Visibility with One Reflection

Boris Aronov[1]
Alan R. Davis[2]
Tamal K. Dey[3]
Sudebkumar P. Pal[3]
D. Chithra Prasad[3]

October 24, 1997

Abstract

We extend the concept of the polygon visible from a source point S in a simple polygon by considering visibility with two types of reflection, *specular* and *diffuse*. In specular reflection a light ray reflects from an edge of the polygon according to the rule: the angle of incidence equals the angle of reflection. In diffuse reflection a light ray reflects from an edge of the polygon in all inward directions. Several geometric and combinatorial properties of visibility polygons under these two types of reflection are described, when at most one reflection is permitted. We show that the visibility polygon $Vs(S)$ under specular reflection may be non-simple, while the visibility polygon $Vd(S)$ under diffuse reflection is always simple. We present a $\Theta(n^2)$ worst-case bound on the combinatorial complexity of both $Vs(S)$ and $Vd(S)$ and describe simple $O(n^2 \log^2 n)$ time algorithms for constructing the sets.

1 Introduction

Visibility problems have a long history [21], with one of the first questions, the "art gallery problem," posed in 1973 by Klee [14], who asked how many stationary guards are needed to guarantee that every point of a polygonal room is visible from at least one guard. Many aspects of this problem have been considered, including the use of different types of guards, with different powers of vision. The *visibility polygon* with respect to a given point inside a simple polygon is the set of points that are directly visible to that point. Linear-time algorithms for constructing point visibility polygons were presented by ElGindy and Avis [8], by Lee [16], and by Guibas *et al.* [10].

Horn and Valentine introduced the concept of k-link visibility, defining a point D to be k-link-visible to a point S if there exists a k-link polygonal path between them that remains inside the

[1]Department of Computer and Information Science, Polytechnic University, Brooklyn, NY 11201, USA. Work on this paper by Boris Aronov has been supported by NSF Grant CCR-92-11541.

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

[3]Department of Computer Science and Engineering, Indian Institute of Technology, Kharagpur, Kharagpur 721302, India. Tamal K. Dey acknowledges the support of the NSF grant CCR-93-21799 while he was in Indiana University – Purdue University at Indianapolis, Indiana, USA. Work on this paper by S.P. Pal is partially supported by a research grant from the Jawaharlal Nehru Center for Advanced Scientific Research, Bangalore, India.

polygon [12]. Surveys of link distance work can be found in [19, 18]. O'Rourke's book on art gallery theorems [20] is a comprehensive foundation, while Shermer's and Urrutia's surveys [25, 26] provide a more recent bibliography on art gallery problems.

Billiard paths and illumination of regions bounded by mirrors are two more variations of visibility problems. A good description and summary of some of the open problems can be found in the book by Klee and Wagon [14]. Much of the work on billiard paths has been done in the field of ergodic theory, and the reader can find a full introduction in the paper by Gutkin [11].

Visibility has been studied in computer graphics in relation to ray tracing and hidden surface elimination algorithms. One of the early discussions of the problem can be found in [9]. Reif, Tygar, and Yoshida have investigated the difficulty of ray tracing in optical systems [24].

This paper investigates the geometric and combinatorial structures of visibility polygons under assumption of a single reflection. We model two different types of reflection, specular and diffuse. In specular reflection we assume that energy is reflected along a specific direction and that it abides by the standard law of reflection: the angle of incidence equals the angle of reflection. This type of reflection occurs from a perfectly mirror-like surface. In diffuse reflection, energy is reflected in all inward directions. If the reflecting surface is rough, this type of reflection is more common.

We show that, with a single reflection, the (point) visibility polygon under specular reflection $Vs(S)$ may be non-simple, while the (point) visibility polygon under diffuse reflection $Vd(S)$ is always simple. We prove a tight $\Theta(n^2)$ bound on the worst-case combinatorial complexity (i.e., number of vertices and edges on the boundary) of both types of visibility polygons and describe simple $O(n^2 \log^2 n)$ algorithms for computing them. Our paper only addresses visibility with single reflection. For a discussion of multiple reflections and other extensions, see Section 7.

The remainder of this paper is divided into six sections. Section 2 contains basic definitions. Section 3 is concerned with geometric properties of $Vs(S)$. Section 4 analyzes the complexity of $Vs(S)$. Section 5 discusses the important features of $Vd(S)$ and gives a bound on its complexity. Section 6 presents an algorithm for constructing $Vs(S)$ and $Vd(S)$. Section 7 contains conclusions and discussion.

2 Definitions

Let $P = (v_0, v_1, \ldots, v_{n-1})$ be a simple polygon with no three collinear vertices. Let $int(X)$ and $bd(X)$ denote the relative interior and the boundary of a region X, respectively.

Two points in P are said to be *(directly) visible* to each other if the interior of the line segment joining them lies in $int(P)$. A point p is said to be *weakly visible* from a line segment $st \subset P$, if there is a point z in the interior of st such that p is visible from z. The *(point) visibility polygon* of S, $V(S,P)$, is the set of all points in P that are directly visible to the point light source S. If the polygon is understood in context we shorten this notation to $V(S)$. Let a *chord* be any segment in P whose interior is contained in $int(P)$. $V(S)$ is cut out of P by a collection of disjoint chords, that we call *windows* of $V(S)$.

In *specular* reflection a light ray is reflected according to the standard law of reflection: the angle of incidence equals the angle of reflection. Light rays that hit a vertex of P will terminate without reflection, which is the most common definition in the literature. Letting vertices reflect light off one or both incident edges according to the reflection law, and allowing light to have "grazing contact" with the boundary of P is another possibility. In non-degenerate situations, the two definitions only differ in classifying points on the boundary of the visibility region, both for direct visibility and visibility with one specular reflection. In the case of diffuse visibility, however, allowing light arriving at a vertex to reflect in *all* inward directions produces a dramatic change in the region visible with one reflection, as it becomes possible to light up portions of the polygon lying "around the corner."

In order to make our analysis easier, we will make several non-degeneracy assumptions on the shape of the polygon and position of the source. These assumptions do not change the shape of the visible region in a fundamental fashion. One assumption, already mentioned, is that no three vertices of P are collinear. We also require that the source S is not collinear with a pair of vertices of P. Additional non-degeneracy assumptions related to reflected light rays will be mentioned below.

A point D is said to be *indirectly visible* from the source point S after one *specular* reflection, if there exists a point p lying in the interior of an edge of P, visible from S and D, and such that Sp and pD lie on the opposite sides of the inward normal to $bd(P)$ at p and make the same angle with it. We define a *(specular) visibility path* to be any such polyline SpD between two points, with each segment being a *link* in the path.

In *diffuse* reflection a light ray incident on a point in the interior of an edge of P is assumed to be reflected in all inward directions. A point D is said to be *indirectly visible* from a point S after one *diffuse* reflection, if there exists a point p lying in the interior of an edge of P and directly visible from S and D. We define a *(diffuse) visibility path* to be any such polyline SpD between two points, with each segment being a *link* in the path. In Figure 2.1, D_1 is indirectly visible from S under

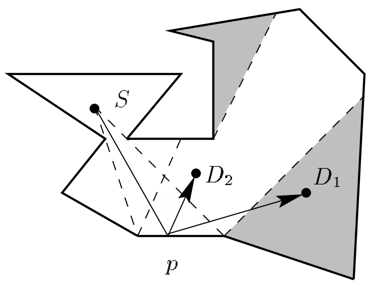

Figure 2.1: Types of indirect visibility.

diffuse reflection. D_2 is indirectly visible from S under specular and diffuse reflection. However, D_1 is not visible from S under specular reflection.

We define the *specular visibility region*, $Vs(S)$, to be the set of all points D in P directly or indirectly visible to S under specular reflection. The *diffuse visibility region*, $Vd(S)$, is the set of all points D in P directly or indirectly visible to S under diffuse reflection.

3 Specular Visibility Regions

In this section, all reflections are assumed to be specular.

Let e_i be an edge of P weakly visible from S. Let m_i be the portion of e_i directly visible from S. We will refer to m_i as a *mirror.* Define the *mirror visibility polygon* of m_i, $V_i(S)$, as the set of points of P that can see S after reflecting off m_i. Let the image of S under reflection about e_i be S^i. Let Δ^i be the triangle formed by S^i and m_i. Then $V_i(S) = P \cap V(S^i, P \cup \Delta^i)$. (More precisely, $P \cup \Delta^i$ denotes here the Riemann surface obtained by "gluing" P to Δ^i along m_i, which handles the possibility that, as a subset of the plane, Δ^i overlaps P.) This is the dark-shaded region in Figure 3.1. (Note that in Figures 3.1, 3.2, 3.7, and 5.4, unlike the remaining figures in this paper, we use shading to indicate presence of light, rather than its absence.) If $a \leq n$ is the number of edges

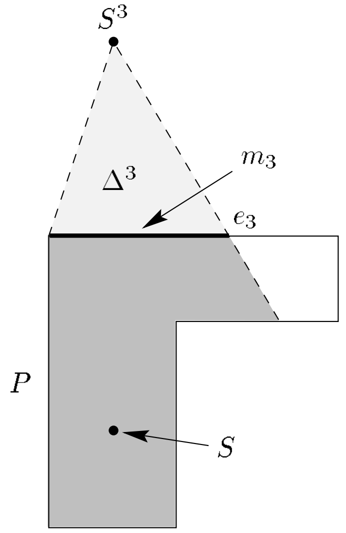

Figure 3.1: Mirror visibility polygon; shading indicates light.

of P weakly visible from S, we let $\{m_0, m_1, \ldots, m_{a-1}\}$ denote the set of mirrors.

By definition, $Vs(S) = V(S) \cup \bigcup_{i=0}^{a-1} V_i(S)$. In Figure 3.2, $V(S) \subset Vs(S)$ is shaded darker,

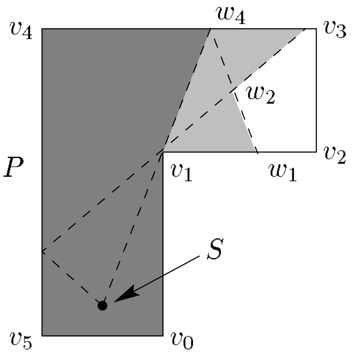

Figure 3.2: Specular visibility region $Vs(S)$; shading indicates light.

$Vs(S) \setminus V(S)$ is shaded lighter. Note that $Vs(S)$ may extend $V(S)$ without completely covering P. We define the regions that are not visible directly or indirectly to be blind spots. More precisely, a *blind spot* is a connected component of $P \setminus Vs(S)$. By computing $Vs(S)$ we mean computing its boundary, so it is sufficient to describe these blind spots in order to have a description of $Vs(S)$.

Recall that our ultimate goal is to prove a tight bound on the maximum combinatorial complexity of $Vs(S)$. This is done by first locally classifying blind spot vertices and then bounding the global number of different types of vertices. Roughly, we prove that blind spots are mostly convex and their vertices are of four different types, three of which can only occur once per pair of mirror visibility polygons, while vertices of the last class have to be counted differently. We now present the details of the argument.

As a finite union of polygons, $Vs(S)$ is a polygonal region contained in P. $Vs(S)$ is connected since every point in it is connected to S by a path consisting entirely of points directly or indirectly visible to S and thus contained in $Vs(S)$. A vertex, such as w_2 in Figure 3.2, generated by the intersection of edges of the relative boundary of two sets among $V(S)$ and $V_i(S)$, $i = 0, \ldots, a-1$, will be called an *interior vertex* of $Vs(S)$. The boundary of $Vs(S)$ may consist of edges of P, such as v_0v_1; portions of edges of P, such as v_1w_1; and segments in the interior of P, such as w_1w_2.

The edges of the relative boundary of $Vs(S)$ come from the edges of the relative boundary of the sets $V(S)$ and $V_i(S)$, $i = 0, \ldots, a-1$. We have assumed that visibility terminates at a vertex, so we define a *(specular) vpath limit* to be any 2-link path in P from S to $bd(P)$ that obeys the reflection property but passes through a vertex of P; this may happen at, before, or after the point of reflection. Refer to Figure 3.3. As mentioned earlier, we assume that another type of degeneracy does not occur, namely that it is not possible for a vpath limit to go through two vertices of P. Informally, we assume that the light source S is not collinear with a pair of vertices of P, even after one reflection. The portion of a vpath limit beyond the occluding vertex is not visible from S, and the region adjacent to one side of that portion of the vpath limit is in shadow, while the opposite side of it is necessarily lit. These subsegments will be called *shadow edges*; they bound $V(S)$ and mirror visibility polygons and may become part of the boundary of $Vs(S)$. However, some portions of shadow edges or entire shadow edges are made visible from other directions, such as v_1w_2 in Figure 3.2, and therefore do not become boundary edges of $Vs(S)$. $V_i(S)$ is a simply connected portion of P cut out by disjoint chords, which are shadow segments contained in second links of appropriate vpath limits. Windows of the direct visibility polygon $V(S)$ are the shadow edges bounding $V(s)$. They cut $V(S)$ out of P. For example, v_1w_4 in Figure 3.2 is a window of $V(S)$.

We have different situations according to whether the shadow after the reflection is inside (*in*) or outside (*out*) of the angle of reflection, see Figure 3.3. In addition, a vpath limit can touch a reflex vertex *before* or *after* reflection, or it can touch the endpoint (*end*) of the edge it reflects off of. Recall that points on a vpath limit are visible to the source up to and including the first reflex vertex on the path. In the figures, small shaded wedges indicate reflex vertices of P. Small circles emphasize that the shadow is generated because the reflection occurs at the end of an edge of P. Shadow edges are defined locally, according to the effects along a single path. They may be lit from other directions, thereby the shadow region along them is reduced. This is illustrated by the dashed lines in Figure 3.3, where the shadow regions above the dashed lines are lit up.

We give without proof the following fundamental facts.

Fact 3.1 *(a) In a simple polygon, the portion of any segment visible from a point is either empty or forms a connected subsegment. (b) Similarly, the portion of any segment weakly visible from another*

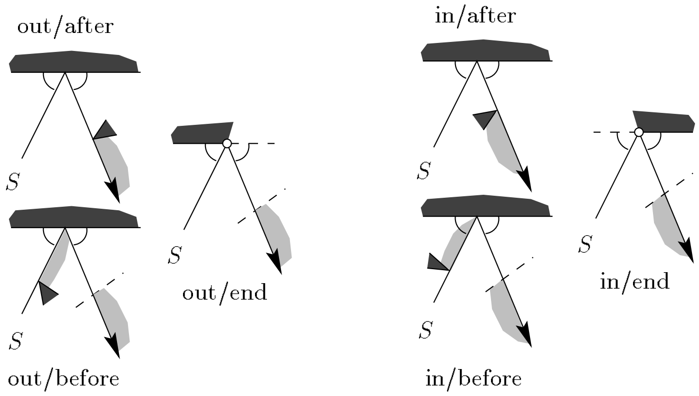

Figure 3.3: Specular vpath limit types; dark shading indicates exterior of P.

segment is empty or forms a connected subsegment.

Fact 3.2 *There can be no part of $bd(P)$ inside any region bounded by a simple curve completely contained in P.*

The first seems to be part of "folklore" that must have been known to the Greeks, though we could not find an explicit published reference. When stated for visibility (from a point) of points lying along a polygon edge, it can be found in [3]. A proof of a stronger statement can be found in [2]. (Briefly, they prove the statement: "The weak visibility region of a relatively convex subset of P is relatively convex". A set is defined to be *relatively convex* if the shortest path connecting any two of its points is contained in the set.) The second fact is equivalent to our assumption that P is a simple polygon and thus is simply connected.

We now turn our attention to classifying blind spots of $Vs(S)$.

A blind spot whose boundary does not intersect $bd(P)$ is called an *interior blind spot.* By placing small mirrors oriented appropriately along edges of our polygon we can produce thin mirror visibility polygons that intersect to create interior blind spots in $Vs(S)$. Refer to Figure 3.4. This shows that $Vs(S)$ is not necessarily simply connected. For example, the interior blind spot labeled 1 in Figure 3.4 was created by encircling it with thin mirror visibility polygons, created by light reflecting off mirrors m_1 and m_2, in an otherwise non-visible region.

We define a *boundary blind spot* to be a connected region of $P \setminus Vs(S)$ which has part of $bd(P)$ as its boundary. Some of the vertices of a boundary blind spot must lie on $bd(P)$ while some may be interior vertices. A boundary blind spot with a single interior vertex is called a *simple boundary blind spot* (labeled 4 in Figure 3.4). As seen in the figure some boundary blind spots, without interior vertices, are generated at the "far ends" of mirror visibility polygons (labeled 2). We can also have boundary blind spots with more than one interior vertex created by shadow edges (such as those labeled 3).

We define a direction along a vpath (or vpath limit) according to increasing distance from S along the path. Suppose two vpath links cross at x. We call the links ℓ and r so that the circular ordering around x in clockwise directions is: ℓ leaving x, r leaving x, ℓ entering x, r entering x. See

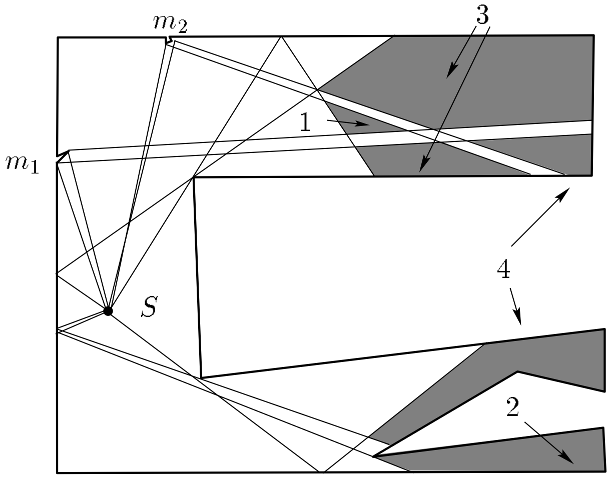

Figure 3.4: Blind spot types.

Figure 3.5(a). Removal of directed link ℓ (r) splits P into two or more pieces, which we classify into those lying locally left of ℓ (r) and those lying locally right of ℓ (r). Removal of ℓ and r together produces four or more connected components, which we classify into left, right, top, or bottom. *Left* (*right*) components lie in the portion of P to the left (resp. right) of both ℓ and r. *Bottom* components lie to the left of ℓ and to the right of r, while *top* components lie to the right of ℓ and to the left of r. Finally, we define a *left (right, top, bottom) quadrant* to be the union of left (resp. right, top, bottom) connected components. Note that this is slightly different from defining a quadrant as the intersection of P with a quadrant of the *plane* created by the two lines; this is apparent when one considers the "handle" on top of Figure 3.5(a). In the figure, left, right, top, and bottom quadrants

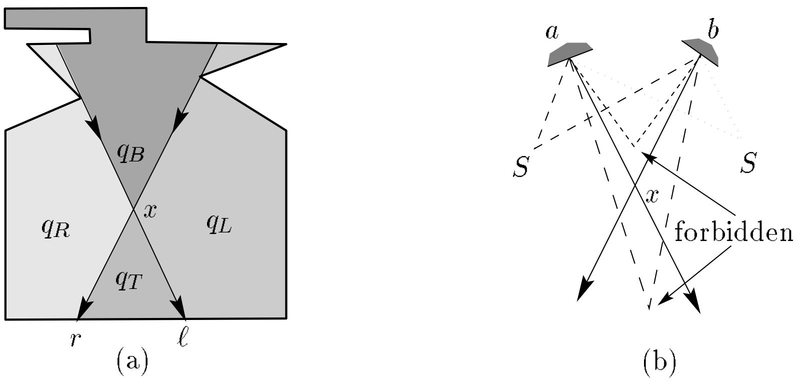

Figure 3.5: Quadrants and position of source. In (a) different quadrants are shaded differently. In (b) shading indicates exterior of P.

are labeled q_L, q_R, q_T, and q_B, respectively.

We now introduce a classification of interior vertices of blind spots. Namely, if x is such a vertex lying at the intersection of two links as above, we call x a *left* (*right, top, bottom*) vertex if the blind spot lies in the left (right, top, bottom, respectively) quadrant with respect to x, locally near x.

Lemma 3.3 *If $x \in int(P)$ is a point not visible to S on the intersection of two vpath limits, then the source S must lie in the closure of either the left or the right quadrant, relative to x.*

PROOF. Consider two segments of vpath limits that intersect at x determining the four quadrants labeled $q_L, \ldots, q_T$, as in Figure 3.5(a). If they are both second links of vpath limits we can without loss of generality indicate the first reflections, a and b, of the vpath limits as in Figure 3.5(b). If the source S lay in q_B, then x would be visible to S since S and x would be opposite vertices of the quadrilateral $Saxb$ whose boundary is completely contained in P. Indeed, in $Saxb$, a and b would necessarily be convex corners, so the open segment Sx would lie in the interior of $Saxb$ and thus in $int(P)$, by Fact 3.2. The case placing S in q_T is similar. Therefore, the only valid placements for S are in q_R or q_L (or on their boundaries) as shown in Figure 3.5(b), as desired.

We now consider the case when one of the intersecting segments is the first link of a vpath limit, i.e., it contains a window. In this case the lemma holds, as S lies on a link which is contained in the closure of either the left or right quadrant.

Finally, x cannot lie on two windows since two windows cannot meet. $\square$

We are interested in counting interior vertices of blind spots. They lie at intersections of the relative boundaries of two (real) mirror visibility polygons, or of a mirror visibility polygon and $V(S)$. Recalling that segments of relative boundaries of $V(S)$ are called "windows," we first examine the vertices obtained as a result of the latter type of intersections.

Lemma 3.4 *The intersection of a fixed edge e of the relative boundary of a mirror visibility polygon with $V(S)$ yields at most one interior blind spot vertex of $Vs(S)$. No bottom vertex can be created in this fashion.*

PROOF. By Fact 3.1(a) e intersects the boundary of $V(S)$ in at most two points, one of which lies on a mirror or at a vertex of P, proving the first part of the lemma.

Now assume the bottom vertex of a blind spot does occur at point x where a window w intersects e. We have defined the positive direction along the vpath limit in such a manner that e must be oriented from $V(S)$ towards $P \setminus V(S)$ near x. Hence if x were a bottom vertex of a blind spot, the blind spot would lie in $V(S)$ near x — a contradiction. $\square$

We will now give a more formal characterization of interior and boundary blind spots.

Lemma 3.5 *An interior blind spot R is a closed convex polygon with each vertex on two vpath limits; it contains exactly one top and one bottom vertex.*

PROOF. R is a component of $P \setminus (V(S) \cup \bigcup_i V_i(S))$. As each of $V(S), V_i(S)$ is a polygonal, relatively open subset of P (except possibly at reflex vertices of P), R is a closed polygonal set. As the relative boundaries of $V(S), V_i(S)$ consist of portions of vpath limits, so does the boundary of R. As $Vs(S)$ is connected, R is simply connected, as a connected component of $P \setminus Vs(S)$. Finally, consider a vertex v of R. As $V(S), V_i(S)$ have no interior vertices, v must lie on the relative boundary of two or

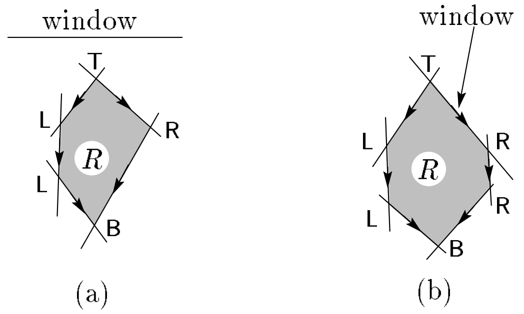

Figure 3.6: Interior blind spots.

more sets from the family $\{V(S), V_1(S), \ldots, V_a(S)\}$ and so on the intersection of two or more vpath limits. Since only one side of a vpath limit is in shadow, only one of the four sectors formed by the two segments can be in shadow, as it will not be lit by either beam of light. Therefore v is a convex vertex of R, so R is convex.

Since R lies outside $V(S)$, there is a window of $V(S)$ between S and R and all vpath limits appearing on $bd(R)$ must intersect that window. The perpendicular distance from the window to any point on a vpath limit increases as we traverse the vpath limit away from S and thus away from the window. Therefore the direction of the edges of R is consistent with increasing perpendicular distance from the line containing the window. As R is convex, it thus must have a unique minimum (top vertex) and a unique maximum (bottom vertex) while the remaining vertices must be side vertices. Figure 3.6(a) illustrates the case where R does not touch the window, while Figure 3.6(b) shows the case when the window overlaps the boundary of R. The proof of the latter case is similar. $\square$

Lemma 3.6 *A boundary blind spot is a closed set (except possibly at reflex vertices of P) bounded by a convex chain and a connected portion of $bd(P)$. It has no bottom vertex and at most one top vertex.*

Proof. Let R be a boundary blind spot. Proceeding as in the proof of Lemma 3.5, we conclude that R is closed (except possibly at reflex vertices of P) and polygonal, and all its interior vertices are convex. If $bd(P)$ appeared on $bd(R)$ more than once, $P \setminus R$ would be disconnected, which is impossible as it would imply that $Vs(S)$ is not connected. Hence R is bounded by a convex chain and a connected portion of $bd(P)$. (Note that a "convex chain" here is just a polygonal curve which makes only left turns as it is traversed in counterclockwise direction around R. We do not claim that it is a portion of the boundary of a convex set, which is the usual definition of a convex chain.)

As for the second claim, for a contradiction, consider a bottom vertex x. By Lemma 3.4, it cannot lie on a window. Let y and z be the points where the (second links of) the two vpath limits containing x intersect the window which separates x from $V(S)$. By definition of a bottom vertex and of a window and by Fact 3.2, the points immediately outside $\triangle xyz$ lie in $Vs(S)$, so the blind spot containing x must be fully contained in $\triangle xyz$ and must thus be an interior blind spot.

Finally, on the bounding convex chain, consider the perpendicular distance to the window separating R from S. Any top vertex is a local minimum on the chain (see the proof of Lemma 3.5). However, any two local minima must be separated by a local maximum, which would have to be a bottom vertex. Thus there is at most one top vertex, as claimed. □

Lemma 3.7 *The relative boundaries of two mirror visibility polygons can not intersect to generate two bottom, two left, or two right blind spot vertices.*

PROOF. Consider the situation depicted in Figure 3.7 where two vpath limits from the source point S reflecting off edge e_j are intersected at points a and b by a third vpath limit reflecting off a different edge e_i. We call this a *core configuration* and observe the following. Consider the four angles adjacent to the segment ab. By Fact 3.2 no reflex vertex of P can occur inside the "triangular region" abS^j

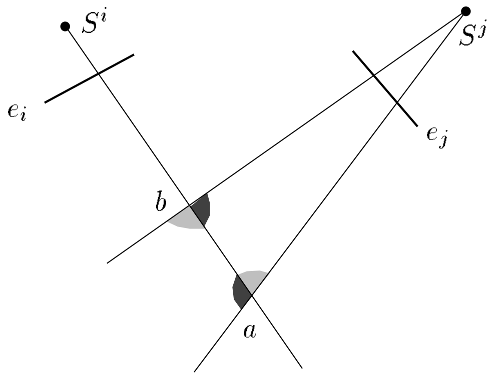

Figure 3.7: Angles lit up and thus not containing incident blind spots; shading indicates presence of light.

whose boundary is contained in P. More precisely, the triangle abS^j can be folded across e_j to form a union of a triangle and a quadrilateral whose interior must not contain a reflex vertex of P by Fact 3.2. In particular, the indicated angles are lit up from S^j and thus a (b) can not be a blind spot vertex with the blind spot lying in one of these angles near a (resp. b). We use this observation about the core configuration repeatedly to prove the lemma.

We proceed by contradiction. If two mirror visibility polygons intersect to generate two bottom, two left, or two right blind spot vertices, the two vertices can either share a vpath limit link, or be created by interaction of four distinct links. Figure 3.8 (3.9) shows all possible configurations of pairs of bottom (resp. left) blind spot vertices. Left and right blind spot spots vertices are handled similarly.

The cases depicted in Figures 3.8(a) and 3.9(a–b) are ruled out immediately by the core configuration observation. The situation in Figure 3.8(b) has two possible candidate pairs of bottom vertices not covered by the other figures: w and y, or x and z. Here the core configuration observation does not allow x, w, or y to be bottom vertices. Note that j_2 and i_1 need not intersect, if we are considering vertices w and y — the argument continues to hold. The case of Figure 3.9(c) is handled similarly.

There is another class of configurations that we do not illustrate with the figures, where among

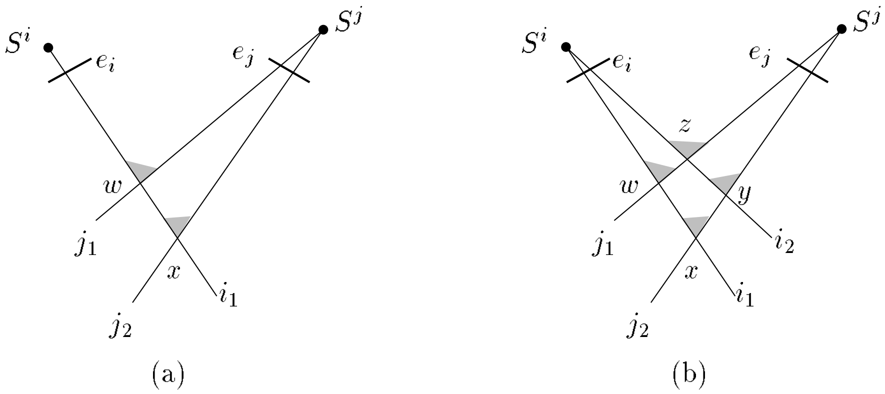

Figure 3.8: One bottom vertex per pair of mirrors.

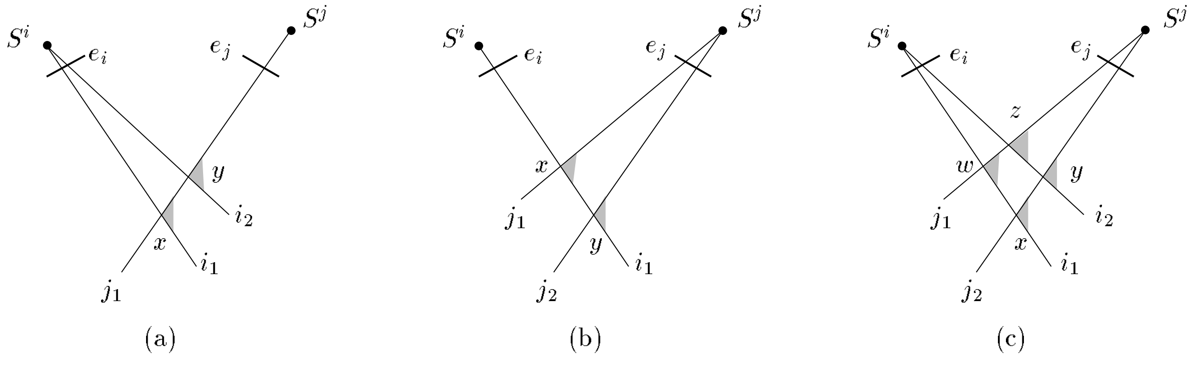

Figure 3.9: One left vertex per pair of mirrors.

the four participating vpath limits there are exactly two intersections. In such a situation, one can also eliminate the possibility of the two vertices being both left, right, or bottom vertices, by inspection. $\square$

At this point we would like to mention that there can be no analogue of Lemma 3.7 for top vertices, as indeed it is possible for the interaction of two mirror visibility polygons to create several top vertices. Refer to Figure 3.10.

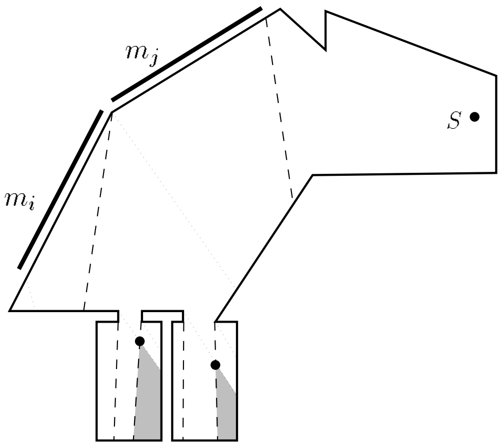

Figure 3.10: Two or more top vertices can be created by the interaction of two mirror visibility polygons. Two blind spots are shaded and their top vertices indicated by dots. The relative boundary of one mirror visibility polygon is drawn dashed, while the second one is outlined in a dotted line.

4 Complexity of $Vs(S)$

Recall the use of thin mirror visibility polygons in Figure 3.4 to create a blind spot. By placing $\Theta(n)$ small mirrors appropriately along the boundary of P we can produce a $Vs(S)$ with $\Theta(n^2)$ components; see Figure 4.1. The details of the mirror M_1 are shown separately in Figure 4.2; mirror M_2 is similar. More precisely, we use collections of mirrors organized into two "supermirrors" which have the property that they split a light beam into many subbeams. Placing two such supermirrors, M_1 and M_2, with $\Theta(n)$ submirrors each, as indicated in the figure produces a grid-like arrangement of beams with $\Theta(n^2)$ blind spots.

Note that rays reflected off edges of type a in Figure 4.2 are unchanged from reflections off the original top edge, and therefore do not intersect the vpath limit i_1 in Figure 4.1. There are no reflections off edges of type b since they are not directly visible to the source. Reflections off edges of type c generate the beams that escape the direct visibility polygon and create one set of "grid lines." Analogous reasoning applies to reflections off the other supermirror. This produces an $\Omega(n^2)$ lower bound on the number of blind spots and on the worst-case complexity of $Vs(S)$.

We bound the combinatorial complexity of $Vs(S)$ by counting its vertices, as the number of

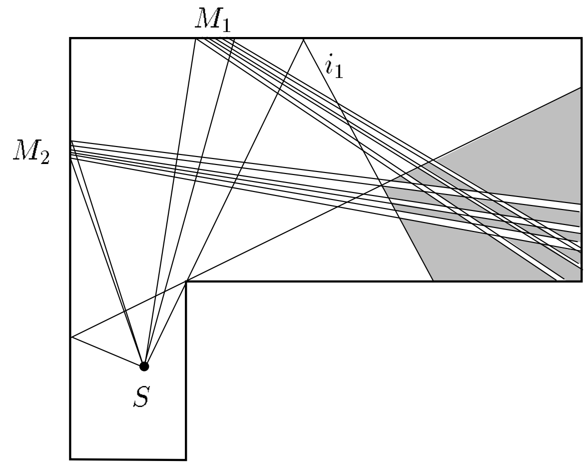

Figure 4.1: Lower bound for $Vs(S)$.

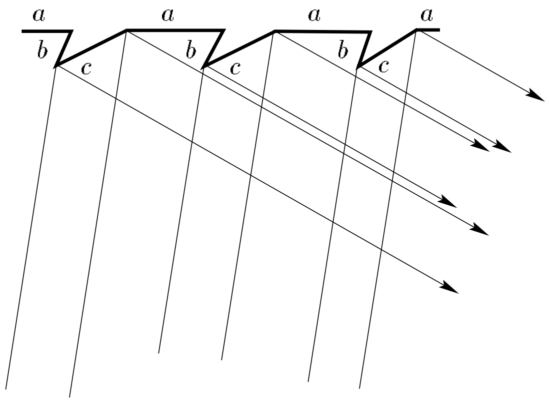

Figure 4.2: Magnified view of supermirror M_1.

its edges is equal to the number of vertices. Vertices of $Vs(S)$ are the vertices of the interior and boundary blind spots and some of the vertices of P. Recall that P has n edges and vertices. S has at most n edges weakly visible from it and thus there are at most n mirror visibility polygons. $V(S)$ is known to have $O(n)$ vertices and edges, and thus windows. A vertex of $Vs(S)$ on the boundary of P is one of the n original vertices of P, an endpoint of a window, or an endpoint of a vpath limit. There are $O(n)$ vertices of the former two types, as each is a vertex of $V(S)$. A vpath limit must emanate from S and then pass through one of the n vertices of P, before, during, or after reflection. If it passes through a vertex of P before or during the reflection, the vertex is visible to S and there is one such vpath limit per vertex. If it passes through a vertex after the reflection, it is fully determined by the mirror and the vertex through which it passes, limiting the number of the second type of vpath limits to no more than n^2. Therefore altogether there are $O(n^2)$ vpath limits and $O(n^2)$ vertices of $Vs(S)$ lying on the boundary of P.

Interior vertices can only be generated by intersections of vpath limits occurring in $P \setminus V(S)$ and must be of type top, bottom, left, or right. Interior vertices occurring on a window must be generated at most in pairs by interaction of $V(S)$ and an edge of the relative boundary of a mirror visibility polygon, by Lemma 3.4. Since each of the at most n mirror visibility polygons has $O(n)$ edges, their number is $O(n^2)$.

Lemmas 3.4 and 3.7 imply that any bottom vertex can be uniquely charged to the pair of mirrors generating it, and that no other bottom vertex can be charged to the same two mirrors. Therefore the number of such vertices is $O(n^2)$. Lemma 3.7 similarly implies that the number of left and right vertices not on a window is $O(n^2)$.

Every top vertex must be incident to a left or a right vertex unless it is the top vertex of a simple boundary blind spot. In the former case, we charge it to its neighboring vertex. The neighbor is charged no more than once. Hence the number of top vertices of non-simple boundary blind spots is $O(n^2)$. A simple boundary blind spot has a top vertex and exactly two boundary vertices created by vpath limits or windows, and we have shown above that there are $O(n^2)$ such boundary vertices. Therefore there can be at most $O(n^2)$ top vertices of simple boundary blind spots. Combining these estimates, we obtain an $O(n^2)$ upper bound on the complexity of $Vs(S)$.

Theorem 4.1 *The combinatorial complexity of $Vs(S)$ is $O(n^2)$, and $\Omega(n^2)$ can be achieved in the worst case.*

A slightly more careful argument yields that the worst-case complexity of $Vs(S)$ is in fact $\Theta(ar + a^2 + n)$, where a is the number of edges of P visible from S and r is the number of reflex vertices of P [6].

Observe that in all of our arguments we have only examined the interaction of $V(S)$ and individual mirror visibility polygons (Lemma 3.4), and the interaction of pairs of mirror visibility polygons (Lemma 3.7). No global argument appears to be necessary. So we have proven a statement stronger than Theorem 4.1. It will be useful in Section 6 where we address algorithmic issues and need to manipulate subsets of mirror visibility polygons.

Observation 4.2 *The combinatorial complexity of the union of $V(S)$ and any $k \geq 1$ mirror visibility polygons is $O(kn)$.*

5 Diffuse Visibility Regions

We now consider the diffuse visibility region $Vd(S)$. We take the same approach as in the specular case and only comment on the differences in analysis — remarkably, there are relatively few differences.

The notion of a mirror remains the same. The (diffuse) mirror visibility polygon $V_i(S)$ of mirror m_i is just the weak visibility polygon of m_i, i.e., the set of points that can see at least one point in the interior of m_i. $Vd(S) = V(S) \cup \bigcup_{i=0}^{a-1} V_i(S)$. The boundary of $Vd(S)$ consists of portions of windows, portions of the relative boundaries of the $V_i(S)$, and (portions of) edges of P. Similar to the case of specular reflection, $V_i(S)$ is cut out of P by (portions of) limiting cases of visibility paths which we call *(diffuse) vpath limits*. The types of segments (besides windows) that arise in the vpath limits are illustrated in Figure 5.1 (see, for example, [10]). For simplicity of presentation, in addition

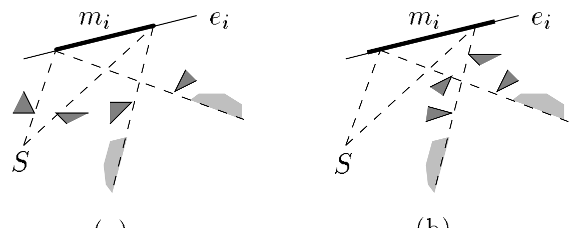

Figure 5.1: Diffuse vpath limit types; darker shading indicates exterior of P.

to the assumption that P does not contain three collinear vertices, we also suppose that the endpoint of a mirror is not collinear with any pair of vertices of P.

The second link of a diffuse vpath limit either (a) emanates from an endpoint of m_i and passes through a single reflex vertex, or (b) emanates from an interior point of m_i and passes through two reflex vertices. In the former case the line containing the second link has the mirror and the reflex vertex on the same side. In the latter case, it separates the two reflex vertices. See Figure 5.1. The *shadow edge* of the diffuse vpath limit starts after the *last* reflex vertex on the path and this vertex determines on which side the shadow lies. We define this vertex to be the *key reflex vertex* for the diffuse vpath limit. Our non-degeneracy assumptions ensure that the above classification of diffuse vpath limits is well defined. If we were to allow two vertices collinear with the endpoint of a mirror, it would be possible for a portion of a vpath limit to have both sides in shadow.

For specular reflections, distinct rays (or vpath limits) reflecting from a common mirror must diverge. However, for diffuse reflections distinct rays or vpath limits reflecting from a given edge may intersect. Analysis of blind spot vertices using diffuse vpath limits will therefore need to consider cases involving intersecting diffuse vpath limits from a given mirror. These are the only new cases that require a change of the arguments we have used for specular reflection.

Lemma 5.1 *If two distinct diffuse vpath limits reflecting from a common mirror intersect, they can not meet their respective key reflex vertices until after that intersection.*

Proof. Let two diffuse vpath limits from S through points y and z of mirror m_i intersect at x as in Figure 5.2.

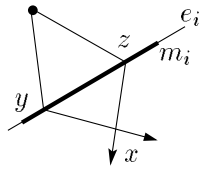

Figure 5.2: Intersecting diffuse vpath limits.

Without loss of generality consider the diffuse vpath limit through y and its link containing x and y. If y is an endpoint of m_i, the corresponding key reflex vertex must occur on the same side of the link as m_i. If it occurred before x it would lie inside $\triangle yxz$ contradicting Fact 3.2. Thus y must be interior to m_i and we attempt to place the key reflex vertex before x. Placing it inside $\triangle yxz$ contradicts Fact 3.2, so it must lie outside. If the key reflex vertex is placed on yx outside $\triangle yxz$, the corresponding first reflex vertex must be placed on the inside, again contradicting Fact 3.2.

Therefore in all cases the key reflex vertices generating the diffuse vpath limits must occur after the intersection at x. ◻

Note that no interior vertex of a blind spot can be formed by two vpath limits reflected off the same edge, since $V_i(S)$, as the region of weak visibility from m_i, is cut out of P by disjoint chords (see, e.g., [10]). Alternatively, it follows from Lemma 5.1 that two vpath limits from the same m_i intersect before the start of their shadow edges, and therefore can not form an interior vertex of any blind spot.

Lemma 5.2 *Blind spots of $Vd(S)$ have no bottom vertices.*

Proof. We argue by contradiction. If x is a bottom vertex of a blind spot it must lie at the intersection of two vpath limits i_1 and j_1, by the reasoning of Lemma 3.4. As already observed, these vpath limits must come from two different edges e_i, e_j of P. Let these vpath limits intersect e_i and e_j at y and z respectively. As x is a bottom vertex of a blind spot, it must be shaded locally as in Figure 5.3. Then there must exist two key reflex vertices as indicated in the figure.

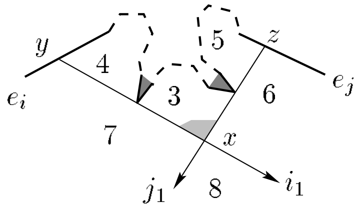

Figure 5.3: No interior blind spots; darker shading indicates exterior of P.

Now consider the placement of the source S. The boundary of P can not cross a vpath limit, or cross itself, so the only possible configuration for it is shown as the dashed line in the figure. S can not be placed in regions 3, 4, or 7 since then z would not be directly visible to S. S can not be placed in regions 3, 5, or 6 since then y would not be directly visible to S. S can not be placed in region 8 since then x would be directly visible from S and could not be a blind spot vertex.

Since there is no valid placement for S, there can be no bottom vertex of a blind spot. $\square$

Corollary 5.3 *$Vd(S)$ has no interior blind spots.*

Proof. Arguing as in Lemma 3.5, if an interior blind spot existed, it would have to be convex and have a bottom vertex. $\square$

Corollary 5.4 *A boundary blind spot of $Vd(S)$ is a closed (except possibly at reflex vertices of P) set bounded by a convex chain and a connected portion of $bd(P)$. It has no bottom vertices and at most one top vertex.*

Proof. Arguing as in Lemma 3.6, a boundary blind spot must be cut out by a connected convex curve. Further following the proof of Lemma 3.6, we conclude that a boundary blind spot can have at most one top vertex, as the presence of two top vertices would imply the existence of a bottom vertex between them, leading to a contradiction. $\square$

As in Lemma 3.4, interaction of an edge of a mirror visibility polygon with $V(S)$ yields at most one interior blind spot vertex, as the edge can only intersect $V(S)$ in a connected segment, by Fact 3.1(b), and one of its endpoints lies on $bd(P)$.

We now show how to count the other types of blind spot vertices in a manner similar to that used for specular reflections. Note however, that for diffuse reflections it is possible for vpath limits reflected from the same edge to intersect, but their key reflex vertices must occur after that intersection.

Lemma 5.5 *Interaction of two mirror visibility polygons can not generate two left blind spot vertices.*

Proof. As in Lemma 3.7 we first consider the situation where two vpath limits from the source point S reflecting off edge e_j are intersected at points a and b by a third vpath limit reflecting off a different edge e_i. We call it a *core configuration.* Consider the four angles adjacent to the segment ab. See Figure 5.4(a). By Fact 3.2 no reflex vertex of P can occur inside the triangular region abz whose boundary is contained in P. In particular, the indicated angles are lit up by rays reflecting off e_j and thus a (b) can not be a blind spot vertex with the blind spot lying in one of these angles near a (resp. b). Similarly, the same reasoning applies when there is no intersection of vpath limits, as it Figure 5.4(b), where quadrilateral $aa'b'b$ is fully lit up. We call this fact *core configuration observation* and use it repeatedly to prove the lemma.

To derive a contradiction, assume that such a pair of vertices exists. As in the proof of Lemma 3.7, we separately consider the case of the two vertices sharing a vpath limit (Figure 5.5) and the case where there are four vpath limits involved (Figure 5.6). In either case, vpath limits from one edge must cross before arriving at the vertices, if they cross at all, by Lemma 5.1. If they do not cross, the argument of Lemma 3.7 can be repeated verbatim here, with the role of the single reflex vertex of the

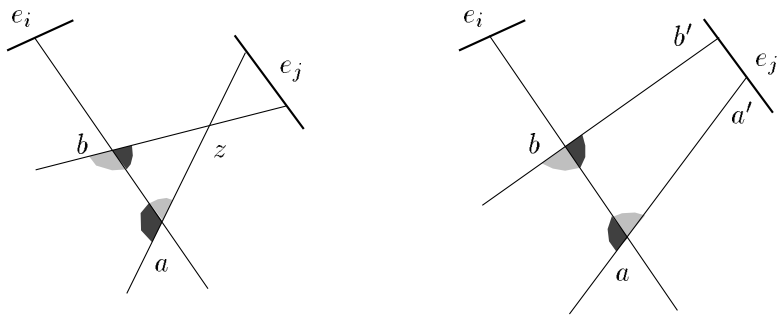

(a) (b)

Figure 5.4: Shading indicates angles not containing incident blind spots.

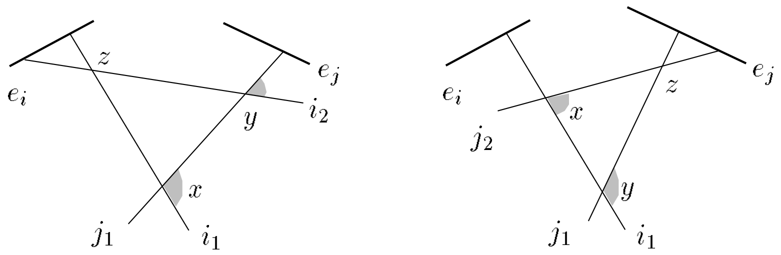

(a) (b)

Figure 5.5: Two left vertices sharing a vpath limit.

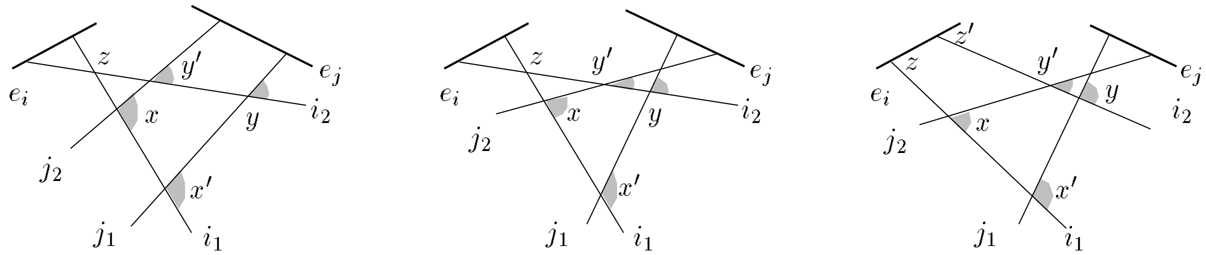

Figure 5.6: Two left vertices not sharing a vpath limit.

proof of the specular case played here by the key reflex vertex. Thus we will proceed to handle only those cases where there is at least one crossing among vpath limits reflecting off the same mirror.

The cases depicted in Figures 5.5(a) and (b) are ruled out immediately by the core configuration observation, as there can be no left blind spot vertex at x. In each subfigure of Figure 5.6, either x, y or x', y' is the pair of candidate left vertices. Here the core configuration observation disallows x (y') as a left blind spot vertex if the first (respectively, second) pair is assumed to be present.

There is another class of configurations that we do not illustrate with figures, where among the four participating vpath limits there are exactly two intersections. In such a situation, one can eliminate the possibility of the two vertices both being left blind spot vertices, by an analogous process. □

Therefore a left blind spot vertex can be associated with a unique pair of mirrors, and no pair is charged more than once, so there are at most a quadratic number of left blind spot vertices. Similarly for right vertices.

Theorem 5.6 *The complexity of $Vd(S)$ is $O(n^2)$ and can be $\Omega(n^2)$ in the worst case.*

PROOF. In Figure 5.7, each of the $\Theta(n)$ pockets generates $\Theta(n)$ boundary blind spots, for a total of

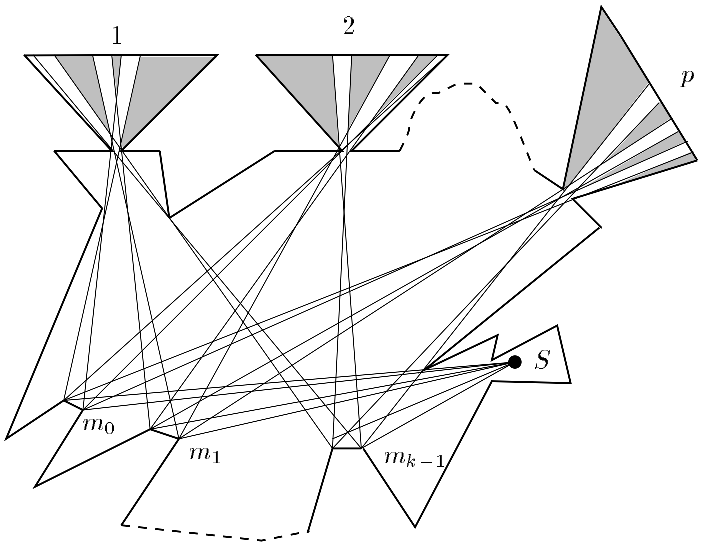

Figure 5.7: Lower bound construction for $Vd(S)$.

$\Omega(n^2)$ blind spots. This give an $\Omega(n^2)$ lower bound on the worst-case complexity of $Vd(S)$.

We now derive an upper bound. We begin by recalling that $Vd(S)$ is the union of $V(S)$ and the at most n mirror visibility polygons. Each of the at most $n+1$ sets is a polygonal set cut out of P by at most n chords [8, 16, 10]. Thus only $O(n^2)$ vertices can occur on the boundary of P. Only $O(n^2)$ vertices can occur on windows of $V(S)$, at most two per edge of the relative boundary of a mirror visibility polygon. $Vd(S)$ is a polygon whose interior vertices can only be left, right, or top vertices of boundary blind spots. By Lemma 5.5, we can only have $O(n^2)$ left and right vertices. Since any boundary blind spot has at most one top vertex by Corollary 5.4 and exactly two vertices on $bd(P)$, there are at most $O(n^2)$ top vertices. Therefore the complexity of $Vd(S)$ is $O(n^2)$. □

As in the case of specular visibility, our analysis proves a stronger statement than the above theorem.

Observation 5.7 *The combinatorial complexity of the union of $V(S)$ and $k \geq 1$ mirror visibility polygons is $O(kn)$.*

The analysis can also be refined to produce an upper bound of the form $O(a^2 + ar + n)$ and a lower bound of the form $\Omega(\min\{a^2, ar\} + n)$ on the worst-case complexity if $Vd(S)$, where a is the number of edges of P weakly visible from S and r is the number of reflex corners of P.

6 Algorithms

In order to compute the visibility polygon from point S in P:

1. Compute $V(S)$.
2. Compute $V_i(S)$, for $i = 0, \ldots, a-1$.
3. Compute $P_i = V(S) \cup V_i(S)$, for $i = 0, \ldots, a-1$.
4. Merge P_i's in pairs, then merge the resulting sets in pairs, etc., obtaining the desired set in $\lceil \log a \rceil$ rounds.

Each of the $a+1$ sets in steps (1) and (2) can be computed in linear time by a standard visibility polygon algorithm, or weak visibility for diffuse mirror visibility polygons [8, 16, 10]. (More precisely, one first computes $V(S)$ by any standard algorithm. At this point, mirrors have been computed. Specular mirror visibility polygons can be computed from the equation $V_i(S) = P \cap V(S^i, P \cup \Delta^i)$. Triangles Δ^i are easy to compute, so one forms a Riemann surface by gluing P to Δ^i along m_i and runs any triangulation-based visibility polygon algorithm on it, such as that of Guibas *et al.* [10]. It is easy to check that their algorithm works as well on a triangulated simply-connected Riemann surface, though we only need it for a very restricted surface, with limited overlap. The diffuse case is easier, as it only requires a computation of a standard weak visibility polygon from a mirror, inside P. The pairwise merges in steps (3) and (4) are accomplished by a sweepline algorithm [4], at the cost of $O(\log n)$ per vertex of the polygons being merged and per vertex of their union. Fact 3.1 easily implies that each P_i has linear complexity, so computing all P_i's takes $O(an \log n)$ time.

The running time of one merge step is $O((u+v)\log v)$, where u is the number of new intersection points of the input segments discovered by the algorithm, and v is the number of edges bounding the regions being merged. In our case, u is exactly the number of vertices of the union of k P_i's discovered by the algorithm that were not vertices of the two smaller unions being merged. Note that using a $O(u + v \log v)$ time algorithm, such as [5], does not improve the running time, as a newly discovered vertex of a generic merge step participates in the next round of merges as an old vertex, and thus incurs a logarithmic cost. Also, observe that merging all sets P_i simultaneously is inefficient, as it is easy to construct cases where the boundaries of sets P_i intersect among themselves in $\Theta(n^3)$ points.

Observations 4.2 and 5.7 guarantee that the union of $k \geq 1$ P_i's has complexity $O(nk)$. On the other hand, v is the output of the two merge algorithms at the previous level of the merge tree, each of which has complexity $O(nk/2)$. Hence the cost of merging the two sets is $(2 \times O(nk/2) + O(nk)) \times \log(nk) = O(nk \log n)$.

Let $T(k)$ be the amount of time it takes to produce the union of k P_i's. As discussed above, sweepline-based merge costs $O(kn \log n)$, so $T(k)$ has to satisfy

$$T(k) = \begin{cases} O(n), \text{ if } k = 1 \\ 2T(k/2) + O(kn \log n), \text{ otherwise.} \end{cases}$$

Recalling that there are a mirrors altogether, we conclude that the entire algorithm runs in time

$$T(a) = O(an \log a \log n).$$

Thus we conclude

Theorem 6.1 *Both* $Vs(S)$ *and* $Vd(S)$ *can be constructed in* $O(an \log a \log n)$ *time, which is* $O(n^2 \log^2 n)$ *in the worst case; here* $a \leq n$ *is the number of edges of* P *weakly visible from* S.

7 Conclusions

We have defined the concepts of a specular visibility region $Vs(S)$ and a diffuse visibility region $Vd(S)$ based on the type of reflection that is allowed along 2-link visibility paths. We have shown that $Vs(S)$ may be non-simple with interior blind spots, while $Vd(S)$ is always simple. We have presented a $\Theta(n^2)$ worst-case bound on the combinatorial complexity of both $Vs(S)$ and $Vd(S)$ and described simple $O(n^2 \log^2 n)$ time algorithms for constructing both sets.

The *art gallery problem* with reflections is to find a set of a minimum number of point light sources positioned at vertices of a polygon that can illuminate the entire polygon. Lee and Lin [15] have shown that this problem is NP-hard for direct visibility. Prasad [22] shows that this problem remains NP-hard for visibility with at most one reflection.

A variant of the art gallery problem is to find the set of all points (if any) in the polygon, called the *kernel*, each of which can illuminate the entire polygon. Extending this notion of kernel, we define K^i, the kernel under i levels of reflection, to be the set of points each of which can illuminate the entire polygon with at most i reflections. This definition applies naturally to both types of reflections, specular or diffuse. It is obvious from the definition that if K^i is non-empty, so is K^ℓ for any $\ell > i$. Hence, the interesting problem is to find the non-empty K^i where i is minimum. A related problem of finding the *link center* of a polygon has been solved in [17, 13, 7]. Link center of a polygon is defined to be the set of points in the polygon, each of which can reach the entire polygon with at most k-link paths, where no point can reach the entire polygon with fewer links. The algorithms of

[17, 13] use the crucial property of the link center that it is the intersection of link-k neighborhoods (i.e., loci of points reachable by a k-link path) of vertices of the given polygon. Unfortunately this is no longer true for visibility with reflection. Let $N = \bigcap_v Vd(v)$, with the intersection taken over all vertices v of polygon P. Clearly $K^1 \subseteq N$. However, we can construct examples showing that $K^1 \neq N$; an analogous statement holds for specular reflection. This poses a difficulty in computing kernels under reflections. The only progress we could make in this regard is to characterize the kernel under single diffuse reflection as follows.

Let $S(e)$ denote the set of all points that are visible directly or indirectly with one reflection from every point of an edge e. We call $S(e)$ the *strong visibility region* of e. We can establish that K^1 under diffuse reflection coincides with the intersection of strong visibility regions of all edges. Details of these results on the art gallery problem and its relatives are reported in [22]. Finally, we observe that the above condition is *not* sufficient to describe K^1 under specular reflection, as it is possible for a light source to illuminate (with at most one specular reflection) the boundary of the polygon without fully illuminating its interior.

In another direction of research we consider multiple levels of reflection. The results are reported in a companion paper [1]. It establishes asymptotically tight bounds on combinatorial complexity of the portion of a simple polygon visible from a single point with at most k specular reflections, for any fixed $k > 0$. Computing these regions efficiently is discussed there as well. Further extensions of multiple specular reflection work to general polygonal, possibly disconnected sources of light, and to presence of obstacles ("holes") are discussed in [6]. Diffuse visibility with multiple reflections is considered in [23].

Acknowledgments

The authors wish to thank Thomas Shermer and Joseph O'Rourke for a helpful discussion and an anonymous referee and Joseph O'Rourke for constructive criticism of an earlier version of this paper.

References

[1] B. Aronov, A. R. Davis, T. K. Dey, S. P. Pal, and D. C. Prasad. Visibility with multiple reflections. *Discrete Comput. Geom.* To appear.

[2] B. Aronov and M. Teichmann. Online maintenance of visibility and shortest-path information. Manuscript, 1994.

[3] D. Avis and G. T. Toussaint. An optimal algorithm for determining the visibility of a polygo n from an edge. *IEEE Trans. Comput.*, C-30:910–1014, 1981.

[4] J. L. Bentley and T. A. Ottmann. Algorithms for reporting and counting geometric intersections. *IEEE Trans. Comput.*, C-28:643–647, 1979.

[5] B. Chazelle and H. Edelsbrunner. An optimal algorithm for intersecting line segments in the plane. *J. ACM*, 39:1–54, 1992.

[6] A. R. Davis. *Visibility with reflections in flat triangulated surfaces.* Ph.D. thesis, Dept. Comput. Inf. Sci., Polytechnic Univ., Brooklyn, NY. In preparation.

[7] 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.

[8] H. ElGindy and D. Avis. A linear algorithm for computing the visibility polygon from a point. *J. Algorithms*, 2:186–197, 1981.

[9] H. Fuchs, Z. M. Kedem, and B. Naylor. On visible surface generation by a priori tree structures. *Comput. Graph.*, 14:124–133, 1980.

[10] 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.

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

[12] A. Horn and F. Valentine. Some properties of l-sets in the plane. *Duke Mathematics Journal*, 16:131–140, 1949.

[13] Y. Ke. An efficient algorithm for link-distance problems. In *Proc. 5th Annu. ACM Sympos. Comput. Geom.*, pages 69–78, 1989.

[14] V. Klee and S. Wagon. *Old and New Unsolved Problems in Plane Geometry and Number Theory.* Mathematical Association of America, 1991.

[15] D. Lee and A. Lin. Computational complexity of art gallery problems. *IEEE Trans. Inform. Theory*, 32:276–282, 1986.

[16] D. T. Lee. Visibility of a simple polygon. *Comput. Vision Graph. Image Process.*, 22:207–221, 1983.

[17] 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. *Discr. Comput. Geom.*, 3:281–293, 1988.

[18] J. Mitchell. Shortest paths and networks. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 24, pages 445–466. CRC Press LLC, 1997.

[19] J. S. B. Mitchell and S. Suri. Geometric algorithms. In M. O. Ball, T. L. Magnanti, C. L. Monma, and G. L. Nemhauser, editors, *Network Models*, Handbook of Operations Research/Management Science, pages 425–479. Elsevier Science, Amsterdam, 1995.

[20] J. O'Rourke. *Art Gallery Theorems and Algorithms.* Oxford University Press, New York, NY, 1987.

[21] J. O'Rourke. Visibility. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, pages 467–479. CRC Press LLC, 1997.

[22] D. C. Prasad. *New Visibility Problems: Combinatorial and Computational Complexities.* Ph.D. thesis, Dept. Comput. Sci. Engg., Indian Institute of Technology, Kharagpur, India, 1997.

[23] D. C. Prasad, S. P. Pal, and T. K. Dey. Visibility with multiple diffuse reflections. *Comput. Geom. Theory Appl.*, 1997. To appear.

[24] J. Reif, J. Tygar, and A. Yoshida. Computability and complexity of ray tracing. *Discr. Comput. Geom.*, 11:265–288, 1994.

[25] T. C. Shermer. Recent results in art galleries. *Proc. IEEE*, 80(9):1384–1399, Sept. 1992.

[26] J. Urrutia. Art gallery and illumination problems. In J.-R. Sack and J. Urrutia, editors, *Handbook on Computational Geometry*. North-Holland, Amsterdam, Netherlands. To appear.