Discrete Comput Geom 20:61–78 (1998)

Discrete & Computational Geometry

Visibility with Multiple Reflections*

B. Aronov,[1] A. R. Davis,[2] T. K. Dey,[3] S. P. Pal,[3] and D. C. Prasad[3]

[1]Computer and Information Science Department,
Polytechnic University,
Brooklyn, NY 11201, USA
aronov@ziggy.poly.edu

[2]Division of Computer Science, Mathematics, and Science,
St. Johns University,
Jamaica, NY 11439, USA
davisa@stjohns.edu

[3]Department of Computer Science and Engineering,
Indian Institute of Technology,
Kharagpur, India 721302
{dey,spp}@cse.iitkgp.ernet.in

Abstract. We show that the region lit by a point light source inside a simple n-gon after at most k reflections off the boundary has combinatorial complexity $O(n^{2k})$, for any $k \geq 1$. A lower bound of $\Omega((n/k - \Theta(1))^{2k})$ is also established which matches the upper bound for any fixed k. A simple near-optimal algorithm for computing the illuminated region is presented, which runs in $O(n^{2k} \log n)$ time and $O(n^{2k})$ space for $k > 1$, and in $O(n^2 \log^2 n)$ time and $O(n^2)$ space for $k = 1$.

1. Introduction

Visibility-related problems have been extensively studied, in the diverse disciplines in which they naturally arise, in different ways. In computational geometry and associated research areas alone, O'Rourke [27] reports over 300 articles related to various aspects of visibility. Visibility topics include, among others, problems concerning computation,

* Work on this paper by Boris Aronov has been supported by NSF Grant CCR-92-11541 and a Sloan Research Fellowship. Tamal K. Dey acknowledges the support of NSF Grant CCR-93-21799 while he was at the Department of Computer Science, Indiana University–Purdue University, Indianapolis, IN 46202, USA. S. P. Pal acknowledges the support of a research grant from the Jawaharlal Nehru Centre for Advanced Scientific Research, Bangalore, India.

characterization, and realization of visibility graphs, art gallery problems, shortest path problems, and ray shooting and hidden surface elimination; see, for example, [3], [23], [25], [27], [28], and [31].

Visibility is obviously intimately related to geometric optics, so that not only the issue of direct (straight-line) visibility, but also of visibility with reflection naturally occur here. Indeed, there is a large literature on geometric optics (such as [24], [12], and [5]) and on the chaotic behavior of a reflecting ray of light or a bouncing billiard ball (see, e.g., [4], [14], [18], and [21]). The field is very interesting and has accumulated a surprising number of long-standing open problems, some of which are startlingly simple to state [20]. For example, it is not known if every obtuse triangle admits a cyclic billiard ball path. More closely related to the issues addressed in this paper is the question: "Can any simple polygon bounded by mirrors be completely lit up by a single light bulb placed at an arbitrary point in its interior?" This question was first published in 1969 [19], and was finally settled only very recently [33].

Remarkably, there has been almost no investigation of visibility *with reflection* in the geometric complexity or algorithmic context. For example, reflection is a natural issue in computer graphics, where a common rendering technique is to trace the path of light arriving on each pixel of the screen, backward through multiple reflections [11]. Some simple rendering problems can be solved at the pixel level with hardware assistance, such as using a hardware Z-buffer algorithm. Any application requiring further manipulation of the scene would benefit from this output being placed in a readily accessible data structure rather than an array of pixels. However, it seems that non-pixel-based algorithms have so far involved only direct visibility (see, e.g., [3]).

Reif et al. [30], on the other hand, address the problem of tracing a *single* light ray through a complicated optical system with the purpose of detecting if it ever arrives at the specified destination point. They show that several versions of the problem in three dimensions are undecidable. In addition, they list several restricted two- and three-dimensional versions that are decidable, but provably hard. While they do not deal with the situation considered in this paper, we feel that the problem might be hard even in this simple context. We do not address this question directly, however.

Direct visibility has been investigated extensively over the past several years, and a number of linear-time direct visibility algorithms for simple polygons are known [10], [13]. Among different alternative notions of visibility, k-link visibility comes closest to what we study in this paper. Horn and Valentine introduced this concept, where a point y inside a given polygon is *k-link-visible* from another point x if there exists a k-link polygonal path between them inside the polygon [15]. Link visibility has been extensively studied since then; see, for example, [22], [16], [17], [32], and [8]. However, in contrast to k-link visibility, we further restrict the path so that it may only turn at the boundary of the polygon and, moreover, must obey laws of geometric optics at these reflection points. As a result, the two notions of visibility produce drastically different behavior. For example, we show that, for small k, the complexity of the region lit up with at most k reflections is exponential in k, while the corresponding region of k-link-visible points is bounded by at most n edges, for any k; see, e.g., [32]. A model in which the path may only turn at the boundary, but need not obey the reflection laws corresponds to so-called "diffuse reflection." It was analyzed, for a single reflection, in [1], and for multiple reflections, by Prasad et al. [29].

In a companion paper [1], we investigated the region visible from a point in a simple n-gon bounded by mirrors, when at most one reflection is permitted. We obtained a tight $\Theta(n^2)$ worst-case geometric complexity bound and described a simple $O(n^2 \log^2 n)$ algorithm for computing this set. In the current paper we investigate the case where at most k reflections are permitted. We are interested in both the worst-case complexity of the resulting lit region and in an efficient algorithm for computing the region. We produce an $O(n^{2k})$ upper bound and an $\Omega((n/k - \Theta(1))^{2k})$ worst-case lower bound on this complexity and construct an algorithm with $O(n^{2k} \log n)$ running time, for $k > 1$. The combinatorial complexity bound involves a careful counting argument (it turns out that an upper bound of $O(n^{2k+1})$ is easy, but the proof of the stronger bound is more involved), while the algorithm uses a standard divide-and-conquer approach and follows almost immediately from previous analysis.

Two aspects of our analysis deserve special mention. The first is that the approach to the complexity analysis needed for the case $k > 1$, as described in the current paper, is different from that needed for $k = 1$, as described in [1], and applying the current analysis to the case of at most one reflection yields a bound which is not tight; a unified approach is proposed in [7]. Secondly, the lower and upper complexity bounds match for any fixed k, but diverge when k grows as a function of n. In fact, the lower bound construction breaks down completely for k comparable with n. We have no construction where the complexity of the region lit up with at most k reflections is superquadratic in k, for large k and fixed n.

The remainder of this paper is organized as follows. Section 2 presents some preliminary definitions. Sections 3 and 4 establish the upper and lower bounds, respectively. Finally, Section 5 describes a near-optimal algorithm that computes the visibility polygon with at most k reflections.

2. Preliminaries

Let $P \subset \Re^2$ be a simple n-gon with no three collinear vertices. Let $int(P)$ and $bd(P)$ denote the interior and the boundary of P, respectively. Two points in P are said to be 1-*visible* (or *directly visible*[1]) if the interior of the line segment joining them lies in $int(P)$.

We consider visibility with reflection where the angle of incidence is equal to the angle of reflection. This type of reflection is termed *specular reflection* in computer graphics. For $k > 1$, a point y is said to be k-*visible* from a point x (under specular reflection), if there exist points $p_1, p_2, \ldots, p_{k-1}$ lying in the interiors of edges of P such that a ray emitted from x reaches y after $k - 1$ stages of specular reflection at $p_1, p_2, \ldots, p_{k-1}$, in this order. Since specular reflection at a vertex is not well defined, we disallow reflection at vertices of P, which is a standard assumption in the literature. In Fig. 1, y is 2-visible from S and z is 4-visible from S.

For a point $S \in P$, let $\mathcal{V}_0(S)$ denote the polygonal region consisting of points in P that are directly visible from S, and, for $k \geq 1$, let $\mathcal{V}_k(S)$ denote the polygonal region

[1] This is called "clear visibility" in [26].

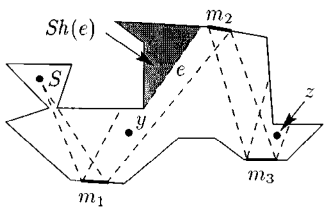

Fig. 1. Visibility under reflection.

consisting of points that are ℓ-visible from S, for some $\ell \le k+1$. Informally, $\mathcal{V}_k(S)$ is the set of points that receive light from S after at most k reflections off the boundary of P. For $\ell > 1$, let D be an ℓ-visible point. By definition, there exist points $p_1, p_2, \ldots, p_{\ell-1}$ lying in the interiors of edges $e_1, e_2, \ldots, e_{\ell-1}$ of P, respectively, on the ℓ-link path from S to D. The maximal $(\ell-1)$-visible portion of $e_{\ell-1}$ consisting of points lit up by rays reflected off edges $e_1, \ldots, e_{\ell-2}$, in this order, constitutes a *mirror* m at the $(\ell-1)$st stage of reflection. All points of $\mathcal{V}_k(S)$ that are ℓ-visible to S through m, via the same sequence of reflections, constitute the *mirror visibility polygon*, $V(m)$, of that mirror. $V(m)$ is a relatively open subset of P (except possibly for some reflex vertices of P) since we assume that vertices of P absorb any light incident on them. As the light rays are reflected off edges $e_1, \ldots, e_{\ell-1}$, a corresponding sequence of virtual images of the source is also created. At the first stage, the light rays reflected off e_1 emerge (when extended backward) from a virtual image S^1 of $S^0 = S$ with respect to the mirror on e_1; S^1 is a reflection of S^0 in the line containing e_1. To define the sequence of virtual images, denote the virtual image with respect to the mirror on e_{i-1} by S^{i-1}, for $1 < i \le \ell$. S^i, the next virtual image in the sequence, is the reflection of S^{i-1} through the mirror on e_i.

Let $\{m_1, m_2, \ldots, m_a\}$ be the set of all mirrors, up to and including stage k, *in any order*. Let $V_i = V(m_i)$. Let the corresponding set of virtual images be $\{S^1, S^2, \ldots, S^a\}$ where S^i is created with respect to the mirror m_i. Slightly abusing the notation, we let m_0 represent a "dummy mirror" so that $\mathcal{V}_0(S) = V_0 = V(m_0)$ is the set of points directly visible from S in our collection $\{V_i\}$.

Clearly, $\mathcal{V}_k(S) = \bigcup_{i=0}^{a} V_i$. Put $\Delta_i = \bigcup_{j=0}^{i} V_j$, for $0 \le i \le a$ so that $\mathcal{V}_k(S) = \Delta_a$. The connected components of the complementary region $P \backslash \Delta_i$ are called *blind spots* of Δ_i (see Fig. 2). These are the regions of P that do not receive light when we consider only mirror visibility polygons $V(m_0), V(m_1), \ldots, V(m_i)$. In particular, blind spots of Δ_a do not receive light after k stages of reflections, as a is the total number of (real) mirrors. The blind spots of Δ_i that are adjacent to $bd(P)$ are called *boundary blind spots* of Δ_i; the remaining blind spots are *interior*.

3. Upper Bound

It has already been shown that the complexity of $\mathcal{V}_1(S)$ is $O(n^2)$ [1]. Here, we aim to prove that the complexity of $\mathcal{V}_k(S) = \Delta_a$ is $O(n^{2k})$ for $k > 1$. To estimate the size of $\mathcal{V}_k(S)$ we start by showing that there can be at most $O(n^{2k})$ blind spots in $\mathcal{V}_k(S)$. First we state several crucial properties of the mirror visibility polygons.

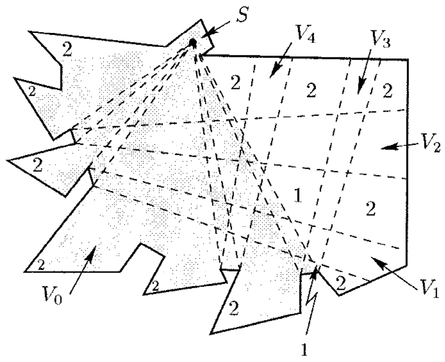

Fig. 2. Interior blind spots (1) and boundary blind spots (2) of $\Delta_4 = V_0 \cup V_1 \cup V_2 \cup V_3 \cup V_4$.

Lemma 3.1. *Each V_i is a simple polygon with no more than n edges.*

Proof. Consider the polygon $P \cup T$ where T is the triangle formed by S^i and m_i (see Fig. 3). Although, as subsets of the plane, T and P may overlap, we view $P \cup T$ as a Riemann surface, with the two polygons identified along m_i. It is easily observed that V_i is the direct visibility polygon of S^i in $P \cup T$ minus the triangle T. Since a direct visibility polygon in a polygon of size n cannot have more than n edges, the result follows. □

Lemma 3.2. *There are a total of $O(n^k)$ mirrors if k stages of reflection are allowed.*

Proof. The bound follows directly from the fact that a mirror at stage i can generate at most $n - 1$ mirrors for the next stage since each mirror visibility polygon is bounded by at most n edges (Lemma 3.1). □

Now we make some simple observations about how each V_i decomposes P. The edges of the relative boundary of V_i are called *shadow edges*. It is a straightforward, nevertheless crucial, observation that V_i has no vertex in $int(P)$. Let e be a shadow edge

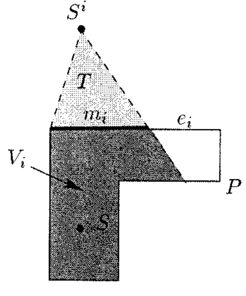

Fig. 3. Mirror m_i, its visibility polygon V_i, and virtual source S^i.

of V_i (see Fig. 1). It divides P into two subpolygons with disjoint interiors. The interior of one of these two subpolygons does not meet V_i. This portion of P, denoted $Sh(e)$, is called the *shadow* of e. The following lemma is an immediate consequence of this definition.

Lemma 3.3. *If e is a shadow edge of V_i, no other edge of V_i can lie in* $\mathrm{Sh}(e)$.

Lemma 3.4. *No segment contained in the interior of P can intersect the relative boundary of V_i more than twice.*

Proof. If a segment intersects at least three shadow edges of V_i, it has to intersect at least one edge of V_i lying in the shadow of another edge, violating Lemma 3.3. □

Lemmas 3.1 and 3.2 immediately imply an $O(n^{2k+2})$ bound on the complexity of $\mathcal{V}_k(S)$, as the desired set is the union of $O(n^k)$ polygons with $O(n)$ edges each. The claim follows from the observation that each vertex of $\mathcal{V}_k(S)$ is a vertex of the resulting arrangement of $O(n^{k+1})$ segments. The bound can be further strengthened to $O(n^{2k+1})$ by observing that the arrangement is "special" due to Lemma 3.4. However, this bound is still an order of magnitude larger than our target bound, which we proceed to establish.

Lemma 3.5. *The region bounded by a simple closed curve contained in P cannot have any point of $bd(P)$ in its interior.*

Proof. If a point of $bd(P)$ lay inside such a region in P, then $bd(P)$ would intersect the interior of P, an impossibility. □

Lemma 3.6. *If two mirror visibility polygons V_i, V_j intersect in such a way that each of two shadow edges e_1, e_2 of V_i intersects each of two shadow edges f_1, f_2 of V_j, then there are no other intersections between the relative boundaries of V_i and V_j.*

Proof. First, observe that by Lemma 3.4 neither e_1 nor e_2 meet the boundary of V_j again, and a symmetric statement holds for f_1 and f_2. Thus, by Lemma 3.3, the intersection of V_i and V_j lies completely in the quadrilateral Q delimited by e_1, e_2, f_1, and f_2 (see Fig. 4). Therefore, if the boundaries of the two visibility polygons ever meet at points

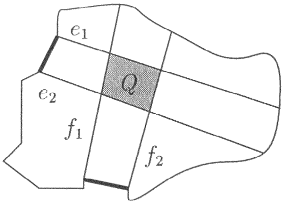

Fig. 4. A pair of edges of one mirror visibility polygon intersects a pair of edges of another mirror visibility polygon.

other than the four corners of Q, they have to do so inside Q. Then there is a shadow edge of V_i entirely in Q. However, such edges connect two points on the boundary of P. Hence there are two points on the boundary of P in Q. This contradicts Lemma 3.5. □

The combinatorial complexity of $\mathcal{V}_k(S)$ is determined by the complexity of its blind spots. Therefore, we concentrate on counting first the number of blind spots present in $\mathcal{V}_k(S)$ and then their total combinatorial complexity. We prove that there are only $O(n^{2k})$ blind spots in $\mathcal{V}_k(S)$. To show this we add the mirror visibility polygons $V_0, V_1, \ldots, V_a$, one by one, and count the increase in the number of blind spots. Recall that, by definition, $\Delta_i = \Delta_{i-1} \cup V_i$ and $\mathcal{V}_k(S) = \Delta_a$. First we observe some important properties of blind spots. Each interior blind spot is convex and each boundary blind spot is bounded by a connected portion of $bd(P)$ and a convex chain formed by portions of shadow edges. The proof of this fact for the case of at most one reflection is given in [1]. It applies here as it only uses the fact that V_i's are polygonal and have no vertices in the interior of P. We will need another property of blind spots:

Lemma 3.7. *The intersection of a blind spot h and any segment $s \subset P$ connecting two points on $bd(P)$, if nonempty, is a connected subsegment of s.*

Proof. The only points at which s can enter or leave h are points where s intersects a shadow edge e of some mirror visibility polygon V_j. Since h must be contained in $Sh(e)$, once s leaves h, it can never re-enter it, as it cannot re-enter $Sh(e)$. □

The next property follows from Lemma 3.3 and the fact that all lines containing the shadow edges of V_i must pass through the single image point S^i. We say that a segment s *cuts across* a blind spot if s crosses the relative boundary of the blind spot twice.

Lemma 3.8. *Any fixed blind spot in Δ_{i-1} can be cut across by at most two shadow edges of V_i.*

Proof. Suppose to the contrary that three shadow edges e_1, e_2, e_3 of V_i do cut across blind spot h.

Lemma 3.7 implies that each e_j, for $j = 1, 2, 3$, intersects h in an interval between its two points of intersection with the relative boundary of h.

Let V_i be a mirror visibility polygon associated with mirror m_i and source S^i, for some $i > 0$; the case that it is the direct visibility polygon V_0 of S is considered below. Then the edges e_1, e_2, e_3 can be ordered according to the order in which the rays emanating from virtual source S^i and containing the edges cross m_i; see Fig. 5(a,b). Without loss of generality, suppose the order is e_1, e_2, e_3. Consider the ray r_2 emanating from S^i and containing e_2. Notice that e_2 cannot emanate from m_i (see Fig. 5(a)), as otherwise it would cut P into two parts, each containing one of e_1, e_3, thereby contradicting Lemma 3.3. Thus the situation is as in Fig. 5(b).

We have shown that e_2 cannot emanate from the mirror m_i. Let s_2 be the portion of r_2 between m_i and e_2. By construction, s_2 cuts P into two parts P_1 and P_3, containing e_1 and e_3, respectively. Recall that we have assumed that e_2 crosses the relative boundary

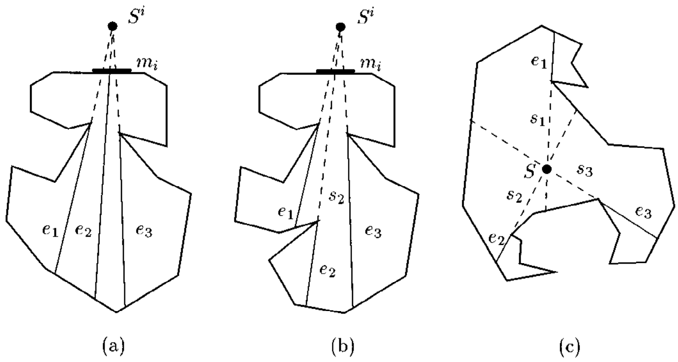

Fig. 5. Three shadow edges of V_i cannot cut across a blind spot.

of h twice. Applying Lemma 3.7 to $s_2 \cup e_2$, we see that $(s_2 \cup e_2) \cap h = e_2 \cap h$, so s_2 never meets h. However, h must lie on both sides of s_2, as it is crossed by both e_1 and e_3. Contradiction. Therefore, if e_1, e_2, e_3 bound V_i, for some $i > 0$, they cannot cut across a common blind spot h.

Finally, suppose that V_i is the direct visibility polygon of S. (For the sake of simplicity in the following argument, we assume that S is not collinear with any two vertices of P.) Let ℓ_j, $j = 1, 2, 3$, be the line containing e_j. Let s_j, for each $j = 1, 2, 3$, be the minimal line segment of ℓ_j connecting two boundary points of P and containing S. See Fig. 5(c). By the above assumption one of the endpoints of s_j is the reflex vertex of P from which e_j emanates. (Note that $S \notin e_j$.) Each s_j cuts the polygon in two. It is easily seen that (possibly after permuting the indices) the following must hold: s_2 cuts P into two parts so that each part contains one of e_1, e_3. On the other hand, arguing as above, $(s_2 \cup e_2) \cap h = e_2 \cap h$. Thus s_2 does not meet h, so there is no point of h on one side of s_2, contradicting the choice of e_1, e_2, e_3. □

Blind spots of Δ_i are obtained by removing points of V_i from blind spots of Δ_{i-1}. Thus we have:

Corollary 3.9. *Any blind spot of Δ_{i-1} can generate at most two new* interior *blind spots in Δ_i as a result of overlapping with V_i. In other words, a blind spot of Δ_{i-1} contains at most two interior blind spots of Δ_i.*

Lemma 3.10. *Let h_1, h_2 be two blind spots (interior or boundary) of Δ_{i-1} that are simultaneously intersected by a pair of shadow edges e and e' of V_i and such that no other blind spot intersects both e and e' between h_1 and h_2. Such an event can be associated with a pair (m_i, m_j) of mirrors, for some $j < i$, which is "charged" only once throughout the incremental construction of $\mathcal{V}_k(S) = \Delta_a$ from Δ_0.*

Proof. Let e exit h_1 through the edge g. Consider the region Q bounded by the two portions of relative boundaries of h_1, h_2 and the two portions of e and e' as in Fig. 6. By

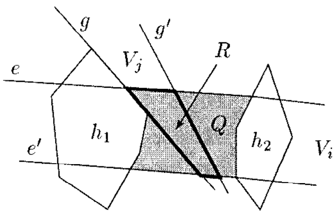

Fig. 6. A pair of shadow edges intersecting two blind spots.

Lemma 3.5, the interior of Q does not contain any point of $bd(P)$. Let $j < i$ be such that $g \subset bd(V_j)$. Since e enters another blind spot after entering V_j through g, it must exit V_j through another shadow edge, say g'. We claim that g and g' must intersect e'. Since Q cannot contain a boundary point of P (Lemma 3.5), g and g' must exit Q through an edge other than e. These exit points cannot lie on the relative blind spot boundaries since the blind spots are contained in $P \backslash V_j$. Hence g, g' of V_j must intersect e and e' of V_i, and by Lemma 3.6 no other edge of V_j can intersect V_i. Thus V_i and V_j intersect in a quadrilateral R bounded by (portions of) two shadow edges e, e' of V_i and two shadow edges g, g' of V_j.

The event in question can be assigned to the distinguished pair of mirrors (m_i, m_j), with $j < i$. This pair of mirrors cannot be charged for another pair of blind spots of Δ_{i-1}. Indeed, if a different pair, say, (h'_1, h'_2) charge (m_i, m_j), we must have $R' = V_i \cap V_j$ between h'_1 and h'_2, in the above sense. However, that is impossible since $R' = R$ by Lemma 3.6 and only h_1, h_2 have R between them. □

We now count the number of interior blind spots that can be present in Δ_i. The number of boundary blind spots is determined separately. To count the interior blind spots in Δ_i, we enumerate them as they are generated throughout the incremental construction of Δ_i starting from Δ_0. For this we concentrate on a generic incremental step of constructing Δ_j from Δ_{j-1} by overlaying V_j on it, for $1 \leq j \leq i$.

Let H^j denote the set of blind spots of Δ_{j-1}. We enumerate blind spots in H^j in successive steps and count the contributions of each group of blind spots to the increase of the number of *interior* blind spots during construction of $\Delta_j = \Delta_{j-1} \cup V_j$.

Step (i). First consider all blind spots from H^j that do not contribute to the increase in the number of interior blind spots as a result of the intersection with V_j. These include (a) the blind spots that are not intersected by $bd(V_j)$, since they either light up completely, or remain completely dark, (b) the interior blind spots that are intersected by only one edge of $bd(V_j)$, since each of them yields a single (though smaller) blind spot, and (c) the boundary blind spots that do not generate any interior blind spot as a result of intersection with $bd(V_j)$—they may be split into two or more boundary blind spots, or they may simply get smaller. Let H_1^j denote the set of *remaining* blind spots.

Step (ii). Consider those boundary blind spots in H_1^j whose convex chains are intersected by only one edge of $bd(V_j)$. Each such edge intersects the convex chain of a boundary blind spot twice since otherwise the blind spot would have been considered in step (i).

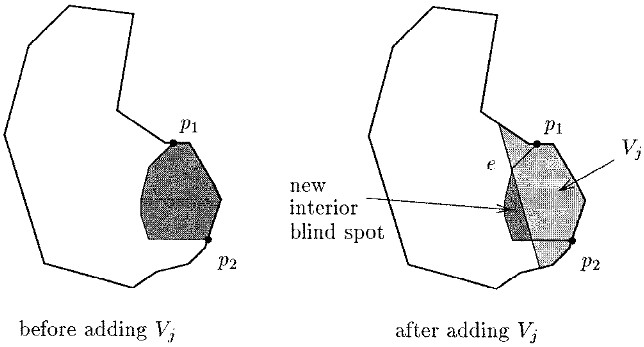

Fig. 7. The interior blind spot created from a boundary blind spot is charged to the endpoints of its convex chain.

Since only one new interior blind spot is created from an existing boundary blind spot considered in step (ii), it is enough to count the total number of boundary blind spots considered in step (ii) for determining the number of new interior blind spots created from them.

Lemma 3.11. *At most ni boundary blind spots are considered in step* (ii) *over all V_j, $1 \le j \le i$.*

Proof. A boundary blind spot considered in step (ii) must have both endpoints of its convex chain covered by the interior of V_j, $1 \le j \le i$, for otherwise it would not yield an interior blind spot. This means that these endpoints do not appear in any other boundary blind spot later. We charge the contributions of these boundary blind spots to these endpoints. Figure 7 illustrates the two endpoints p_1, p_2 of a convex chain. Since there are no more than n endpoints of shadow edges bounding every V_j, for each $j \le i$, we have ni charges in total. □

Now we proceed to count the other interior blind spots. Let H_2^j be the set of blind spots in H_1^j that are not considered in step (ii). The increase in the number of interior blind spots due to generation of new blind spots from H_2^j is bounded by twice the size of H_2^j, because each blind spot in H_2^j, whether interior or boundary, is replaced by at most two new interior blind spots; see Corollary 3.9.

Lemma 3.12. $\sum_{j=1}^{i} |H_2^j|$ *is at most* $\binom{i+1}{2} + 3ni$.

Proof. First, we prove that H_2^j has at most $j + 3n$ blind spots.

Due to eliminations in steps (i) and (ii), the relative boundary of every blind spot in H_2^j is intersected by exactly two edges of V_j. Consider the following planar graph G: Its nodes are the shadow edges of V_j that meet the relative boundary of at least one blind spot; recall that there are fewer than n such edges. The arcs of G correspond to blind

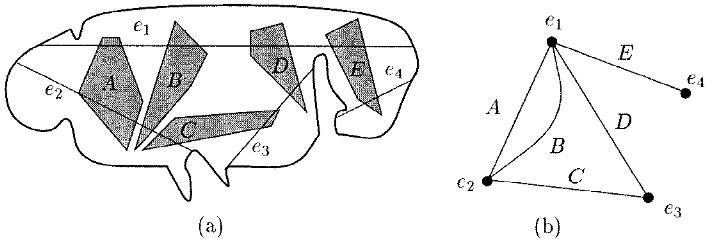

Fig. 8. (a) Blind spots remaining in H_2^j. (b) Planar graph G.

spots of H_2^j—two nodes are adjacent if the corresponding shadow edges meet a common blind spot. We allow multiple arcs between two nodes, corresponding to multiple blind spots meeting the same pair of edges of V_j. See Fig. 8. G is clearly planar, as blind spots are disjoint and each meets exactly two edges of V_j. By Euler's formula, the number of arcs in G is proportional to the number of its nodes (fewer than n) plus the number of what we refer to as 2-*sided faces*, which are faces in the embedding of G bounded by two nodes (shadow edges) and two arcs connecting them (two consecutive blind spots). For example, in Fig. 8(b) we have a 2-sided face between e_1 and e_2, bounded by A and B. The 2-sided faces correspond exactly to events counted by Lemma 3.10; each event is associated with a pair of mirrors $(m_j, m_{j'})$, with $j' \leq j$. Thus the contribution of this quantity is at most j since m_j can be paired with at most j mirrors processed before it. By Euler's formula, the number of the arcs not incident to any 2-sided face is at most $3n$, proving the claim.

Summing over all H_2^j, $j = 1, \ldots, i$, we obtain the desired bound. □

Lemma 3.13. *Δ_i has at most $\binom{i+1}{2} + 5ni$ blind spots.*

Proof. All interior blind spots in Δ_i have been created from an existing blind spot. Lemma 3.11 provides the total number of interior blind spots created from boundary blind spots considered in step (ii), while Lemma 3.12 provides the total number of remaining interior blind spots. Combining these two counts there are $\binom{i+1}{2} + 4ni$ interior blind spots in Δ_i. Finally, since there are at most ni endpoints of shadow edges in i mirror visibility polygons, there are at most ni boundary blind spots in Δ_i. Summing the two estimates, we obtain the desired $\binom{i+1}{2} + 5ni$ bound. □

Lemma 3.14. *Δ_i has at most $O(i^2 + ni \log i)$ edges.*

Proof. Divide the set of mirror visibility polygons into two subsets $M_1 = \{V_0, V_2, \ldots, V_{\lfloor i/2 \rfloor}\}$ and $M_2 = \{V_{\lfloor i/2 \rfloor + 1}, V_{\lfloor i/2 \rfloor + 2}, \ldots, V_i\}$. Put $R_1 = \bigcup M_1$ and $R_2 = \bigcup M_2$. Notice that all arguments used in Lemma 3.13 can be applied to *any* ordering of the mirrors. Thus, the bound in Lemma 3.13 can be used for both collections of mirror visibility polygons, R_1 and R_2. Accordingly, they each have $O(i^2 + ni)$ blind spots. We can think of the blind spots of R_1 (resp. R_2) as a collection of faces in the arrangement formed by the boundaries of the polygons in M_1 (resp. M_2) and the boundary of P. The blind spots of Δ_i are a subset of faces in the merged arrangement. We mark each blind spot of Δ_i with a point. To bound the complexity of these marked faces in Δ_i we determine the combined complexity of the marked blind spots (such a blind spot may contain more than one marker) individually in R_1 and R_2 and then consider the effect of merging the two. Note that each blind spot of Δ_i is a marked face in the resulting overlaid arrangement. Let $c(m, \ell)$ denote the complexity of m marked cells in an arrangement $A(L)$ of ℓ line segments. Let $L = L_1 \cup L_2$ where L_1 and L_2 have ℓ_1 and ℓ_2 segments, respectively. The combination lemma of [9] expresses the complexity of the marked faces in $A(L)$ in terms of the complexities $c(m, \ell_1)$, $c(m, \ell_2)$ of the marked faces in $A(L_1)$, $A(L_2)$, respectively, and the effect of merging the two, as follows:

$$c(m, \ell) = c(m, \ell_1) + c(m, \ell_2) + O(m + \ell)\,.$$

For Δ_i we have $m = O(i^2 + ni)$ (Lemma 3.12) and $\ell = O(ni)$. Denoting the worst-case complexity of blind spots in Δ_i as $f(i)$, over all possible orderings of the mirrors, we obtain the recurrence

$$f(i) = \begin{cases} 2f(i/2) + O(i^2 + ni), & \text{for} \quad i > 1, \\ O(n), & \text{for} \quad i = 1. \end{cases}$$

This recurrence solves to $O(i^2 + ni \log i)$. □

Theorem 3.15. *$\mathcal{V}_k(S)$ has combinatorial complexity $O(n^{2k})$, for any $k \geq 1$.*

Proof. Recall that $\mathcal{V}_k(S) = \Delta_a$. Using Lemma 3.14, $\mathcal{V}_k(S)$ has $O(a^2 + na \log a)$ edges. Since there are $O(n^k)$ mirrors involved in constructing $\mathcal{V}_k(S)$ we have $a = O(n^k)$. Plugging in this value of a we obtain an $O(n^{2k} + n^{k+1} \log n)$ bound for the complexity of $\mathcal{V}_k(S)$. For $k > 1$, the first term dominates the second and thus $\mathcal{V}_k(S)$ has $O(n^{2k})$ edges. By a different argument we proved in [1] that $\mathcal{V}_1(S)$ has size $O(n^2)$. Combining these two results we obtain the desired bound for all $k \geq 1$. □

Note that we had to use two different proof techniques, one for $k = 1$ and another for $k > 1$. An alternate approach to estimating the complexity of $\mathcal{V}_k(S)$ that applies uniformly to all $k \geq 1$ and also handles more general light source shapes is described in [7].

4. Lower Bound

In this section we describe the construction of a simple n-gon P with a point light source S so that the region $\mathcal{V}_k(S)$ lit up with at most k reflections has combinatorial complexity $\Omega((n/k - \Theta(1))^{2k})$. The construction can be carried out for any $k < n/c$, where $c > 1$ is an absolute constant. This lower bound asymptotically matches the upper bound of Theorem 3.15, if k is considered fixed.

We use a series of k "gadgets" that we call *convex mirrors* (CMs) each consisting of $N = \lfloor n/k \rfloor - \Theta(1)$ segments that we call *facets*. In this section the term *beam* refers to

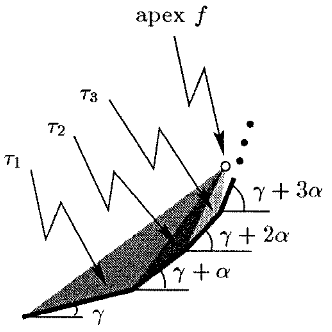

Fig. 9. CM(α, N).

the collection of all rays emanating from S in a contiguous interval of directions and then reflecting off the same sequence of polygon edges, before arriving at the portion of the polygon under consideration. We arrange for a single beam of light to emerge from the source S and fall on the first CM, which splits it into N beams. These beams converge on the second CM, which splits them into N^2 beams. Repeating this process k times, we obtain a set of N^k very thin beams, most pairs of which intersect, obtaining a pattern of complexity $\Omega(N^{2k})$. Details are provided below.

Let $N \geq 1$ be an integer and let $\alpha < \pi/(2N)$ be a positive number. We define an *N-faceted convex α-mirror*, CM(α, N), as a convex chain of N segments (*facets*), with turn angle α at each vertex, and such that there exists a point f (we call it an *apex* of the mirror) from which each facet subtends an angle of measure 2α. When the precise values of α and N are unimportant or understood from the context, we refer to this object as a *convex mirror*, or simply CM.

Observation 4.1. *For sufficiently small values of* α, CM(α, N) *exists.*

Proof. Fix a parameter γ, $0 < \gamma < \pi/2$. Refer to Fig. 9. Without loss of generality, assume that the first facet of CM emanates from the origin at the angle of $+\gamma$ to the positive x-axis and every subsequent edge turns counterclockwise from the line containing the previous one by α. Let m_i be the ith facet of the mirror. Using m_i as the base, construct the triangle τ_i above it with interior angles $\gamma + (i-1)\alpha$, $\pi - \gamma - (i+1)\alpha$, and 2α, respectively, at the left, right, and top vertex. Then the left edge of τ_i overlaps the right edge of τ_{i-1}. Hence, an appropriate choice of relative sizes of the facets guarantees that the top vertices of all triangles τ_i coincide. This is the desired apex f of the CM.

The construction can be carried out if $0 < \gamma < \pi/2$ and $\alpha < (\pi - \gamma)/(N+1)$. □

Note that in the above construction α, N, and γ determine the CM, up to scaling.

We define a set of beams to be *β-parallel in a disk* if each beam lights the entire disk, and the range of directions of incoming light rays, over all beams and all points of the disk, fits in an angular interval of measure β. We refer to the length of the smallest such interval as the (*angular*) *spread* of the beams covering the disk.

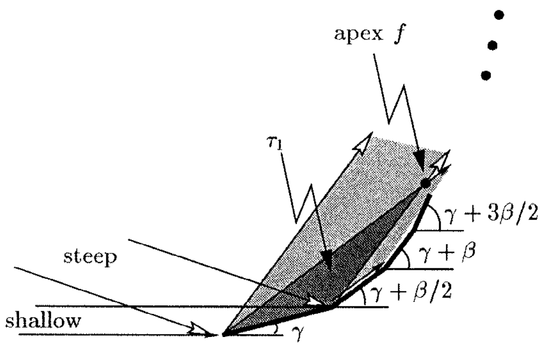

Fig. 10. CM and the light beams.

Lemma 4.2. *Let D be a disk lit up completely by a collection of s* incoming *β-parallel beams. Then there exists another disk D' with the following properties*: CM($\beta/2, N$) *can be constructed in D so that the incoming beams split into sN* outgoing *beams which are $(N+1)\beta$-parallel in D' and the ratio of the distance between the centers of D and D' to the radius of the larger of the two disks can be made arbitrarily large as $\beta \to 0$.*

Proof. Without loss of generality, suppose that D is centered at the origin and that the directions from any point in D to the source of any incoming beam are all in the range from $\pi - \beta$ to π. Let γ, $0 < \gamma < \pi/2$, be a parameter. We orient CM $=$ CM($\beta/2, N$) as before, aiming the first facet m_1 at angle $+\gamma$, second m_2 at angle $+\gamma + \beta/2$, and so forth. In general, facet m_i, for $i = 2, \ldots, N$, emanates from the rightmost endpoint of facet m_{i-1} at the angle of $\beta/2$ to the line containing m_{i-1} and thus at the angle of $\gamma + (i-1)\beta/2$ to the positive x-axis. Refer to Fig. 10. The entire construction is scaled so that the rightmost point of the CM lies on the boundary of D.

Recall that each incoming beam emanates from a point (its real or virtual source). Consider two parallel beams of light emanating from directions $\pi - \beta$ ("steep") and π ("shallow"), respectively; a *parallel* beam here is a beam the source of which lies at infinity. The intersection of reflections of the two beams off facet m_i is precisely the triangle τ_i. The reflection of any parallel beam emanating from a direction in the interval $[\pi - \beta, \pi]$ in m_i covers τ_i. Any beam emanating from a point source at a finite distance in some direction in this range has the property that its reflection off m_i covers τ_i, so that f lies in the interior of the reflection. In particular, reflections off every incoming beam in every facet of the CM form sN beams, every one of which covers a sufficiently small neighborhood of the apex f of the CM. We place a disk D' centered at f in that neighborhood. An easy calculation shows that the angular spread of the sN reflected beams at D' is at most $(N+1)\beta$.

Let x_0 be the length of the left edge of τ_1 and let x_N be the length of the right edge of τ_N. Straightforward calculation shows that

$$\frac{x_0}{x_N} = \frac{\sin(\gamma + (N+1)\beta/2)\sin(\gamma + N\beta/2)}{\sin(\gamma + \beta/2)\sin\gamma},$$

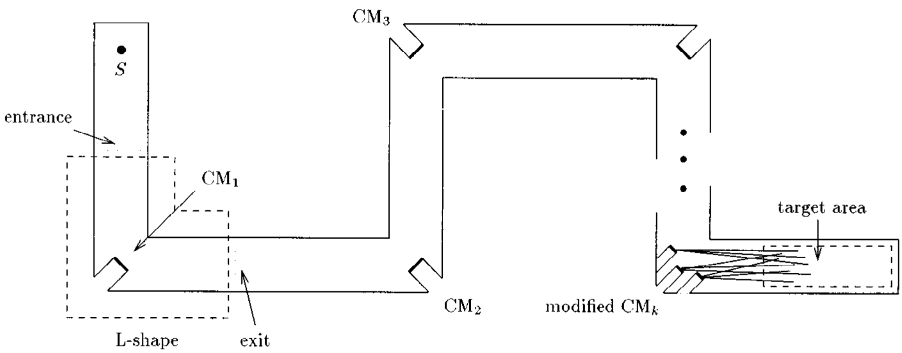

Fig. 11. Placement of mirrors in the polygon.

which is larger than one, for fixed γ and small enough β, and approaches one as $\beta \to 0$, since γ is an acute angle.

As the CM is scaled to just fit in D, the radius of D is the distance δ from the beginning of m_1 to the end of m_N. As $x_0 \geq x_N$, for small β, $\delta \leq 2x_0 \sin(N\beta/2)$, since $N\beta$ is the angle subtended from f by the entire CM. On the other hand, D' is an arbitrarily small disk centered at f. Hence the ratio of the distance x_0 between the centers of D and D' to the larger radius is $x_0 : \delta = 1 : (2\sin(N\beta/2)) \to \infty$ as $\beta \to 0$. This completes the proof of the lemma. □

Putting CMs together, we finally obtain the lower bound construction.

Theorem 4.3. *There exists a simple n-gon with a source point S and k CMs, such that the combinatorial complexity of the resulting visibility region,* $\mathcal{V}_k(S)$, *is* $\Omega((n/k - \Theta(1))^{2k})$.

Proof. We start with a snake-like polygon; the width of every leg of the "corridor" is much smaller than its length; refer to Fig. 11, not drawn to scale. Pick a disk D_1 visible to S, and place the first convex mirror, $\mathrm{CM}_1 = \mathrm{CM}(\delta, N)$, inside it. There is only one beam coming in to CM_1 and N going out. The angular spread of the single incoming beam can be made less than δ, for any $\delta > 0$ of our choice, by shrinking D_1. The corridor is constructed to be a little wider than D_1. The outgoing beams will overlap in a disk D_2 which we construct by Lemma 4.2. The spread among the beams in that disk will be at most $(N+1)\delta$. Place $\mathrm{CM}_2 = \mathrm{CM}((N+1)\delta, N)$ in D_2. This produces N^2 beams overlapping in disk D_3 with angular spread bounded by $(N+1)^2\delta$. Repeat k times, getting N^k beams. The N^k beams overlap in a common region ("target area"); in fact they overlap at least in a common disk D_{k+1}, by construction, and have angular spread $(N+1)^k\delta$ in it. The beams make a near-right-angle turn at each CM, which corresponds to setting the γ parameter of CM to a value close to $\pi/4$. Putting $\delta = \pi/(100(N+1)^k)$ ensures that angular spread in every CM is smaller than $\pi/100$. This is sufficient to make all the CM constructions work as described since it satisfies the constraint $\beta N \ll \pi$ assumed in the definition of a CM.

We now modify CM_k as follows. Pick a generic point on each of its N facets and consider each point to be a submirror of the facet, with infinitesimal length. This produces N fans of N^{k-1} reflected ray-like beams each. At least about $\frac{1}{2}\binom{N^{k-1}}{2}\binom{N}{2} = \Theta(N^{2k})$ pairs of these rays intersect in the target area.

Indeed, if one numbers the rays in each fan (see rays coming out of CM_k in Fig. 11), from 1 to N^{k-1} (clockwise in the figure), and numbers fans from 1 to N (top left to bottom right in the figure), then ray i in fan j must intersect ray i' in fan $j' > j$ if $i' \geq i$. The case $i' = i$ corresponds to rays originating from the same incoming beam—these rays must meet as the mirror is "convex" and distinct facets are rotated by at least the angle to compensate for the angular spread of the incoming rays, which is $(N+1)^{k-1}\delta$ for CM_k. Consider the case $i' > i$. By the above reasoning ray i in fan j' meets ray i in fan j. However, ray i' in fan j' lies clockwise of ray i and hence meets ray i in fan j even earlier. Hence at least approximately half of the ray pairs eventually intersect.

By the generic choice of points for infinitesimal mirrors, no three rays emanating from different mirrors have a point in common. Thus the above $\Theta(N^{2k})$ estimate on the number of pairs of intersecting rays also estimates the number of vertices of the resulting arrangement. Therefore, each infinitesimal mirror can be expanded to a sufficiently short, but positive-length, mirror without reducing the complexity of the union of resulting beams. This produces a family of $\Theta(N^k)$ beams of light each of which has encountered k reflections and whose union has complexity $\Theta(N^{2k})$. It remains to check that no light from S other than that reflecting off $CM_1, \ldots, CM_k$, in this order, is allowed to arrive at the target area with k or fewer reflections.

The link distance between S and the target area is $k+1$, so no ray can reach the target area with fewer than k reflections. No "unauthorized" ray can reach the target area with exactly k reflections since, in order to do so, such a ray would have to make at most one turn inside each L-shaped region indicated in the figure. The L-shape is not drawn to scale—its "legs" are much longer than they are wide. However, the only points directly visible from both the "entrance" and the "exit" of an L-shape are those of the CM contained in it, points in the interior of P near the turn, points of $bd(P)$ on either side of the CM, and the reflex corner of the L-shape. Among those, however, none but the points of the CM can be used to reflect a ray of light from the entrance, so that it arrives at the exit. This proves the claim. □

5. Algorithm for Computing $\mathcal{V}_k(S)$

We first compute $V_0(S)$ and all mirror visibility polygons V_i successively as follows. For computing the mirror visibility polygon V_i we first determine the image S^i of S with respect to the mirror m_i. Then considering the direct visibility polygon of S^i in the polygon $P \cup T$, T being the triangle formed by S^i and m_i, we can obtain V_i as described in Lemma 3.1. Direct visibility polygons can be computed by any one of the known linear-time algorithms [10], [13]. (More precisely, it is easy to check that a triangulation-based algorithm will work correctly on the Riemann surface $P \cup T$, even if, as subsets of the plane, P and T overlap.) After computing all V_i's we apply a divide-and-conquer technique to compute the final visibility polygon $\mathcal{V}_k(S)$. Let $V' = V_0 \cup \cdots \cup V_{\lfloor \ell/2 \rfloor}$ and $V'' = V_{\lfloor \ell/2 \rfloor+1} \cup \cdots \cup V_\ell$ be computed recursively. To compute $V = V' \cup V''$ we merge

V' and V'' employing the sweep-line algorithm of Bentley and Ottmann [2]. We sweep V' and V'' from left to right and construct V as the sweep proceeds. Each intersection between $bd(V')$ and $bd(V'')$ detected by the algorithm must contribute a vertex on the boundary of V. By Lemma 3.14 V', V'', V have only $O(\ell^2 + n\ell \log \ell)$ vertices and segments on their boundary. However, $\ell \le a = O(n^k)$ so $O(\log \ell) = O(\log n^k) = O(\log n)$, as k is fixed. Thus sweeping takes $O(\ell^2 \log n + n\ell \log^2 n)$ time. Note that using an optimal output-sensitive algorithm, such as [6], for merging does not improve the running time since V', V'', V have the same order of complexity in the worst-case. The time $T(\ell)$ for our divide-and-conquer algorithm can be expressed as

$$T(\ell) = \begin{cases} 2T(\ell/2) + O(\ell^2 \log n + n\ell \log^2 n), & \text{for} \quad \ell > 1, \\ O(n), & \text{for} \quad \ell = 1, \end{cases}$$

which solves to $T(\ell) = O(\ell^2 \log n + n\ell \log^3 n)$. Observing that the required space remains bounded by the sizes of the unions of mirror visibility polygons, we obtain the following result.

Theorem 5.1. *The visibility polygon $\mathcal{V}_k(S)$ can be computed in $O(n^{2k} \log n + n^{k+1} \log^3 n)$ time and $O(n^{2k})$ space, for any $k \ge 1$.*

Note that, with the sharper $O(n\ell)$ bound on the complexity of the union V of ℓ mirror visibility polygons for the case $k = 1$, previous analysis yields a stronger $O(n^2 \log^2 n)$ bound on the running time of the algorithm, when it is used to compute $\mathcal{V}_1(S)$. This is essentially the algorithm given in [1].

References

1. B. Aronov, A. R. Davis, T. K. Dey, S. P. Pal, and D. C. Prasad. Visibility with one reflection. *Discrete Comput. Geom.*, 19(4):553–574, 1998.
2. J. L. Bentley and T. A. Ottmann. Algorithms for reporting and counting geometric intersections. *IEEE Trans. Comput.*, 28:643–647, 1979.
3. M. de Berg. *Ray Shooting, Depth Orders and Hidden Surface Removal*. Lecture Notes in Computer Science, vol. 703. Springer-Verlag, Berlin, 1993.
4. C. Boldrighini, M. Keane, and F. Marchetti. Billiards in polygons. *Ann. Probab.*, 6:532–540, 1978.
5. M. Born and E. Wolf. *Principles of Optics*, 6th edn. Pergamon Press, Oxford, 1980.
6. B. Chazelle and H. Edelsbrunner. An optimal algorithm for intersecting line segments in the plane. *J. Assoc. Comput. Mach.*, 39:1–54, 1992.
7. A. R. Davis. Visibility with Reflections in Triangulated Surfaces. Ph.D. thesis, Dept. Computer and Information Science, Polytechnic University, Brooklyn, NY, 1998.
8. H. N. Djidjev, A. Lingas, and J. Sack. An $O(n \log n)$ algorithm for computing the link center of a simple polygon. *Discrete Comput. Geom.*, 8(2):131–152, 1992.
9. H. Edelsbrunner, L. J. Guibas, and M. Sharir. The complexity and construction of many faces in arrangements of lines and of segments. *Discrete Comput. Geom.*, 5:161–196, 1990.
10. H. ElGindy and D. Avis. A linear algorithm for computing the visibility polygon from a point. *J. Algorithms*, 2:186–197, 1981.
11. J. Foley, A. van Dam, S. Feiner, J. Hughes, and R. Phillips. *Introduction to Computer Graphics*. Addison-Wesley, Reading, MA, 1994.
12. R. Guenther. *Modern Optics*. Wiley, New York, 1990.
13. L. J. Guibas, J. Hershberger, D. Leven, M. Sharir, and R. E. Tarjan. Linear-time algorithms for visibility and shortest path problems inside triangulated simple polygons. *Algorithmica*, 2:209–233, 1987.

14. E. Gutkin. Billiards in polygons. *Phys. D*, 19:311–333, 1986.
15. A. Horn and F. Valentine. Some properties of *l*-sets in the plane. *Duke Math. J.*, 16:131–140, 1949.
16. Y. Ke. An efficient algorithm for link-distance problems. *Proc. 5th Annu. ACM Sympos. Comput. Geom.*, 1989, pp. 69–78.
17. Y. Ke. Polygon Visibility Algorithms for Weak Visibility and Link Distance Problems. Ph.D. thesis, Dept. Computer Science, Johns Hopkins University, Baltimore, MD, 1989.
18. S. Kerckhoff, H. Masur, and J. Smillie. Ergodicity of billiard flows and quadratic differentials. *Ann. of Math.*, 124:293–311, 1986.
19. V. Klee. Is every polygonal region illuminable from some point? *Amer. Math. Monthly*, 76:180, 1969.
20. V. Klee and S. Wagon. *Old and New Unsolved Problems in Plane Geometry and Number Theory.* Mathematical Association of America, Washington, DC, 1991.
21. V. V. Kozlov and D. V. Treshchev. *Billiards: A Genetic Introduction to the Dynamics of Systems with Impacts*. Translations of Mathematical Monographs, vol. 89. American Mathematical Society, Providence, RI, 1991.
22. W. Lenhart, R. Pollack, J.-R. Sack, R. Seidel, M. Sharir, S. Suri, G. T. Toussaint, S. Whitesides, and C. K. Yap. Computing the link center of a simple polygon. *Discrete Comput. Geom.*, 3:281-293, 1988.
23. J. Mitchell. Shortest paths and networks. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, Chapter 24, pp. 445–466. CRC Press, Boca Raton, FL, 1997.
24. I. Newton. *Opticks, or A Treatise of the Reflections, Refractions, Inflections and Colours of Light*, 4th edn. London, 1730.
25. J. O'Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, New York, 1987.
26. J. O'Rourke. *Computational Geometry in C*. Cambridge University Press, New York, 1994.
27. J. O'Rourke. Visibility. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, pp. 467–479. CRC Press, Boca Raton, FL, 1997.
28. M. Pellegrini. Ray shooting and lines in space. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, pp. 599–614. CRC Press, Boca Raton, FL, 1997.
29. D. C. Prasad, S. P. Pal, and T. K. Dey. Visibility with multiple diffuse reflections. *Comput. Geom. Theory Appl.*, to appear, 1998.
30. J. Reif, J. Tygar, and A. Yoshida. Computability and complexity of ray tracing. *Discrete Comput. Geom.*, 11:265-288, 1994.
31. T. C. Shermer. Recent results in art galleries. *Proc. IEEE*, 80(9):1384–1399, 1992.
32. S. Suri. On some link distance problems in a simple polygon. *IEEE Trans. Robot. Automat.*, 6:108–113, 1990.
33. G. W. Tokarsky. Polygonal rooms not illuminable from every point. *Amer. Math. Monthly*, 102:867–879, 1995.

Received March 14, 1996, *and in revised form December* 22, 1997, *and January* 5, 1998.