The union of billiard-ball paths: Structure and complexity

Sudebkumar Prasant Pal [1]
Department of Computer Science and Engineering
Indian Institute of Technology, Kharagpur, 721302, India
email: spp@cse.iitkgp.ernet.in

`url: http://www.facweb.iitkgp.ernet.in/~spp/`

Abstract

Consider the piece-wise linear path followed by a billiard-ball starting at an internal point s in a simple n-sided polygonal enclosure P with reflecting walls. The union of all possible billiard-ball paths with $k+1$ links and suffering at most $k>0$ reflections, is identical to the region $V_k(s)$ of the polygon P visible from the point light source at s due to at most k reflections. The nature of reflection, *diffuse* or *specular*, and, the number of reflections k, determine the properties of the union of the paths or (equivalently) of the visible region $V_k(s)$. In diffuse reflection, an incident ray is reflected back in any possible direction. In specular reflection, the angle of reflection is equal to the angle of incidence. Multiple reflections make larger portions of P visible, creating visible regions with more complex structure and combinatorial complexity.

We survey recent results on the combinatorial complexity of $V_k(s)$ for specular and diffuse reflections. For specular reflection, $V_k(s)$ has at most $O(n^{2k})$ edges. In the case of diffuse reflection, the complexity drops down to $O(n^{2\lceil(k+1)/2\rceil+1})$. We also mention matching lower bounds for the specular case and a loose $\Omega(n^2)$ lower bound for the diffuse case.

It is easy to show that the region $V_0(s)$ of P directly visible from s, is simply connected and has at most n bounding edges. The region $V_1(s)$ of points visible from s after at most one specular reflection may be multiply connected. Being multiply connected, $V_1(s)$ may have numerous *islands* or *holes*– dark regions surrounded by illuminated regions. These are distinct from *shadows*– dark regions adjacent to the boundary of P. We mention the non-trivial construction of an n-sided polygon P where the region $V_k(s)$ has $\Omega(n/k)^{2k}$ islands, for fixed k.

The scenario is quite contrasting (as expected) when we change the nature of reflection from the above-mentioned specular to diffuse (or scattering) reflection. One expects fewer and smaller shadows and possibly no islands of darkness. We mention the result in [2] that $V_1(s)$ has no islands (it is simply connected). We state the $\Omega(n^2)$ lower bound on the number of edges of $V_1k(s)$. We also demonstrate a remarkable and rich construction where two successive diffuse reflections on a light path starting at s, permits islands of darkness amidst illumination in $V_2(s)$. Interestingly enough, there are plenty of islands or holes in $V_2(s)$– proportional to n in the worst case but possibly no more than a fraction of n.

Finally, we conclude with a discussion on a few open problems in combinatorial and computational geometry of visibility with reflection.

Keywords: visibility, diffuse reflection, specular reflection, union, connectedness, combinatorial complexity, computational complexity.

[1] A part of this work was done when the author was visiting the University of Miami.

1 Introduction

Billiard paths are well studied in the field of ergodic theory. Gutkin's paper [9] and the book of Klee and Wagon [11] deal with problems pertaining to billiard paths. Billiard paths (called billiard-ball paths in this paper) are similar to paths followed by light undergoing multiple reflections on bounding walls. A remarkable recent result of Tokarsky [18] has settled Klee's question in [10] asking whether a simple polygon bounded by mirrors can be completely lit up by a single light bulb placed at an arbitrary point in its interior. Regions illuminated by light sources due to multiple reflections can be visualized as unions of billiard-ball paths. In our study of the structure and complexity of these regions, we restrict the number of reflections to a constant $k > 0$. We survey results on the connectedness properties and combinatorial complexities of these unions. We also discuss algorithms and their computational complexities for computing such regions.

In recent years, visibility problems have been studied extensively (see [14, 17]). Visibility computations abound in computer graphics, motion planning, robotics and computer vision. Two points inside a simple polygon are mutually *visible* if the line segment joining them is not obstructed by any edge of the polygon. Several algorithms for computing the region visible from a point light source inside a simple polygon exist [7, 12, 8]. The computation of the region of a simple polygon that is weakly visible from (visible from some point on) an internal segment [13, 8], or a convex set [4], are also well studied. Weak visibility polygons succumb to linear time triangulation algorithms that use simple data structures [5]; they have efficient recognition algorithms too [6, 3].

Certain portions of the polygon that are not directly illuminated from the source may become visible due to one or more reflections on the bounding edges of the polygon. Reflection at a point is called *specular* if the reflected ray is directed as per the standard law of reflection; the angle of incidence is the same as the angle of reflection. Most reflecting surfaces cause another type of reflection, called *diffuse* reflection, where light incident at a point is reflected in all possible interior directions. An edge that reflects by diffuse reflection contributes to the total visible region; the union of such regions generated due to reflection from individual edges of the enclosing polygonal boundary is the final visible region due to reflections. Note here that the contribution of each reflecting edge e is the region of the enclosing polygon that is weakly visible from the illuminating or reflecting portion of the edge e. This shows the close connection between visibility with diffuse reflection and weak visibility. In this paper, we study properties of the region visible from a point light source inside a simple polygon due to multiple reflections on the edges of the polygon. We consider both diffuse as well as specular reflections. We assume that light incident at a vertex is absorbed and not reflected. We restrict our attention to the region visible due to at most k reflections within a simple polygon with n edges, where $k \geq 1$ is a constant.

The combinatorial complexity of visibility polygons due to reflection was first studied by Aronov *et al.* [2]. A tight $\Theta(n^2)$ worst-case bound was established for visibility polygons when at most one diffuse reflection is permitted. They also studied the combinatorial complexity of the region visible due to at most one specular reflection. A tight $\Theta(n^2)$ bound was also established for at most one specular reflection in [2]. In a later paper [1], they investigated the case where at most k specular reflections are permitted. They produced an $O(n^{2k})$ upper bound and a lower bound of $\Omega((n/k - \Theta(1))^{2k})$ on the combinatorial complexity of the visible region when at most a constant number of k specular reflections are permitted. The upper bound improves to $O(n^{2\lceil (k+1)/2 \rceil + 1})$, when specular reflections are replaced by diffuse reflections [16].

In this paper we present a brief survey of results on the structure, combinatorial complexity, connectivity and computational complexity of regions visible due to multiple reflections. In addition to a summary of the above-metioned results, we also delve into results pertaining to the connectedness of visible regions formed by multiple reflections. Consider a point light source s inside an n-sided simple polygon P. It is easy to see that the portion of P visible from s is also simply connected. In [2], Aronov et al. present the construction of an n-sided polygon P where the portion of P visible from s due to at most one specular reflection has $\Omega(n^2)$ *islands* or *holes*. A *hole* is a connected invisible region bounded by visible points. So, holes are distinct from *shadows*; shadows are invisible regions adjacent to the boundary of P. In the same paper, the authors also show that the region of any simple polygon P, visible from any internal point s due to at most one diffuse reflection cannot have a hole. In this paper we show that the region visible from a point due to at most two diffuse reflections may have a hole (see [15]). We establish this by constructing a simple polygon P and by placing a point s inside it so that the portion of the polygon visible from s due to two diffuse reflections does not include the interior of a triangle tqr (see Figures 5.1 and 5.2). Here, the triangle tqr is totally contained inside P and the boundary of triangle tqr as well as its exterior neighbourhood are visible from s after two diffuse reflections.

We also claim that there can be numerous holes. More precisely, we sketch an argument showing an $\Omega(n)$ lower bound on the number of holes in the region of an n-sided simple polygon visible due to at most two diffuse reflections. The number of holes can be much larger in visibility regions in the case of multiple specular reflections [1]. For one specular reflection, there can be $\Theta(n^2)$ holes [2].

The remainder of this paper is organized as follows. In section 2 we some notation and discuss some preliminaries. In section 3 we discuss bounds on the combinatorial and computational complexity of regions visible due to multiple diffuse reflections. In section 4 we briefly state results on the combinatorial and computational complexity of regions visible due to multiple specular reflections. In section 5 we present the construction of a simple polygon with an interior point source of light to show the existence of a hole in the the region visible from the source due to two diffuse reflections. In section 6 we show a lower bound of $\Omega(n)$ on the number of holes in the visibility polygon of an n-sided simple polygon due to two diffuse reflections. We conclude in section 7 with a discussion on open problems.

2 Preliminaries and notation

Let $P= (v_1, v_2, ..., v_n)$ be a simple polygon with no three collinear vertices. Let $int(R)$ and $bd(R)$ denote the relative interior and boundary of a region R, respectively. If p and q are two points on $bd(P)$, let $bd(p, q)$ denote the part of $bd(P)$ traversed from p to q in counterclockwise order.

Two points in P are said to be *(mutually) visible* if the interior of the line segment joining them lies entirely in $int(P)$. A point p is *weakly visible* from an edge or an internal line segment st, if there is a point z in the *interior* of st such that p and z are mutually visible. A point y is visible under *diffuse reflection* from a point x, if there exists a point p lying in the interior of an edge of P, such that p is visible from both x and y. A point y is visible under *specular reflection* from a point x, if

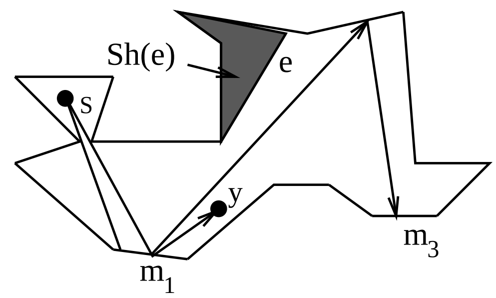

Figure 2.1: Definition of visibility with reflection

there exists a point p lying in the interior of an edge e of P, such that p is visible from both x and y and the angle made by segments xp and yp are equally inclined to the edge e.

Let s denote a point light source inside P. For $k > 2$, a point $y(= p_{k+1})$ is said to be visible *after k diffuse reflections* from a point $s(= p_0)$, if there exist points $p_1, p_2, ..., p_k$ in the interiors of edges of P, such that p_i and p_{i+1} are mutually visible for $0 \leq i \leq k$. We say that the ray emitted from s reaches y after k *diffuse reflections* at points $p_1, p_2, ..., p_k$. For visibility between s and $y = p_{k+1}$ *after k specular reflections* the additional condition is that $p_{i-1}p_i$ and p_ip_{i+1} must make the same inclination with the edge of P containing the point p_{i+1} of the ith reflection, $0 \leq i \leq k$. In Figure 2.1, y is visible from s after one diffuse reflection. In Figure 2.1, the shadow $Sh(e)$ is the dark region, comprising points not visible from s after one specular reflection. The shadow edge e is defined by the unique point z on the edge m_1 (z is not shown in the figure), such that sz and line containing edge e are equally inclined with m_1. We assume that the light incident at a vertex is absorbed and not reflected further.

Let $V(s)$ denote the portion of P (directly) visible from s. We know that $V(s)$ has $O(n)$ edges and at most one edge of $V(s)$ lies on each edge of P [7, 8]. For $k \geq 0$, let $V_k(s)$ denote the polygonal region consisting of points that are visible from s after k reflections. Informally, $V_k(S)$ is the set of points that receive light from s after at most k reflections. We have $V_0(s) = V(s)$.

3 Visibility by multiple diffuse reflections

We know that $V_1(s)$ has combinatorial complexity upper bounded by $O(n^2)$ [2]; in addition $V_1(s)$ has no islands or holes. In section 5, we show that $V_2(s)$ may have holes. In this section we summarize results on the complexity $V_k(s)$ in the case of multiple diffuse reflections as established in [16].

We view $V_1(s)$ as the union of $V(s)$ and the regions visible from edges of V(s) lying on $bd(P)$. These edges of $V(s)$ cause the first reflection. These edges are also called *mirrors at the 1st stage of reflection*. We can similarly and inductively define a *mirror at the kth stage of reflection* as follows. Let m be a mirror at the $(k-1)$st stage of reflection, and $V(m) \subseteq P$ be the set of points weakly visible from the mirror m. We know that $V(m)$ has at most one edge lying on each edge of P [8]. We call $V(m)$ the *visibility polygon* of m. An edge of $V(m)$ lying on $bd(P)$ is a *mirror at the kth stage* of reflection. In Figure 2.1, m_3 is a mirror at the third stage of reflection.

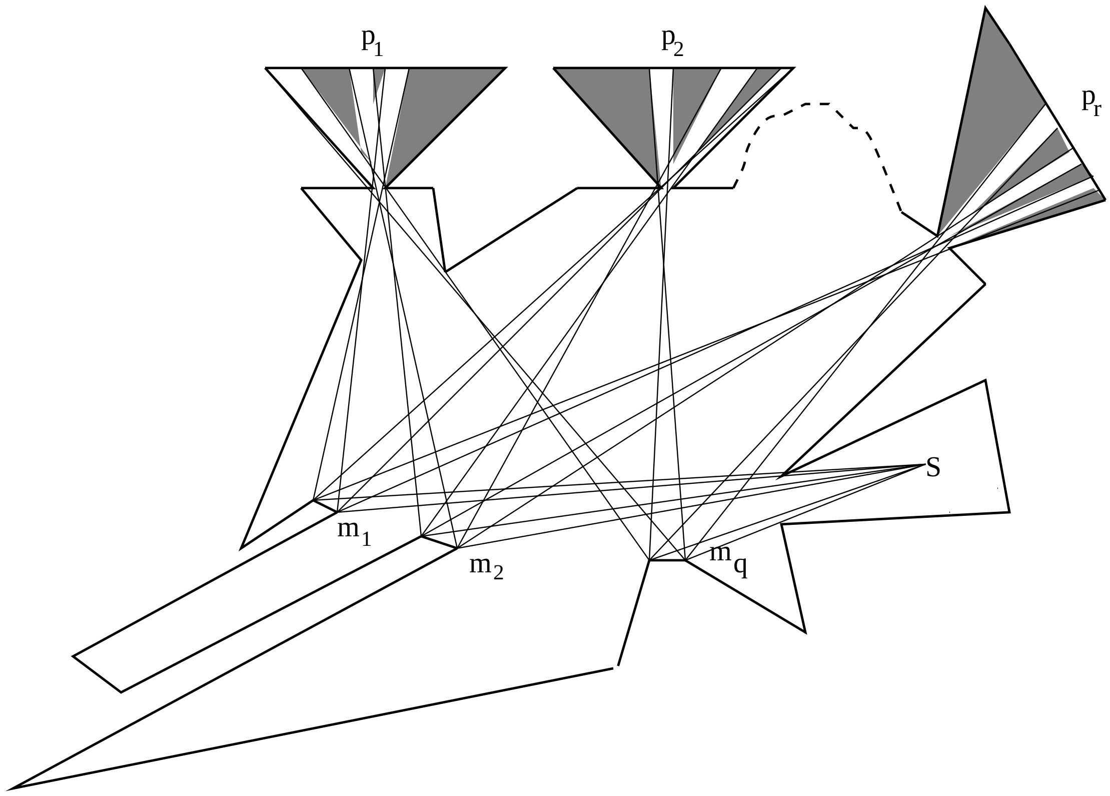

Figure 3.2: $V_1(S)$ has $\Omega(n^2)$ edges

It is worthwhile considering the union of mirrors lying on an edge e or P at the kth stage of reflection; we are interested naturally in the connected components of this union. Each point in a connected component is weakly visible from some mirror at the $(k-1)$st stage of reflection. We call such a connected component a *reflecting segment at the kth stage* of reflection [16]. The region $V(c)$, weakly visible from a reflecting segment, is called its *visibility polygon*. Recall that the mirrors at the first stage of reflection are created on the edges that are directly lit up by the source S. Since the number of edges of $V(S)$ is $O(n)$ (see [7, 8]), the number of mirrors that cause the first stage of reflection is n. Each of these n mirrors can create at most one reflecting segment on each edge of P [8]. Therefore, there are $O(n^2)$ mirrors involved in the second stage of reflection. Indeed, this worst case can be realized as illustrated in Figure 3.2, where $\Omega(n)$ mirrors are created on each of $\Omega(n)$ edges, thus creating $\Omega(n^2)$ mirrors for the second stage of reflection (see [2, 16]). We observe that each mirror at a certain stage of reflection can create at most n mirrors for the next stage of reflection. So, an obvious loose upper bound on the number of mirrors at the kth stage of reflection, is $O(n^k)$. Let $m_1, m_2, ..., m_t$ be the set of all mirrors up to the kth stage of reflection in any order, and $m_0 = S$. Clearly, $V_k(S) = \cup_{i=0}^{t} V(m_i)$, a union of $(t+1) = O(n^k)$ visibility polygons. Alternatively, $V_k(S)$ may be viewed as the union $\cup_{i=0}^{a} V(c_i)$, where $c_1, c_2, \ldots, c_a$ are reflecting segments over all the k stages of reflection and $c_0 = S$. The major contribution in [16] was showing that the number of reflecting segments at the kth stage of reflection is $O(n^{\lceil (k+1)/2 \rceil})$. Prasad et al. [16] established that $V_k(S)$ is the union of $O(n^{\lceil (k+1)/2 \rceil})$ visibility polygons, leading to an $O(n^{2\lceil (k+1)/2 \rceil + 1})$ bound on the number of vertices of $V_k(S)$. We state the results from [16], some without proof.

Suppose f has a reflecting segment at the $(l-1)$st stage of reflection. The following lemma implies that the entire edge f is a reflecting segment at the $(l+1)$st stage of reflection.

LEMMA 3.1 [16] Let e be an edge of P that has a reflecting segment at the lth stage of reflection. Then, the entire edge e has is a reflecting segment at the $(l+2)$nd stage of reflection.

PROOF. Let the *extension* of e be the maximal connected segment in P that contains e. The extension of e is an open set and the end points of its closure lie on $bd(P)$. The lth reflection from the edge e illuminates points of $bd(P)$ in the neighbourhoods of these endpoints. The $(l+1)$st reflections from these neighbourhoods flood the entire edge e with light. □

The following lemma is a consequence of a novel string of arguments in [16], limiting the number of reflecting segments causing reflections at the lth stage of reflection, well below a naive bound of $O(n^l)$. We omit the proof here.

LEMMA 3.2 [16] Any edge f of P has at most $g(n,l) = 3n^{\lceil l/2 \rceil} + 2\Sigma_{i=1}^{\lceil l/2 \rceil -1} n^i$ reflecting segments at the $(l+1)$st stage of reflection.

Following the above results, $V_k(S)$ is the union of of $O(n^{\lceil (k+1)/2 \rceil})$ visibility polygons. We quote this result and its proof from [16].

LEMMA 3.3 [16] $V_k(S)$ is the union of $O(n^{\lceil (k+1)/2 \rceil})$ visibility polygons.

PROOF. We know that $V_1(S)$ is the union of $O(n)$ visibility polygons including $V(S)$. $V_2(S)$ is the union of $V(S)$, $V_1(S)$, and the visibility polygons in $\{V(e)|\ e$ is a mirror at the second stage of reflection$\}$. Since $V_1(S)$ has $\Theta(n^2)$ edges [2], the number of such mirrors at the second stage of reflection is $O(n^2)$. So, $V_2(S)$ is the union of $O(n^2)$ visibility polygons. The lemma therefore holds for $k = 1$ and 2.

Now consider cases where $k \geq 3$. Edges of P that get lit up for the first time due to at most $k-3$ reflections, get fully lit up after the $(k-1)$st reflection (see Lemma 3.1). Each such edge e of P contributes the region $V(e)$ to $V_k(S)$ in the kth stage of reflection. There are at most n such visibility polygons, which are subsets of $V_k(S)$. The remaining edges of P have reflecting segments only at the $(k-1)$st and kth stages of reflection. Applying Lemma 3.2, it follows that $V_k(S)$ is the union of the visibility polygons of at most $n + n * g(n,k-2) + n * g(n,k-1) = O(n^{\lceil (k+1)/2 \rceil})$ reflecting segments.

□

COROLLARY 3.4 [16] The number of edges of $V_k(S)$ that lie on $bd(P)$ is $O(n^{\lceil (k+3)/2 \rceil})$.

PROOF. $V_k(S)$ is the union of $O(n^{\lceil (k+1)/2 \rceil})$ visibility polygons, each having $O(n)$ edges. □

Corollary 3.4 gives a bound on the number of edges of $V_k(S)$ that lie on $bd(P)$. However, $V_k(S)$ also has edges and vertices lying inside P.

LEMMA 3.5 [16] Let m be a segment lying on an edge of P. Any segment $s \subseteq int(P)$ can intersect at most two shadow edges of $V(m)$.

PROOF. The segment s can not intersect any edge of $V(m)$ that lies on $bd(P)$. So, s must intersect a shadow edge. If s is a subset of a shadow edge then it clearly can not intersect another shadow edge because no two shadow edges of $V(m)$ intersect. Otherwise, s may intersect a shadow edge e at a single point p. The part of s that lies inside the shadow $Sh(e)$, does not intersect $V(m)$ or any other shadow edge of $V(m)$ because a shadow can not contain another shadow edge in its interior. The part of s that lies outside $Sh(e)$ may intersect another shadow edge f of $V(m)$ and enter the shadow $Sh(f)$. Using similar arguments we can show that the part of s that lies inside this shadow too can not intersect any other shadow edge of $V(m)$.

□

LEMMA 3.6 [16] Let c_1,c_2,$\ldots$,c_a be reflecting segments on the edges of P. The number of vertices in the union of the visibility polygons $V(c_i), 1 \leq i \leq a$ is $O(a^2n)$.

PROOF. The total number of vertices of the union of these a visibility polygons include the vertices of these polygons and the vertices created by the intersection of these polygons. The total number of edges in these visibility polygons is $O(an)$. Now consider the intersections between edges of the visibility polygons that lie inside P. By Lemma 3.5, we know that each of these $O(an)$ edges can intersect each of the a visibility polygons at most twice, giving a total of $O(a^2n)$ intersections. □

THEOREM 3.7 $V_k(S)$ has $O(n^{2\lceil(k+1)/2\rceil+1})$ vertices.

PROOF. Follows from Lemmas 3.3 and 3.6. □

The computational complexity of the algorithm for computing $V_k(s)$ as in [16], based on the above theorem is summarized in the following theorem; we omit the details of the algorithm and its analysis.

THEOREM 3.8 $V_k(S)$ can be computed in $O(k * n^{2\lceil(k+1)/2\rceil+1} \log n)$ time.

4 Visibility by multiple specular reflections

It is known that in the case of specular reflection, the region $V_1(s)$ has at most $\Theta(n^2)$ edges and vertices [2]. The construction that achieves the $\Omega(n^2)$ bound for the number of islands or holes in $V_1(s)$ is similar to the example in Figure 4.3. The regions marked 2 are visible due to reflection at some mirror; the region marked 1 is not visible from any mirror. By using two sets of mirrors creating horizontal and vertical beams as in Figure 4.3, $\Omega(n^2)$ islands can be created.

Aronov at al. [1], proved a lower bound of $\Omega((n/k - \Theta(1))^{2k})$ on the combinatorial complexity of the visible region $V_k(s)$ when at most k specular reflections are permitted, for a fixed k. This was done in a novel construction using k "gadgets" called *convex mirrors* (CMs), each consisting of $N = \lfloor n/k \rfloor - \Theta(1)$ segments called *facets*. The k gadgets cause k successive reflections for *beams* starting at s. Each beam is the collection of rays emanating from s in a continuous interval of directions and then suffering reflections off the same sequence of *polygon edges*, before arriving at the portion of the polygon where many (very thin) beams intersect creating numerous islands. The first CM in the construction splits an incident beam into N beams and converges these into the next CM, which in turn creates N^2 beams from N incident beams. Finally, after reflections at the k CMs, we have N^k beams, most of which intersect creating a patter of complexity $\Omega(N^{2k})$. See [1] for details. A long and non-trivial analysis in [1] proves a matching upper bound of $O(n^{2k})$ on the combinatorial complexity of $V_k(s)$; an $O(n^{2k} \log n)$ time and $O(n^{2k})$ space algorithm for computing $V_k(s)$ follows from this bound. Note that this bound is much larger than the combinatorial complexity of $V_k(s)$ in Theorem 3.7 for the case of multiple diffuse reflections. This upper bound in Theorem 3.7 however, is not matched with a lower bound better than that shown in Figure 3.2, for any number k of diffuse reflections.

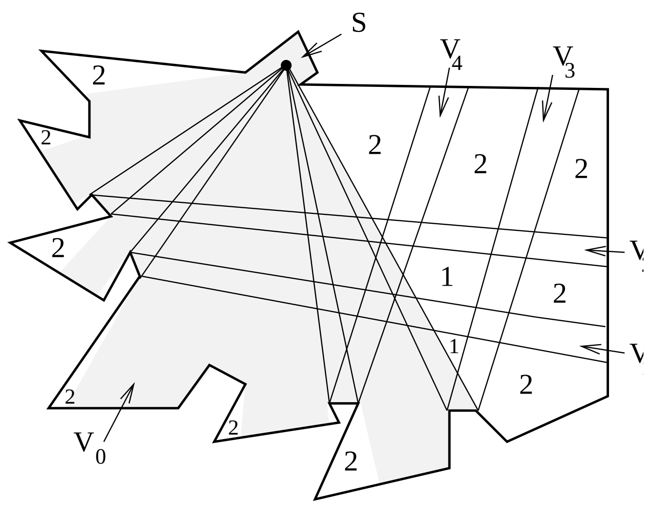

Figure 4.3: $V_1(s)$ has $\Omega(n^2)$ holes.

5 The construction of a hole for two diffuse reflections

Now we show that $V_2(s)$ may have a hole. In what follows, we state this result in detail as in [15]. We show the existence of an island or hole by constructing a simple polygon P containing a point s such that $V_2(s)$ has a triangular hole with vertices t, q and r (see Figure 5.1 and 5.2). Note that rays emanating from s are drawn with dashed lines. Rays that emerge after the first reflection (from points visible to s) are drawn with normal lines. Rays that emerge from points of $bd(V_1(s))$ (after the second reflection) are drawn with lines having alternating dots and dashes. The interior of the triangle tqr is invisible from s, from all points of $bd(V_0(s))$, and, from all points on $bd(V_1(s))$. Note that the first reflection takes place on points of $bd(V_0(s))$ and the second reflection takes place on points of $bd(V_1(s))$.

We need a definition. We say that a point p *blocks* a region R of P from another region S if the shortest path from every point in R to every point in S passes through p.

Clearly, no ray can reach the hole directly from s. This is because b blocks s from the (proposed) hole.

LEMMA 5.1 The triangle tqr is not visible from s after the first reflection.

PROOF.

Consider $bd(V_0(s))$. We start by classifying portions of $bd(P)$ that form $bd(V_0(s))$. There are three portions- $bd(x,y)$, $bd(w,b)$ and $bd(a,v)$. These constitute all points on $bd(V_0(s))$. The sequences of points, $< s,a,y >$, $< s,b,x >$ and $< s,v,w >$ are collinear. How do we guard the hole from the points of $bd(V_0(s))$? The shortest path from any point in $bd(x,y)$ to any point in the hole has to pass through j. In other words, j *blocks* the hole from $bd(x,y)$. So, no point in $bd(x,y)$ can see the hole. Similarly, b blocks $bd(w,b)$ from the holes, and, c blocks $bd(a,v)$ from the hole. So, no point on $bd(V_0(s))$, can see the hole. We conclude that the hole is not visible from any point of $bd(V_0(s))$.

$\square$

All we need to do now is to show that $bd(V_1(s))$ can not see the hole. We start by classifying portions of $bd(P)$ that form $bd(V_1(s))$. These are

1. Points of $bd(a,b)$ visible from $bd(a,v)$, $bd(w,b)$ or $bd(x,y)$. (Note that $bd(x,y)$ is the only portion of $bd(P)$ visible from s outside $bd(a,b)$.)
2. Points of $bd(i,f)$ and $bd(b,c)$, visible from $bd(a,v)$. (Points of $bd(a,v)$ can see points within $bd(i,f)$ because (i) the vertex i is to the right of the ray starting at a and passing through c, and, (ii) the vertex f is to the left of the ray starting at v and passing through a. Observe that $bd(a,v)$ can also see points on $bd(b,c)$. These are the only portions of $bd(P)$, visible from $bd(a,v)$.)
3. Points of $bd(h,a)$ visible from $bd(w,b)$. (Points in $bd(w,b)$ can see points of $bd(P)$ only within $bd(d,a)$, or more precisely, points within $bd(d,k)$, where e,b,a and k are collinear.)
4. Points of $bd(b,c)$ visible from $bd(x,y)$. (The point c is to the left of $ray(y,j)$. So, all points of $bd(V_1(s))$ on $bd(b,j)$ are to the right of $ray(y,j)$).
5. Points of $bd(h,a)$ visible from $bd(x,y)$.
6. Points of $bd(i,h)$ visible from $bd(x,y)$.

The remaining points on $bd(P)$ can not be on $bd(V_1(s))$.

Now we show that none of these six possible portions of $bd(V_1(s))$ can illuminate the triangular hole tqr of $V_2(s)$.

Figure 5.4: Existence of hole tqr in $V_2(s)$

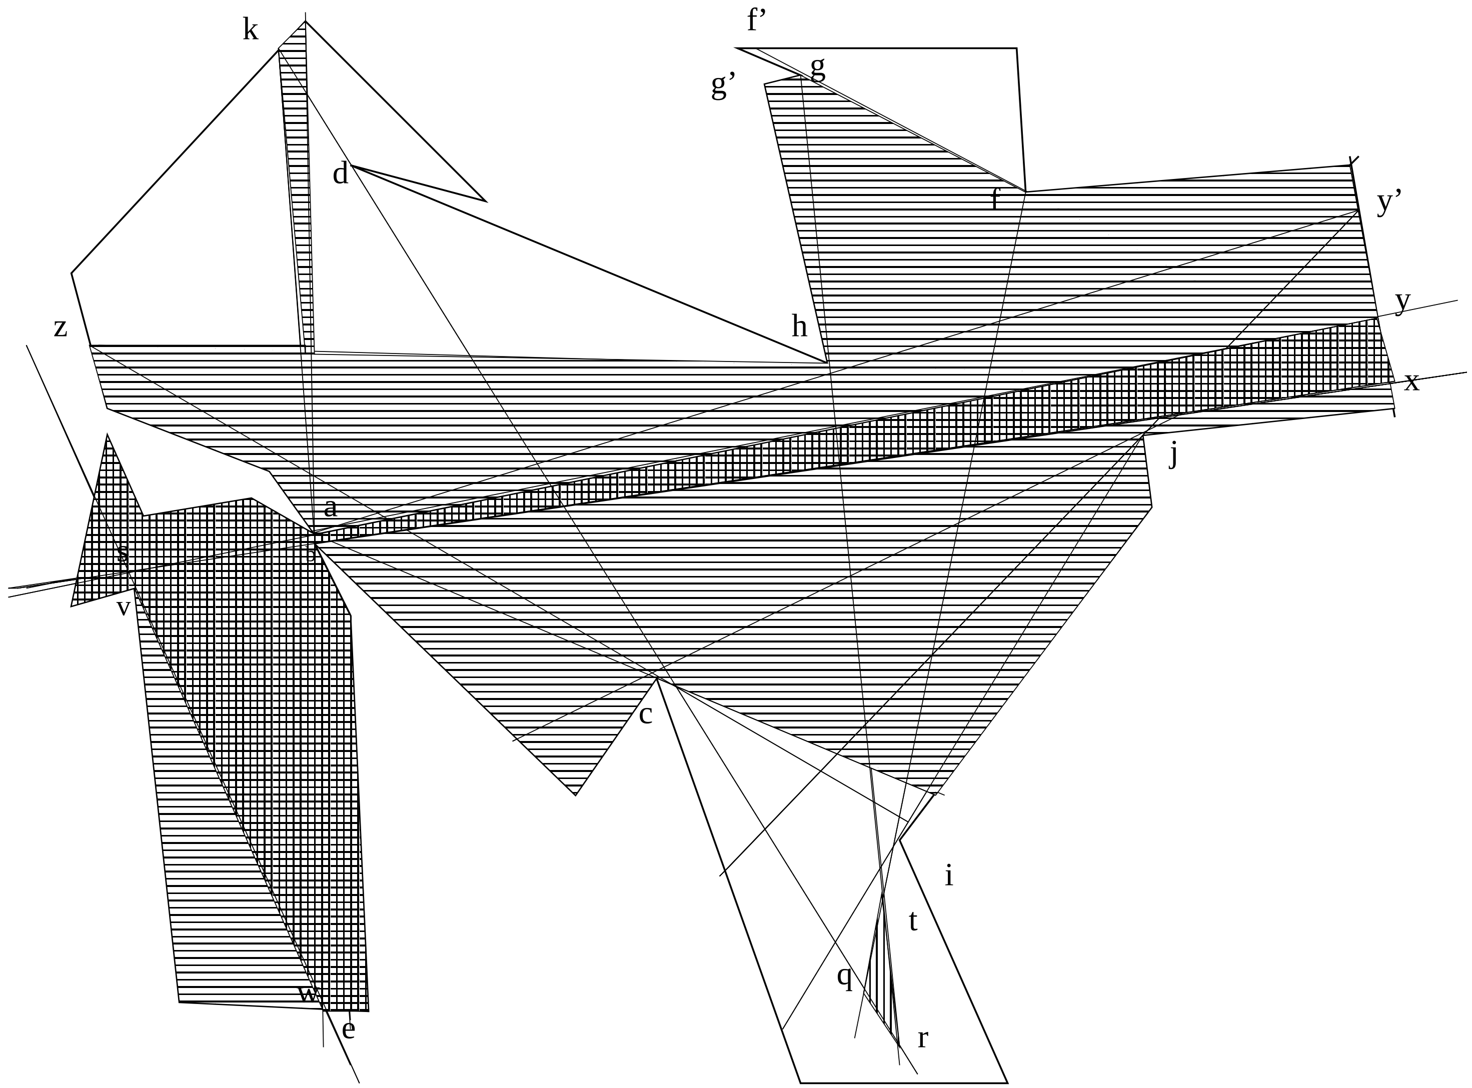

Figure 5.1. Regions visible directly and after one reflection.

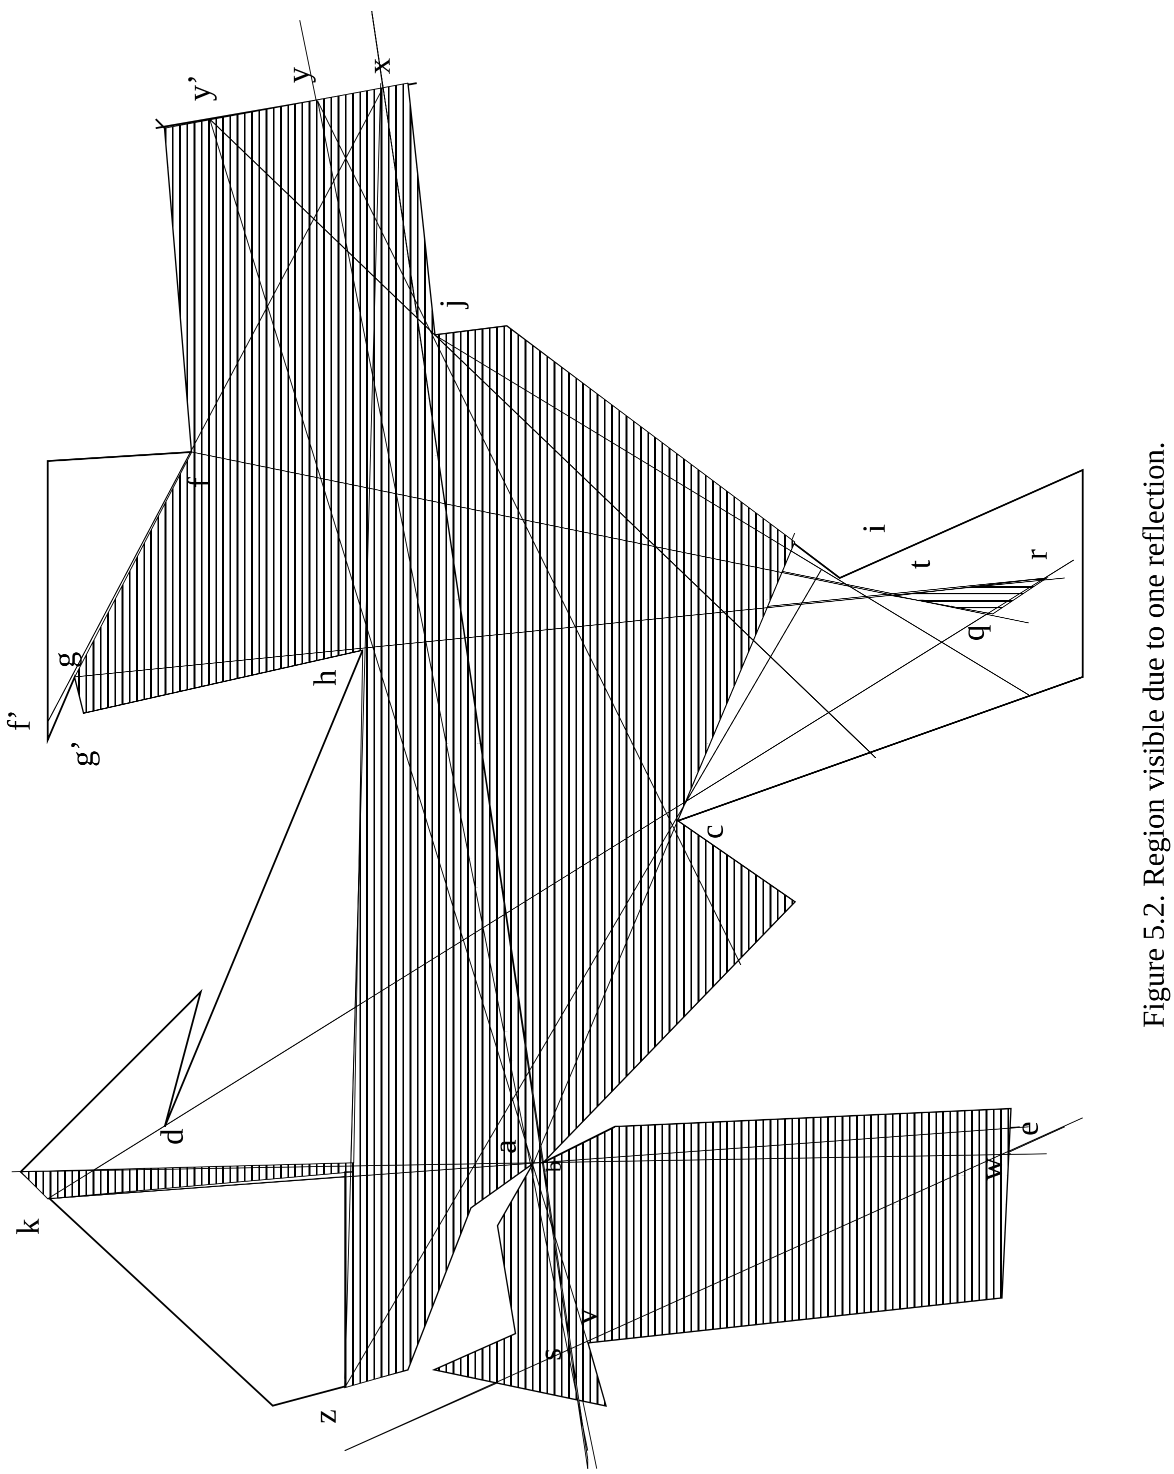

Figure 5.5: Existence of hole tqr in $V_2(s)$

1. The second reflection at $bd(a,b)$ causes rays to miss the hole due to the blocker c. This is due to the fact that the shortest path from any point in $bd(a,b)$ to any point in the hole tqr, has to cross segment ac. In particular, points in $bd(v,b)$ are blocked from the hole by b.

2. Consider points of $bd(V_1(s))$ on $bd(i,f)$. It is sufficient to consider the extreme points y' and c' of $bd(i,f)$ in $bd(V_1(s))$. The point y' is collinear to v and a. The point j blocks y' and the whole of $bd(j,y')$ from the hole. The point c' is collinear with a and c. The point i blocks the whole of $bd(i,j)$ from the hole tqr.

   Now consider points of $bd(b,c)$ that belong to $bd(V_1(s))$ due to first reflection at $bd(a,v)$. These points are blocked from the hole by c.

3. Only points within $bd(d,k)$ in $bd(h,a)$ can cause the second reflection from rays undergoing the first reflection on $bd(w,b)$. All these points in $bd(d,k)$ are blocked by d i from the hole tqr. Indeed, k, d, q and r are collinear, defining the edge qr of tqr. Note that q is the intersection of $ray(k,d)$ and $ray(f,i)$ and, r is the intersection of $ray(k,d)$ and $ray(g,h)$.

4. Consider points of $bd(b,c)$ visible from xy. Points of $bd(b,c)$ can not see tqr because they are blocked by c.

5. Consider $bd(d,a)$. Points in $bd(d,z)$ are blocked by h from points on xy. Since z, h and x are collinear and tqr is to the right of $ray(z,c)$, all points of $bd(z,a)$ are blocked from tqr by c.

6. Consider $bd(i,h)$ and xy. Points on $bd(j,y)$ are blocked from tqr by j. So we need to consider only $bd(y,h)$. Any point on $bd(y,f)$ may reflect into the region to the right of $ray(f,i)$ but not to the left of $ray(f,i)$. In the construction, f, i, t and q are collinear and tq defines an edge of the hole tqr. Note that $ray(g,h)$ and $ray(f,i)$ intersect at vertex t.

   Now consider point f' on $bd(f,h)$ where f' is the point where $ray(x,f)$ meets $bd(P)$. We construct $bd(f,g)$ so that the shortest path from f' to h passes through g and g lies in the segment joining f and f'. Points on xy can not see any point in the interior of $bd(f,f')$. Since g, h, t and r are collinear, (i) the shortest path from any point in $bd(f',g)$ to any point inside tqr has to pass through blockers h and g, and, (ii) the shortest path from any point in $bd(g,h)$ to any point in tqr has to pass through blocker h. Indeed, $ray(g,h)$ defines the edge tr of hole tqr.

   Note that reflections on edge $g'h$ will not reach the interior of triangle tqr because triangle tqr is to the right of $ray(g',h)$.

In summary, we have established that tqr is a hole in $V_2(s)$. We state the result as a theorem.

Theorem 5.2 It is possible to construct a simple polygon P with internal points s, t, q and r such that (i) tqr is a triangle totally contained in $int(P)$ (ii) all points inside the triangle tqr belong to $P \setminus V_2(s)$ and (iii) tq, tr and rq are on $bd(V_2(s))$.

6 Existence of multiple holes for multiple diffuse reflections

We extend the result of the previous section by showing that the portion $V_2(s)$ of an n-sided simple polygon (visible from an internal point s due to two diffuse reflections) can have n/c holes, for some constant $c > 0$. Consider Figures 5.1 and 5.2. We call $bd(b, a)$, a *component.* This component is a part of $bd(w, v)$. The component has the hole tqr. We place more such components like $bd(b, a)$ on $bd(w, b)$, where each component has a hole similar to tqr. Since the number of edges in each component is a constant independent of the number of edges n of the polygon, the total number of holes is proportional to n.

We show that multiple components can indeed be placed on $bd(w, b)$. For placing a new component $bd(b_1, a_1)$, we make an opening $a_1 b_1$ like ab in the interior of $bd(w, b)$. Consider the pair of components say $bd(b, a)$ and $bd(b_1, a_1)$. We can ensure that these two components do not overlap by suitably elongating the $bd(v, b)$ portion of the polygon, by pulling down and enlarging the edge containing w and e. This deformation can be done without affecting the collinearity triples $< s, v, w >$, $< w, b, a >$. $< e, b, a >$ and $< b, a, k >$. Since w is moved down and away from s, we have to suitably shift v without affecting v's collinearity with a and y'. This can be done by moving v along $ray(y', a)$, away from a, so that it becomes aligned with s and (the shifted) w.

It is easy to construct the component $bd(b_1, a_1)$ with a hole that misses the source s; simply keep the hole outside the cone made by $ray(s, a_1)$ and $ray(s, b_1)$. Now consider the rays after the first reflection. The rays that get reflected in one component can not illuminate a hole in another component; such a ray has to undergo the second reflection in $bd(a, b)$ and such reflections can be made to easily miss the hole in any component (as shown in Case 1 in Section 5). If the first reflection is in $bd(a, b)$ and the second reflection is within the component $bd(b_1, a_1)$, then we make the second reflection miss the hole by constructing the $bd(b_1, a_1)$ as in Cases 2 and 3 of the six cases in Section 5. If both the reflections take place inside $bd(b_1, a_1)$, construction of $bd(b_1, a_1)$ must ensure the Cases 4, 5 and 6 of Section 5. We omit the details and summarize our result in the following theorem.

THEOREM 6.1 It is possible to construct an n-sided simple polygon P with an internal point s such that $V_2(s)$ has $\Omega(n)$ holes.

7 Conclusion

The combinatorial complexity of the region of a simple n-gon, visible from a point light source due to $k \geq 1$ diffuse reflections is known to be $O(n^{2\lceil (k+1)/2 \rceil + 1})$. There is a huge gap between the $\Omega(n^2)$ lower bound and the upper bound. Bridging this gap remains an open problem. All our studies prompt us to believe that the real bound is $O(n^2)$. In this paper we have presented the construction as in [15], showing the existence of holes in regions visible due to multiple diffuse reflections. Establishing a tight upper bound on the number of holes is an important open problem. Such upper bounds may help designing faster algorithms for computing regions visible due to multiple diffuse reflections.

Much work needs to be done in three dimensional visibility; no work seems to have been done in three dimensional visibility with reflections. Studying visibility with reflections in three dimensional space with plane or curved reflecting surfaces is likely to be a very fertile area of research.

References

[1] B. Aronov, A. Davis, T. K. Dey, S. P. Pal, and D. C. Prasad. Visibility with multiple reflections. *Discrete and Computational Geometry*, 20:61–78, 1998.

[2] B. Aronov, A. Davis, T. K. Dey, S. P. Pal, and D. C. Prasad. Visibility with one reflection. *Discrete and Computational Geometry*, 19:553–574, 1998.

[3] G. Das and G. Narasimhan. Optimal linear-time algorithm for shortest illuminating line segment in a polygon. In *Proceedings of the ACM Symposium on Computational Geometry*, pages 259–266, 1994.

[4] S. K. Ghosh. Computing the visibility polygon from a convex set and related problems. *Journal of Algorithms*, 12:75–95, 1991.

[5] S. K. Ghosh, A. Maheshwari, S. P. Pal, S. Saluja, and C. E. Veni Madhavan. Computing the shortest path tree in a weak visibility polygon. In *Proceedings of the 11th Conference on Foundations of Software Technology and Theoretical Computer Science, Lecture Notes in Computer Science 560, Springer-Verlag*, pages 369–389, December 1991.

[6] S. K. Ghosh, A. Maheshwari, S. P. Pal, S. Saluja, and C. E. Veni Madhavan. Characterizing and recognizing weak visibility polygons. *Computational Geometry: Theory and Applications*, (3):213–233, 1993.

[7] H. El Gindy and D. Avis. A linear algorithm for computing the visibility polygon from a point. *Journal of Algorithms*, 2:186–197, 1981.

[8] L. Guibas, J. Hershberger, D. Leven, M Sharir, and R. Tarjan. Linear time algorithms for visibility and shortest path problems inside triangulated simple polygons. *Algorithmica*, 2:209–233, 1987.

[9] E. Gutkin. Billiards in polygons. *Phys. D.*, 19:311–333, 1986.

[10] V. Klee. Is every polygonal region illuminable from some point? *Amer. Math. Monthly*, 76:180, 1969.

[11] V. Klee and S. Wagon. Old and new unsolved problems in plane geometry and number theory. *Mathematical Association of America, Washington, D.C.*, 1991.

[12] D. T. Lee. Visibility of a simple polygon. *Computer Vision, Graphics and Image Processing*, 22:207–221, 1983.

[13] D. T. Lee and A. K. Lin. Computing the visibility polygon from an edge. *Computer Vision, Graphics and Image Processing*, 34:1–19, 1986.

[14] J. O'Rourke. *Art gallery theorems and algorithms.* Oxford University Press, 1987.

[15] S. P. Pal and D. Sarkar. On multiple-connectedness of regions visible due to multiple diffuse refelctions. *Technical Report No. TR/IIT/CSE/SPP1, Department of Computer Science and Engineering, IIT Kharagpur, 721302, India*, May 2001.

[16] D. C. Prasad, S. P. Pal, and T. K. Dey. Visibility with multiple diffuse reflections. *Computational Geometry: Theory and Applications*, 10:187–196, 1998.

[17] T. Shermer. Recent results in art galleries. *Proceeding of the IEEE*, 80:1384–1399, September 1992.

[18] George W. Tokarsky. Polygonal rooms not illuminable from every point. 102(10):867–879, December 1995.