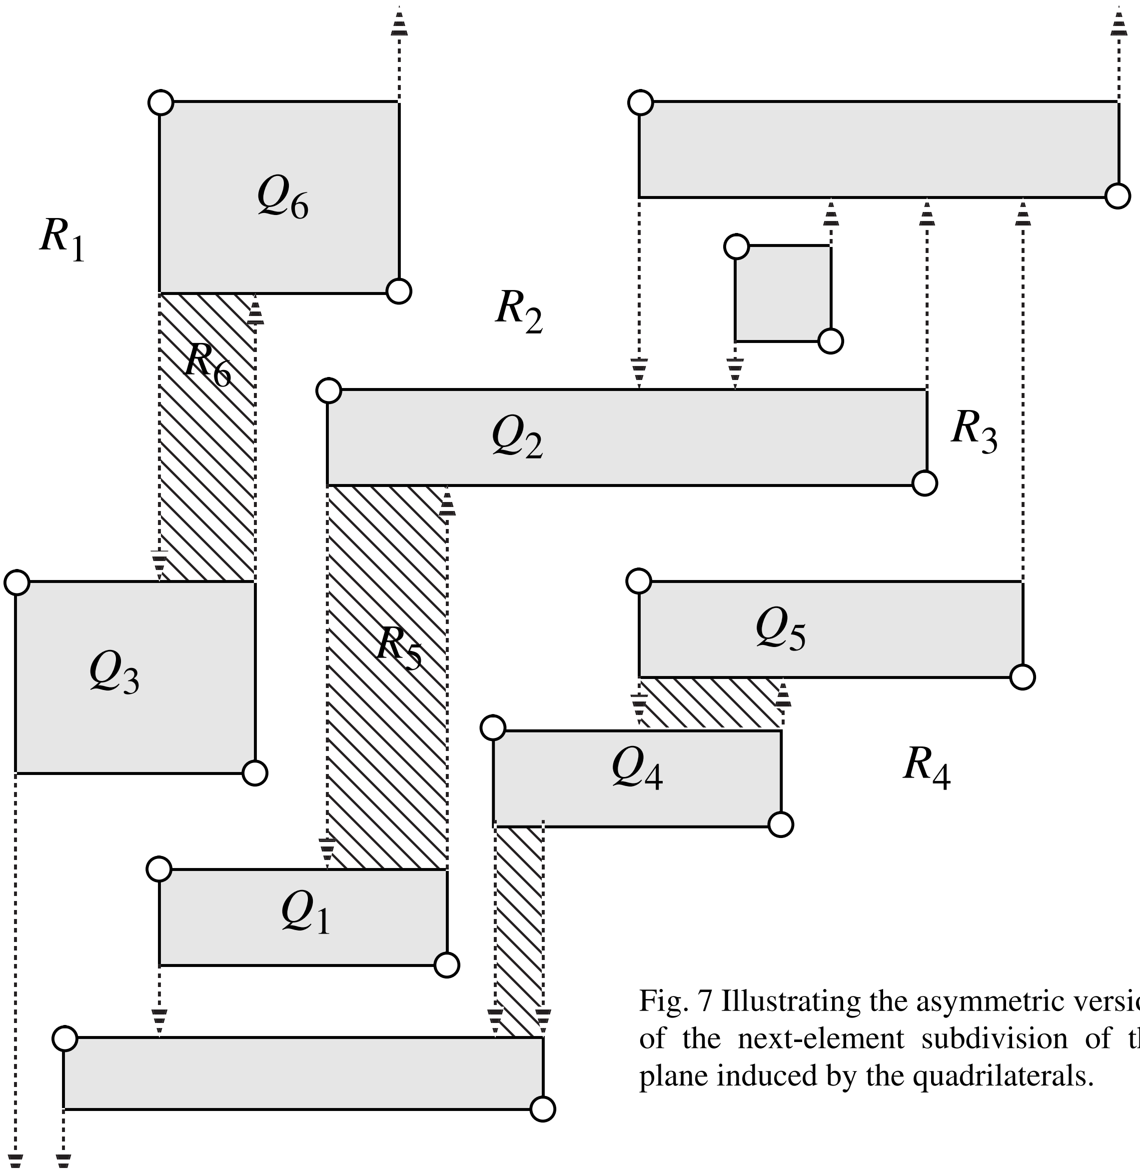

Fig. 7 Illustrating the asymmetric version of the next-element subdivision of the plane induced by the quadrilaterals.

Lemma 3: Given a possibly unbounded simple polygonal region P, if a light is placed on every reflex vertex of P, then the entire region is illuminated.

Proof: Chazelle [Cha] has proved that given a simple polygon of n vertices there exists a partition of the polygon into at most $n+1$ convex pieces by sequentially adding from each reflex vertex a chord that hits either a polygon edge or a chord formed in a previous step. Since each piece is convex and has a vertex coincident on a reflex vertex of the original polygon, it follows that if we place a light on every reflex vertex of the polygon, its entire interior is illuminated. This same argument applies equally to the case of possibly unbounded polygonal regions. Q.E.D.

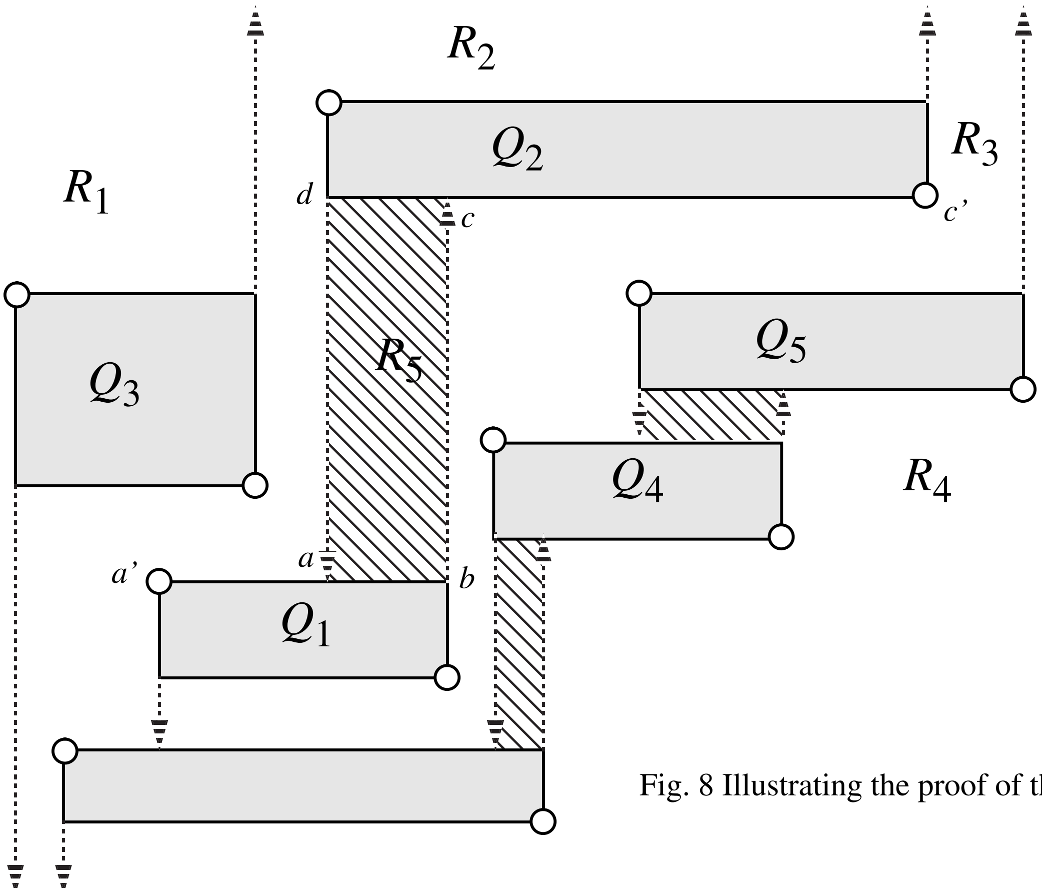

Fig. 8 Illustrating the proof of theorem 8.

First we create a planar subdivision of the free space which is actually an asymmetric version of the next-element subdivision, by extending the left edge of every quadrilateral in a downwards direction and the right edge of every quadrilateral in a upwards direction (see Fig. 7). This subdivision consists of only three types of *staircase* (monotonic in both the x and y directions) polygonal components: (1) unbounded regions such as R_1, R_2, and R_4, (2) bounded regions consisting of more than four edges such as region R_3, and (3) rectangular regions determined by an upper edge of one quadrilateral and a lower edge of another quadrilateral, such as regions R_5 and R_6. Consider the regions of type (1) and (2). By the construction of the subdivision and the placement of the lights it follows that these regions contain a light on every reflex vertex. Therefore by lemma 3 they are illuminated. Consider now a type (3) region such as R_5 in Fig. 8. Region R_5 for example is determined by a right portion of the upper edge of Q_1 and the left portion of the lower edge of Q_2. Denote the vertices of R_5 by a, b, c and d. Let a' be the north-west vertex of Q_1 and let c' be the south-east vertex of Q_2. Now, considering the quadrilaterals as obstacles, construct the geodesic path between a' and d that is *homotopic* with the path $a'ad$. Similarly, construct the geodesic path between b and c' that is *homotopic* with the path bcc'. By construction, the region, denoted by ER_5, bounded by these geodesic paths and the segments $a'b$ and $c'd$ contains R_5. Furthermore, by construction all reflex vertices of ER_5 contain lights. Therefore by lemma 3 the type (3) regions are also illuminated. Therefore we have the following theorem.

Theorem 4: Given a set of n pairwise disjoint isothetic quadrilaterals in the plane a set of $2n$ vertex lights that illuminate all the free space can be located in O(n) time.

5. Arbitrary Quadrilaterals

5.1 Geometric preliminaries

Given a set of n pairwise disjoint but otherwise arbitrary quadrilaterals in the plane $S = \{Q_1, Q_2, \ldots, Q_n\}$ we want to find the number of vertex lights that are always sufficient and sometimes necessary to illuminate the entire free space. That $2n$ vertex lights are sometimes necessary to illuminate the free space among n pairwise disjoint quadrilaterals in the plane is established by the generic example in Fig. 9. In this section we prove that $2n$ are always sufficient and propose an algorithm that computes the positions of these lights. In order to do this it is convenient to divide the free space in two components: the *interior* and *exterior* of the convex hull *CH(S)* of S. The free space in *CH(S)* can be viewed as a polygon with holes, the construction of which will be detailed below. Then we invoke the classical result for illuminating a polygonal art gallery with holes [O'R]. The illumination of the exterior of *CH(S)* is accomplished by exploiting the lights that have already been placed in the interior and by placing some additional lights on the boundary of *CH(S)* in such way that we guarantee that the conditions of the following lemma are satisfied.

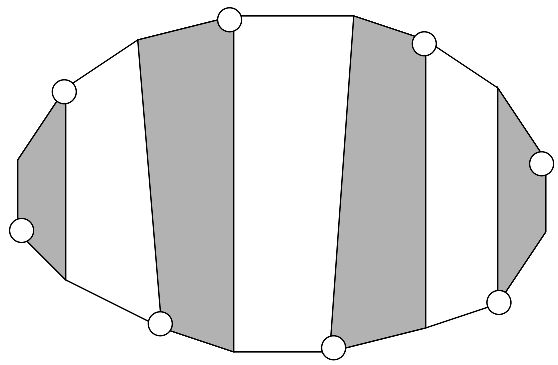

Fig. 9 There exist configurations of n quadrilaterals that require $2n$ vertex lights.

Lemma 4: Let P be a given convex polygon P with n vertices and n non overlapping triangles placed one flush with each edge and internal to the polygon. Assume furthermore that for each triangle we have one vertex light on one of its vertices and that all triangle edges are opaque obstacles except those of P. Then the space outside the polygon is illuminated.

Proof: Consider the boundary of the polygon described by a list of counterclockwise consecutive edges $(e_1, e_2, \ldots, e_n)$. For each $1 \leq k \leq n$ let l_k, r_k denote the "left" and "right" edges of *triangle* T_k, respectively, as viewed from the third edge e_k (see Fig. 10). Let P' denote the simple polygon obtained by substituting each edge e_k of P by l_k and r_k.

Let B_k denote the directed ray, with direction away from P, bisecting the external angle between r_k and l_{k+1} exterior to the polygon P'. Let R_k denote the (possibly bounded) convex region determined by intersecting the left open half-plane determined by the line colinear with B_{k-1} and the right half-planes determined by the lines colinear with e_k and B_k. It follows that $R^2 - P = \cup R_k$. (Note that if some of the B_k intersect each other then some regions of R^2

are covered by even more than one R_k.)

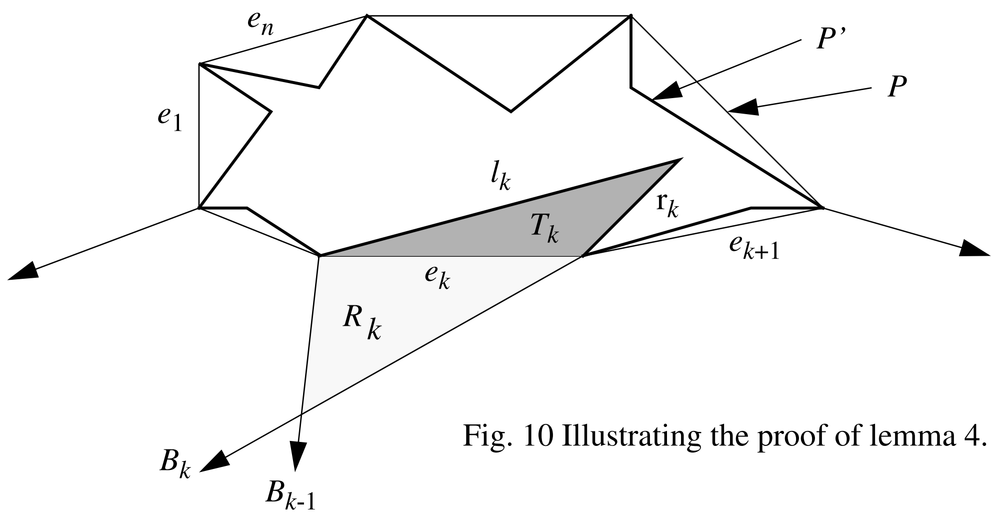

Fig. 10 Illustrating the proof of lemma 4.

Let x be in R^2 - P, so that there exists a j such that $x \in R_j$. Then x is illuminated because $T_j \cup R_j$ is convex, it has no obstacles in its interior and there is a light on the boundary of this region. Q.E.D.

Remark: This result is true even if some triangles are degenerate.

5.2 Illuminating sets that do not contain cut-off quadrilaterals.

Definition: Given a finite set S of quadrilaterals in the plane we say that a quadrilateral Q in S *cuts off S* (or Q is a *cut-off quadrilateral* of S) if there are two *consecutive* vertices v, w of Q on the boundary of $CH(S)$ but the edge vw is not an edge of $bdCH(S)$ (see Fig. 11).

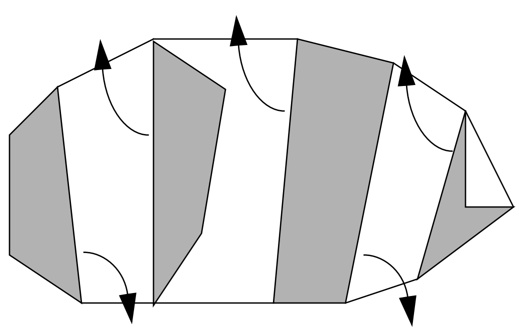

Fig. 11 Some examples of cut-off quadrilaterals.

First consider the case in which there are no cut off quadrilaterals in S. We can classify the quadrilaterals by examining their configuration with respect to the boundary of $CH(S)$ into the following five cases (see Fig. 12):

(1) No vertex is on $bdCH(S)$.

(2) Precisely one vertex is on $bdCH(S)$.

(3) Precisely two vertices are on $bdCH(S)$ *and* they are consecutive vertices of Q.

(4) Precisely two vertices are on $bdCH(S)$ *and* they are non-consecutive on Q.

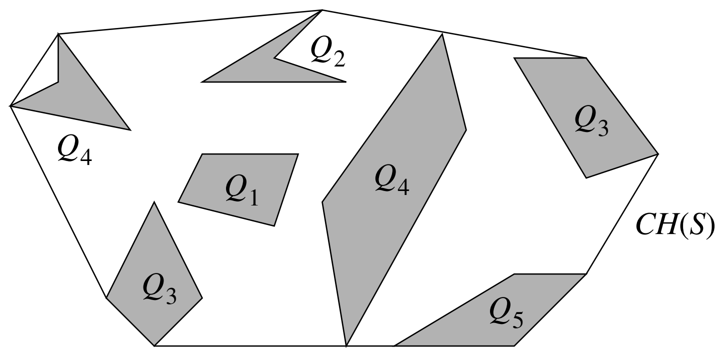

Fig. 12 Illustrating the five types of non cut-off quadrilaterals.

(5) Precisely three vertices are on $bdCH(S)$. Then there must be two pairs of consecutive vertices among them, so that the corresponding edges must also be on $bdCH(S)$ and they are consecutive on Q.

Remark: The case with all four vertices of Q on $bdCH(S)$, and without cut-off quadrilaterals, happens if and only if S has exactly one convex quadrilateral.

First we construct a simple polygon P with holes. In case (1) the holes will be the quadrilaterals of S. We denote by P_1 the exterior boundary of P. We start with $bdCH(S)$ as P_1 and for each vertex q of $CH(S)$ we look at the quadrilateral Q which contains q.

If Q is as in case (2) then we cut P_1 and Q at q and "glue" these pieces together to form the new simple polygon P_1. In order to do this we duplicate q (see Fig. 13a).

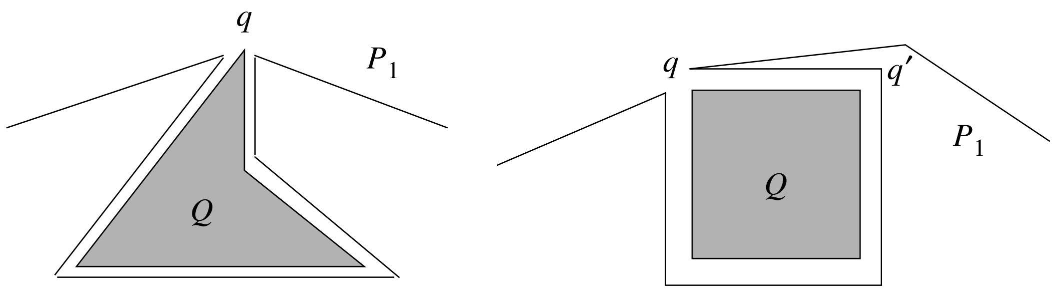

(a) (b)

Fig. 13.

If Q is as in case (3) then we cut P_1 and Q at q and "glue" these pieces together to form the new simple polygon P_1. In order to do this we duplicate q, q' and the edge qq' where q' is the other vertex of Q on $bdCH(S)$ (see Fig. 13 (b)). The point q' will not be considered any more in the construction of P_1.

If Q is as in case (4) then let us cut P_1 and Q at q and "glue" these pieces together to form the new simple polygon P_1. In order to do this we duplicate q and q' where q' is the other vertex of Q on $bdCH(S)$ (see Fig. 14 (a)). The vertex q' is not considered again in the construction of P_1.

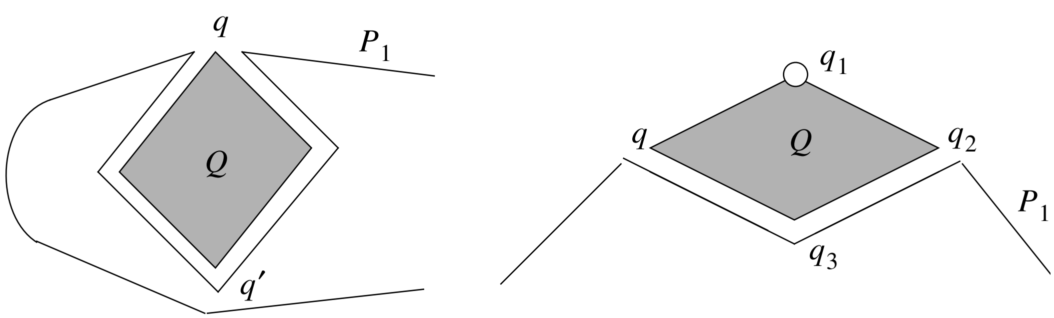

(a) (b)

Fig. 14.

If Q is as in case (5) let q_1 and q_2 be the other vertices of Q such that the edges qq_1 and q_1q_2 are on $bdCH(S)$. Then we substitute the chain qq_1q_2 in P_1 by qq_3q_2 where q_3 is the fourth vertex of Q (see Fig. 14 (b)). Then we place a light on the vertex q_1 in order to illuminate the exterior of $CH(S)$. The point q_1 is not considered again in the construction of P_1.

Let us introduce some notation for the polygon P which has been constructed. H will denote the number of holes in P, V will denote the number of vertices of P and L_1 the set of lights that we have put in the construction.

If n_1, n_2, n_3, n_4 and n_5 are the number of quadrilaterals in cases (1), (2), (3), (4) and (5) respectively then we have $H = n_1$, $V = 4n_1 + 5n_2 + 6n_3 + 6n_4 + 3n_5$ and $|L_1| = n_5$.

Now we turn to the problem of illuminating the free space in the interior of P. In order to do this we use a classical technique due to Fisk [Fi]: triangulate P to obtain $T(P)$ and 3-color the resulting triangulation. Note that a triangulation of a non-simple polygon is not always 3-colorable but by using the cutting technique suggested by O'Rourke [O'R] we can accomplish this by duplicating $2H$ vertices in the polygon P.

For each triangle let us put a light on the vertex with the least frequent color. Therefore we can illuminate the interior of P with the set L_2 of lights, and it is known [O'R] that:

$$|L_2| \leq \left\lfloor \frac{V + 2H}{3} \right\rfloor$$

Theorem 5: Given n pairwise disjoint quadrilaterals in the plane whose convex hull has no cut-off quadrilaterals, then $2n$ vertex lights are always sufficient to illuminate the free space.

Proof: Following the construction described previously we have that the complete set of lights $L = L_1 \cup L_2$ has cardinality:

$$|L| = |L_1| + |L_2| \leq n_5 + \left\lfloor \frac{(4n_1 + 5n_2 + 6n_3 + 6n_4 + 3n_5) + 2n_1}{3} \right\rfloor$$

which is equivalent to:

$$\left\lfloor \frac{6(n_1 + n_3 + n_4 + n_5) + 5n_2}{3} \right\rfloor \leq 2(n_1 + n_2 + n_3 + n_4 + n_5) = 2n$$

To conclude the proof, we have to show that the free space is illuminated with these $2n$ lights. The free space is divided into two parts: the inside of polygon P and the outside of *CH(S)*. The inside of polygon P is illuminated because it is triangulated and each triangle has a light on one of its vertices. We now show that the exterior is also illuminated.

Let e denote an edge of *bd CH*(*S*) then:

(1) If e is an edge of P consider the triangle of *T(P)* such that it has e as an edge. We know that there exists a light on one its vertices.

(2) If e is not an edge of P then it is an edge of a quadrilateral Q in case (5). By construction of P there must be one light on one vertex of the edge e.

Therefore we have a non overlapping triangle on each edge of *CH*(*S*) with a light on one of its vertices. It follows from lemma 4 that the exterior of *CH*(*S*) is also illuminated. Q.E.D.

From the proof of theorem 5, the following algorithm for finding a placement for the lights can be directly derived.

ALGORITHM 1

Input: A set S of n pairwise disjoint quadrilaterals without cutoffs.

Output: A placement of $2n$ lights

begin

Step 1: Compute the convex hull $CH(S)$.

Step 2: Compute the polygon P and the set of lights L_1.

Step 3: Construct the triangulation $T(P)$ of P along with the holes it contains.

Step 4: Construct a polygon P^* where P^* is a simple polygon obtained from P by cutting H diagonals of $T(P)$. Note that since the cuts are made at diagonals of $T(P)$ we also obtain $T(P^*)$.

Step 5: 3-colour $T(P^*)$.

Step 6: Construct the set of lights L_2.

end

Theorem 6: Algorithm 1 computes the position of at most $2n$ vertex lights that illuminate the free space between n pairwise disjoint quadrilaterals without cutoffs in $O(n \log n)$ time.

Proof: Step 1 can be done in $O(n \log n)$ time using one of the classical algorithms [PS]. Step 3 can also be done in $O(n \log n)$ time using a modification of the algorithm of Garey, Johnson, Preparata and Tarjan [AAP]. Since all other steps can be done in linear time, this algorithm has complexity $O(n \log n)$ and computes correctly the position of at most 2n lights by theorem 5. Q.E.D.

5.3 General quadrilaterals

In this sub-section we explore the general situation for quadrilaterals, i.e., the case when cut-off quadrilaterals may be present in S.

Theorem 7: Let S be a given set of n pairwise disjoint quadrilaterals in the plane. Then $2n$ vertex lights are always sufficient to illuminate the free space.

Proof: By induction on the number of quadrilaterals:

If $n=1$ the solution is trivial. Place two lights on one of the two non consecutive vertices of the quadrilateral.

Assume that the induction hypothesis is true for n quadrilaterals, then we consider a set S of $n+1$ quadrilaterals and:

a) If there is no cut-off quadrilateral in S then we have proved that $2(n+1)$ lights are always sufficient by theorem 5.

b) If there exists at least one cut-off quadrilateral Q in S then we divide S into two linearly separated sets S_1 and S_2 with $|S_1|, |S_2| < |S|$. This can be done in the following and natural way. By definition Q contains two consecutive vertices, let us say a and b on $bdCH(S)$. Then the line l through a and b divides S into two sets, each one in one side of l, and neither set is empty because the segment ab crosses the interior of $CH(S)$. See Fig 15.

By the induction hypothesis, the free space among the quadrilaterals of S_1 (respectively S_2) can be illuminated with at most $2|S_1|$ (respectively $2|S_2|$) vertex lights.

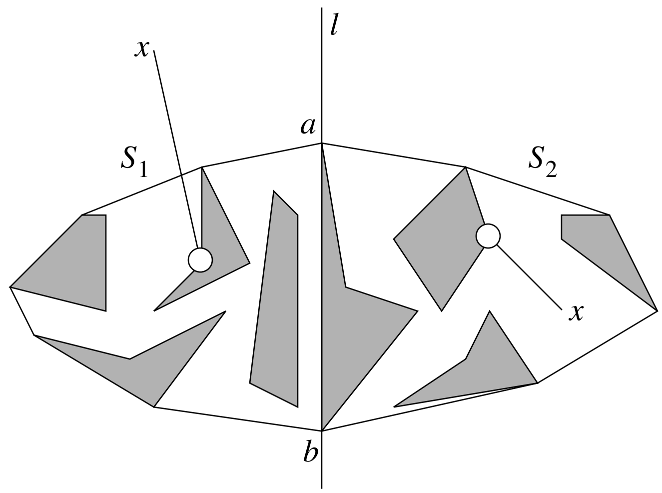

Fig. 15 Illustrating the proof of theorem 7.

Now the remaining question is whether the free space among quadrilaterals in S is illuminated with the $2|S_1| + 2|S_2| = 2(n+1)$ vertex lights that illuminate *indenpendently* the free space among the quadrilaterals in S_1 and S_2. The answer to this question is yes and it is shown as follows. Let x be any point in R^2 - S (see Fig. 15). If x is in the half plane containing the set S_1 then x is illuminated by one light placed in S_1 else x is illuminated by one light placed in S_2. Q.E.D.

The proof of this theorem can be used to obtain an algorithm which finds the positions of these *2n* lights.

ALGORITHM 2

Input: A set S of n pairwise disjoint quadrilaterals in the plane.

Output: The locations of a set of lights L.

begin

Step 0: Initialize the set of lights L to be the empty set.

Step 1: Sort the vertices of the quadrilaterals in S by their x-coordinate.

Step 2: a) Compute the $CH(S)$.

b) Decide the cutoff situation of S:

(1) If $|S| = 1$ add two non consecutive vertices of the unique quadrilateral in S to the set of lights L.

(2) If S has no cut-off quadrilaterals apply Algorithm 1 to S and add the output set of lights to L.

(3) If S contains a cut-off quadrilateral then divide S into two subsets S_1 and S_2 and compute recursively step 2 with S_1 and S_2.

end

Theorem 8: Algorithm 2 computes the position of at most $2n$ vertex lights that illuminate the free space among n pairwise disjoint quadrilaterals in $O(n^2)$ time.

Proof: Step 1 can be done in $O(n \log n)$ time.

Step 2 (a) can be computed in linear time by using Graham's algorithm [PS] calculating the upper and lower hulls using the sorting done in step 1.

Step 2 (b) takes $O(|S|)$ time in case 3 and $O(|S| \log |S|)$ time in case 2, but in this case there are no more recursive calls to step 2, so the worst case is the first one. For this case the recursive formula is

$$T(|S|) = O(|S|) + T(|S_1|) + T(|S_2|)$$

and its solution is $T(n) = O(n^2)$.

Finally, the output is the right set of lights by theorem 3. Q.E.D.

6. Conclusions

We close with several open problems. We showed that given n isothetic equal-width quadrilaterals, the entire free space can be illuminated with $n+2$ *point*-lights that can be placed in $O(n \log n)$ time. This contrasts with the [HK] result that locates $n+1$ lights in $O(n^{1.5} \log^2 (n) \log \log n)$ time. Is it possible to place $n+1$ lights in $O(n \log n)$ time? We showed that given n arbitrary quadrilaterals the free space can be illuminated with $2n$ *vertex*-lights that can be placed in $O(n^2)$ time. Can these lights be placed in $O(n \log n)$ time? Finally, what is a non-trivial lower bound on this problem? We conjecture the answer is $\Omega(n \log n)$.

7. Acknowledgments

The authors thank Jit Bose and the participants of a research seminar on computational geometry for discussions on these topics. The idea of exploring the subject for isothetic squares was suggested by Guylaine Lefebvre.

8. References

[AAP] T. Asano, T. Asano and R. Y. Pinter, "Polygon triangulation: efficiency and minimality," *Journal of Algorithms*, vol. 7, 1986, pp. 221-231.

[AT] D. Avis, and G. T. Toussaint, "An efficient algorithm for decomposing a polygon into star-shaped pieces," *Pattern Recognition,* vol. 13, 1981, pp. 295-298.

[BS] I. Bjorling-Sachs and D. Souvaine, "A tight bound for guarding general polygons with holes," Tech. Report LCSR-TR-165, Rutgers University, Dept. Computer Science, May 1991.

[Cha] B. Chazelle, “Computational geometry and convexity,” Ph.D. thesis, Yale University, 1980.

[Ch] V. Chvatal, "A combinatorial theorem in plane geometry," *Journal of Combinatorial Theory, Series B,* vol. 18, 1975, pp. 39-41.

[CRU91] J. Czyzowicz, I. Rival and J. Urrutia, “Galleries and light matchings: fat cooperative guards," in *Vision Geometry,* Contemporary Mathematics, vol. 119, R. A. Melter, A. Rosenfeld and P. Bhattacharya, Editors, 1991, pp. 21-28.

[CRU89] J. Czyzowicz, E. Rivera-Campo and J. Urrutia, “Illuminating rectangles and triangles in the plane”, Tech. Rept. TR-89-50, Department of Computer Science, University of Ottawa, December 1989.

[CRUZ] J. Czyzowicz, E. Rivera-Campo, J. Urrutia and J. Zaks, “Illuminating lines and circles on the plane”, Tech. Rept. TR-89-49, Department of Computer Science, University of Ottawa, December 1989.

[ELRS] H. Everett, K. Lyons, B. Reed and D. Souvaine, “Illuminating squares on a transversal,” *Proc. Third Canadian Conference on Computational Geometry*, Vancouver, Canada, August 6-10, 1991, pp. 118-121.

[EOS] H. Edelsbrunner, M. H. Overmars, and R. Seidel, "Some methods of computational geometry applied to computer graphics," Tech. Rept. F-117, Technical University of Gratz, June 1983.

[ET] H. Everett and G.T. Toussaint, “On illuminating rectangles in the plane”, McGill University memorandum, 1990.

[ETb] H. ElGindy and G. T. Toussaint, “On geodesic properties of polygons relevant to linear-time triangulation,” *The Visual Computer*, vol. 5, no. 1/2, March 1989, pp. 68-74.

[Fi] S. Fisk, “A short proof of Chvatal’s watchman theorem,” *Journal of Combinatorial Theory,* Series B, vol. 24, 1978, p. 374.

[Fe] L. Fejes-Toth, “Illumination of convex discs,” *Acta Mathematica Academiae Scientiarum Hungaricae*, vol. 29, 1977, pp. 355-360.

[HKK] F. Hoffmann, M. Kaufmann and K. Kriegel, "The art gallery theorem for polygons with holes," in *Proc. Foundations of Computer Science 91,* pp. 39-48.

[HK] F. Hoffmann and M Kaufmann, LNCS 484, Springer-Verlag, pp. 239-250.

[HK93] F. Hoffman and K. Kriegel, “A graph coloring result and its consequences for polygon guarding problems," Tech. Rept. B-93-08, Free University of Berlin, June 1993.

[Hof] F. Hoffmann, "On the rectilinear art gallery problem," in *Proc. Int. Colloq. Automata, Languages and Programming 90,* LNCS 443, 1990, pp. 717-728.

[Ho] Honsberger, R., *Mathematical Gems II*, Mathematical Association of America, 1976, pp. 104-110.

[JL] G. F. Jennings and W. J. Lenhart, "An art gallery theorem for line segments in the plane," *Pattern Recognition Letters,* Special Issue on Computational Geometry, G. T. Toussaint, Ed., 1993.

[LP] D. T. Lee and F. P. Preparata, "An optimal algorithm for finding the kernel of a polygon," *JACM*, Vol. 26, 1979, pp. 415-421.

[O'R92] J. O'Rourke, "Computational geometry column 15," *International Journal of Computational Geometry and Applications*, vol. 2, No. 2, 1992, pp. 215-217.

[O'R] J. O'Rourke, *Art Gallery Theorems and Algorithms*, Oxford U. Press, 1987.

[PS] F.P. Preparata and M.I. Shamos, *Computational Geometry: An Introduction*, Springer-Verlag, 1985.

[Sh] T.C. Shermer, "Recent results in art galleries," *Proceedings of the IEEE*, vol 80, No 9, September 1992.

[TE] G.T. Toussaint and H. ElGindy. "Traditional galleries are star-shaped if every two paintings are visible from some common point". Tech. Rept. SOCS-81.10, School of Computer Science, McGill University, March 1981.

[UZ] J. Urrutia and J. Zaks, "Illuminating convex sets", Tech. Rept. TR-89-31, Department of Computer Science, University of Ottawa, September 1989.