Illuminating Objects with Mirrors (Long Version)

Ferran Hurtado[1] Marc Noy[1] Jean-Marc Robert[2]
Vera Sacristán[1] Steven Skiena[3]

November 5, 1998

1 Introduction

Advanced computer graphics techniques such as ray tracing are being used to visualize complicated three-dimensional surfaces. However, any single image of such a surface can only provide a partial representation of it, since part of the surface is always blocked from view. Multiple images from several different viewpoints can be used to cover the entire surface, however mentally integrating the images to understand of the total surface can be difficult when viewpoints are not correlated in a natural way.

An interesting alternative to using multiple images for visualization is to place synthetic mirrors in the scene with the surface to be visualized, positioned so as to reflect images of the non-visible portions of the surface back to the viewpoint. We get the advantages of multiple viewpoints while maintaining a conception of where these viewpoints come from.

To best realize these advantages, we must carefully position the mirrors in the scene so as to "optimally" view the surface. In this paper, we provide algorithms and geometries providing optimal mirror placements for two natural classes of problems; object-independent and object-dependent mirror placement.

In the object-independent formulation, we seek to design a "camera" with a fixed geometry of mirrors and viewpoint which satisfies certain criteria for all objects in a specific class. The object-independent formulation is well-suited to visualizing scientific data sets, since extensive analysis prior to an initial rendering may be expensive or (as in the case of volumetric data without an explicit geometric model) impossible. In Section 2, we present the design of a camera which guarantees complete visibility for any orientation of any convex object in both 2 and 3 dimensions using the minimum number of mirrors, and which also minimizes the maximum distance between a point in the object and its (virtual) viewpoint. This distance criteria is important because it ensures that the reflected images are as high-resolution as possible.

In addition to these theoretical results on mirror placement, we present ray-traced images showing simulated views from our optimal mirrors configuration. We literally provide a new way to see the world.

In Section 3, we consider problems with a different algorithmic flavor, where we have the freedom to position the mirrors/viewpoint as a function of the object. We seek to find the minimum number of viewpoints sufficient to view the complete surface of the object. In particular, our results include

- An $O(n^3)$ algorithm for positioning a mirror and viewpoint around an n-gon P so as to view the complete surface of P, if possible.
- An $O(n \log n)$ algorithm to determine all placements of pairs of viewpoints sufficient to view the complete surface of P.

[1]Dep. Matemàtica Aplicada II, Universitat Politècnica de Catalunya. (hurtado@ma2.upc.es, noy@ma2.upc.es, vera@ma2.upc.es). Partially supported by Acción CICYT no. TIC-747/93-C02, Acción CICYT no TIC95-0957-E and Projecte UPC PR-9410. The stay of Vera Sacristán at SunySB was supported by 1995BEAI200036.
[2]Dép. d'Informatique et de Mathématique. Université du Québec à Chicoutimi. (jmrobert@ uqac.uquebec.ca).
[3]Computer Science Dept., State University of New York at Stony Brook. (skiena@cs.sunysb. edu).

- An $O(n)$ algorithm to position a viewpoint and mirror around a convex n-gon so as to minimize the maximum viewing distance.

Previous work on visibility with mirrors includes [1] and [5]. These problems are closely related to illumination problems [2, 3], which seek to position lights such that the entire surface of the object is illuminated by at least one light. Das and Narasimhan [4] give a linear time algorithm for finding the shortest line segment to illuminate the interior of a polygon. Klee [9] posed the question of whether all polygons with mirror reflecting edges can be illuminated from a point source, which has since generated a substantial literature [10]. The mirror visibility problems we consider are complicated by the constraints on the position of virtual illumination points imposed by mirror opaqueness and Snell's Law governing light reflection.

2 Object-Independent Problems

In these object-independent problems, we seek a positioning of viewpoint and mirrors to guarantee complete visibility of any orientation of any convex body. These problems are closely related to illumination problems, which seek to position lights such that the entire surface of the object is illuminated by at least one light.

Mirror visibility problems have an additional complication over illumination problems, namely a constraint on the illumination points. There is only one actual illumination point, with k mirrors providing an additional k virtual illumination points. The mirrors are assumed to be opaque but perfectly reflecting, and obey Snell's Law that the angle of incidence equals the angle of reflection for all light rays. Occlusions by the mirrors and the constraints of Snell's Law make it impossible to realize certain sets of viewpoints.

In our model, a scene consists of an eye position, an arrangement of mirrors, and a spherical region which encloses the object to be viewed. Any arrangement of mirrors which completely illuminates the sphere will completely illuminate any convex body placed within it, so without loss of generality such a "camera" suffices for all convex bodies.

2.1 The two dimensional problem

Theorem 2.1 *One viewpoint and two mirrors are necessary and sufficient to illuminate a circle in E^2. Further, these mirrors may be arbitrarily small.*

Proof: Construct the equilateral triangle circumscribed to the circle. The circle can be completely seen from three viewpoints placed in the three vertices. Thus, a viewpoint can be placed in one of the vertices, v; while the other two vertices, p and q, are to be replaced by mirrors.

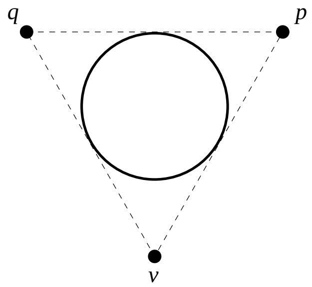

Figure 1: Viewpoints v, p and q together see the entire circle.

We show how to replace p; the replacement of q is analogous. First notice that any point p' lying in the wedge defined by the prolongations of the segment lines vp and qp can see everything that p sees, as shown in Figure 2. Thus, p can be replaced by any of these p'.

We will place the mirror in the line orthogonal to the bisector of pv and pq. It is easy to prove that v sees through the mirror exactly what its reflected image v' sees directly (Figure 3). But because of the position of the mirror, v' lies in the shaded wedge, and therefore sees at least as much as p was able to see. Thus, such a mirror can replace p.

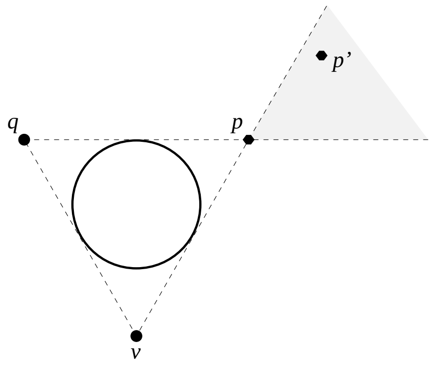

Figure 2: The viewpoint p can be replaced by p'.

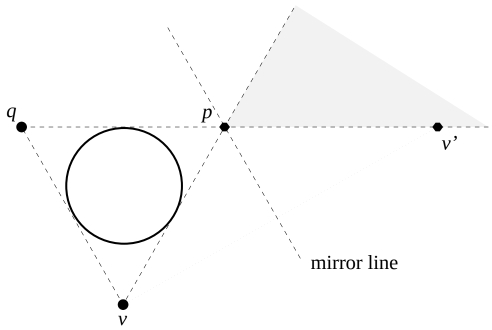

Figure 3: Where to place the mirror line.

Finally, notice that the size of the mirror can be found intersecting the mirror line with the cone determined by the reflected image of the arc seen from p and the viewpoint v, as shown in Figure 4, for there is no need to see with the mirror the entire arc that v' directly sees, but only the arc that p sees.

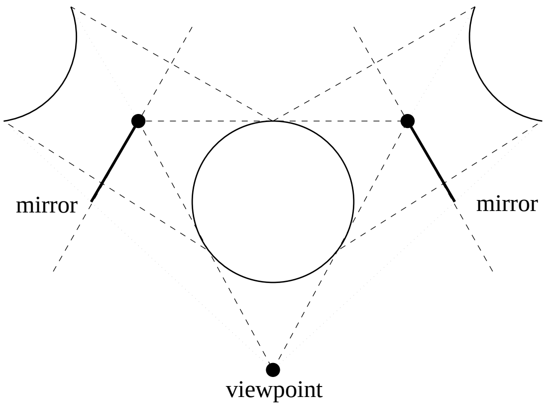

Figure 4: How to find the size of the mirror.

So far, this proves that one viewpoint and two mirrors always suffice to see the entire circle. The necessity comes from the fact that it is impossible to see the complete circle with only two points unless they are in $\pm\infty$.

Finally, observe that the mirrors can be made arbitrarily small by translating v' away from the scene along any line contained in the shaded wedge. See Figure 5. $\square$

Although the mirrors may be made arbitrarily small, for visualization applications we seek mirrors which provide as informative a representation as possible. Specifically, we seek to maximize the resolution of detail we see, and so minimize the maximum distance from the eyepoint to (1) the visible portion of the

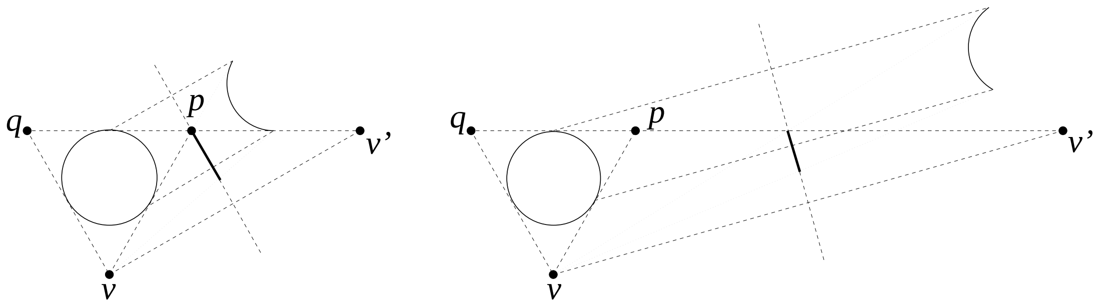

Figure 5: How to decrease the size of the mirrors.

object or (2) the image of its non-visible portions in a mirror. In fact, a careful analysis of the previous construction shows that it is optimal.

Theorem 2.2 *The construction of Theorem 2.1 minimizes the maximum distance between the eyepoint and the image over all arrangements of two covering mirrors.*

Proof: Take any configuration of a viewpoint and two mirrors. The three arcs of the circle that the viewpoint sees (one of them directly, the other two through the mirrors) must overlap in order to cover the whole circle. This proves that any optimal configuration must be a triangle containing the circle. Consider the two arcs reflected in the mirrors. Any tangent to the circle that passes through one of the overlapping points gives a possible location for the mirrors that makes the maximum distance smaller (first step of Figure 6). It should be obvious that a circumscribed one will always be better, as shown in the second step of Figure 6. The next step is to prove that the optimal circumscribed triangle must be isosceles. The reason is that, for a fixed viewpoint, the maximum distance is minimized when the opposite edge of the triangle is perpendicular to the line joining the viewpoint and the center of the circle (third step of Figure 6). For all these triangles, the maximum distance from the viewpoint to the circle

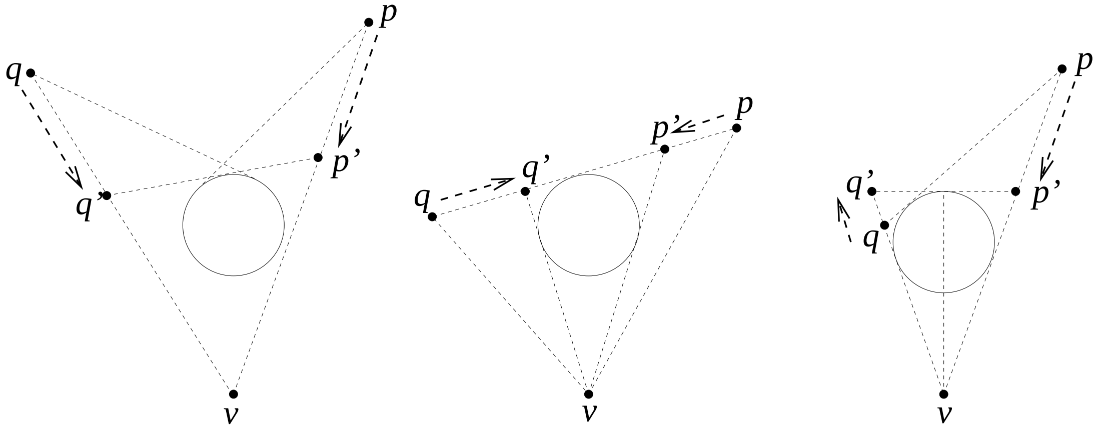

Figure 6: How to decrease the maximum distance: 3 steps.

is equal to the semiperimeter of the triangle. Among all of them, the equilateral triangle minimizes the perimeter and hence is the better one for our purposes. □

2.2 The three dimensional problem

The three dimensional problem is not completely analogous to the two dimensional one. In order to see the entire sphere, the four viewpoints must be placed in the vertices of a tetrahedron containing the sphere, but now the region obstructed by the sphere is not a tetrahedron, but a cone. In the two dimensional case, the viewpoints tetrahedron and the obstruction cone are the same circumscribed triangle. Here, they are not: the viewpoints must lie in the vertices of the tetrahedron but the mirrors can be closer: they only need to lie outside the obstruction cone.

Theorem 2.3 *One viewpoint and three mirrors are necessary and sufficient to illuminate a sphere in E^3. Further, these mirrors may be arbitrarily small.*

Proof: Analogous to the two dimensional construction, we start placing four viewpoints so that they can together see the entire sphere. As the caps they see must overlap (each three of them must share at least one point), the viewpoints must be at the vertices of a tetrahedron containing the sphere. A circumscribed one will always suffice (Figure 7).

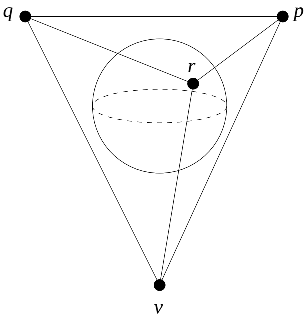

Figure 7: Placement for the four viewpoints.

Fix one of the vertices of the tetrahedron as the viewpoint v, and call the other three vertices p, q and r. We consider one of them, say p. As it happened in two dimensions, all the points p' lying in the region determined by the prolongation of the faces of the tetrahedron converging in p can see at least the same cap of the sphere that p sees, and thus can replace p. We show how to replace p, or one of these p', with a mirror.

Here start to appear the differences with the two dimensional case: observe that the "shadow" generated by the sphere when illuminated from the viewpoint v is a cone. Concretely, the points whose view is obstructed by the sphere lie in a cone circumscribed to the sphere having apex v. In two dimensions, this cone and the tetrahedron correspond to the same object: a circumscribed triangle. We will place the mirrors hanging from one of these cones. Concretely, the bounded cone circumscribed to the sphere, with circular basis tangent to the sphere and aperture of 60° (see Figure 8).

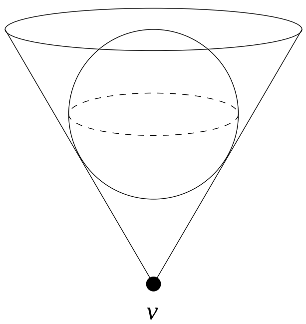

Figure 8: The obstruction cone.

Figure 9 presents a planar section of the solution, in which v is the viewpoint, m is the placement for the mirror plane (a point in the basis of the cone determined by v and the sphere), p is the replaced viewpoint (a vertex of the tetrahedron circumscribed to the sphere having a vertex in point v and an equilateral face opposite to v), v' is the reflected viewpoint and n is the point on the sphere opposite to v. The mirror is to be placed in the plane through m perpendicular to the bisector of mn and mv.

Remember that there is no need to see from v exactly what v' sees, but only what p sees. This means that the shape of the needed mirror can be found by intersecting the mirror plane with the cone determined by the viewpoint and the reflection of the cap seen by p. As the planar section of the sphere determining the cap is not parallel to the mirror, the intersection is an ellipse. Its minor axis lies in the planar section of Figure 9, and the manor one is perpendicular to it.

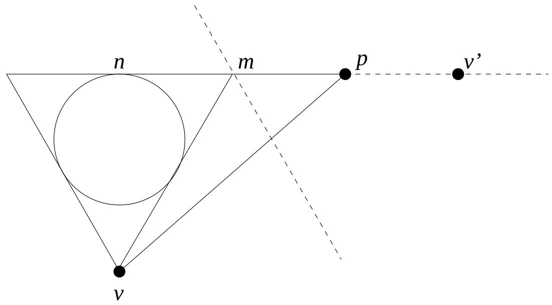

Figure 9: Planar section of the solution.

In Figure 10 you can find a picture of the situation, in where the viewpoints and the mirrors replacing them can be seen.

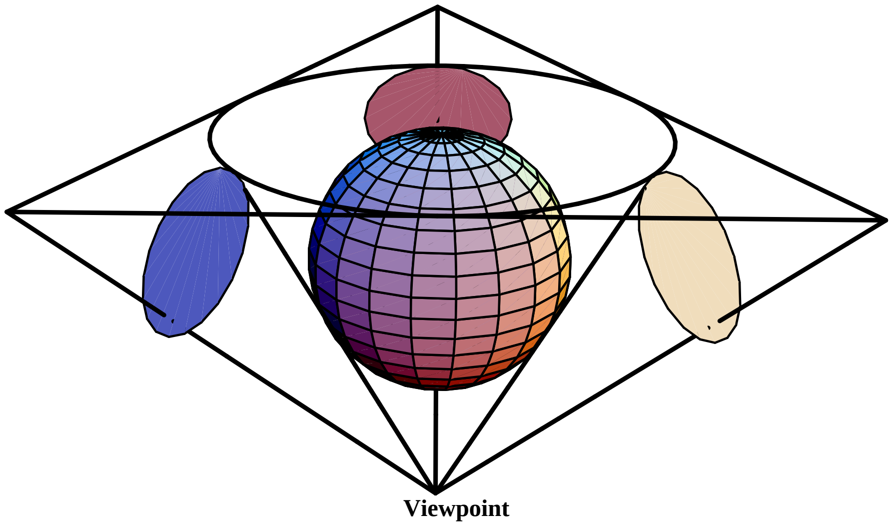

Figure 10: The tetrahedron circumscribed to the sphere and the mirrors replacing its vertices.

As in the two dimensional case, this construction can be modified in order to make the mirrors arbitrarily small, by moving away each mirror from the cone, i.e. by translating away the reflected viewpoint in the unbounded region of the space of all the points that see at least the same cap of the sphere that p sees.

This proves that a viewpoint with three mirrors suffice to see the entire sphere. The necessity comes from the fact that a sphere can not be seen with less than four viewpoints, unless you place them at infinity (in that case, only two are needed). $\square$

Theorem 2.4 *The construction of Theorem 2.3 minimizes the maximum distance between the viewpoint and the image over all arrangements of three covering mirrors.*

Proof: The distance to be minimized is $vm + mn$ (refer to Figure 9), which is again the semiperimeter of the equilateral triangle circumscribed to a circle. The two dimensional argument holds. $\square$

The calculations for an origin-centered unit sphere give the optimal viewpoint at distance 2 from the origin, say $v = (0, 0, -2)$. Thus, the circumscribed cone with apex v has radius $\sqrt{3}$, and $m = (\sqrt{3}, 0, 1)$. The tetrahedron circumscribed to the sphere having v as one of its vertices and an equilateral face opposite to it has the other three vertices at distance $2\sqrt{3}$ from n. Thus, $p = (2\sqrt{3}, 0, 1)$. Finally, the mirror being on the plane through m perpendicular to the bisector of mn and mv, the reflected viewpoint is $v' = (3\sqrt{3}, 0, 1)$.

As said, the mirror needed is elliptic. In this example, the values of its axis are $\frac{16\sqrt{3}}{43}$ and $\frac{4\sqrt{72303}}{1763}$, that is, approximately 0.6445 and 0.6101.

Figures 11, 12, 13 and 14 show simulations of spherical scenes with the three mirrors, ray-traced with rayshade, a ray-tracer created at Princeton University by C. Kolb and successively updated (http://www-graphics.stanford.edu/~cek/rayshade).

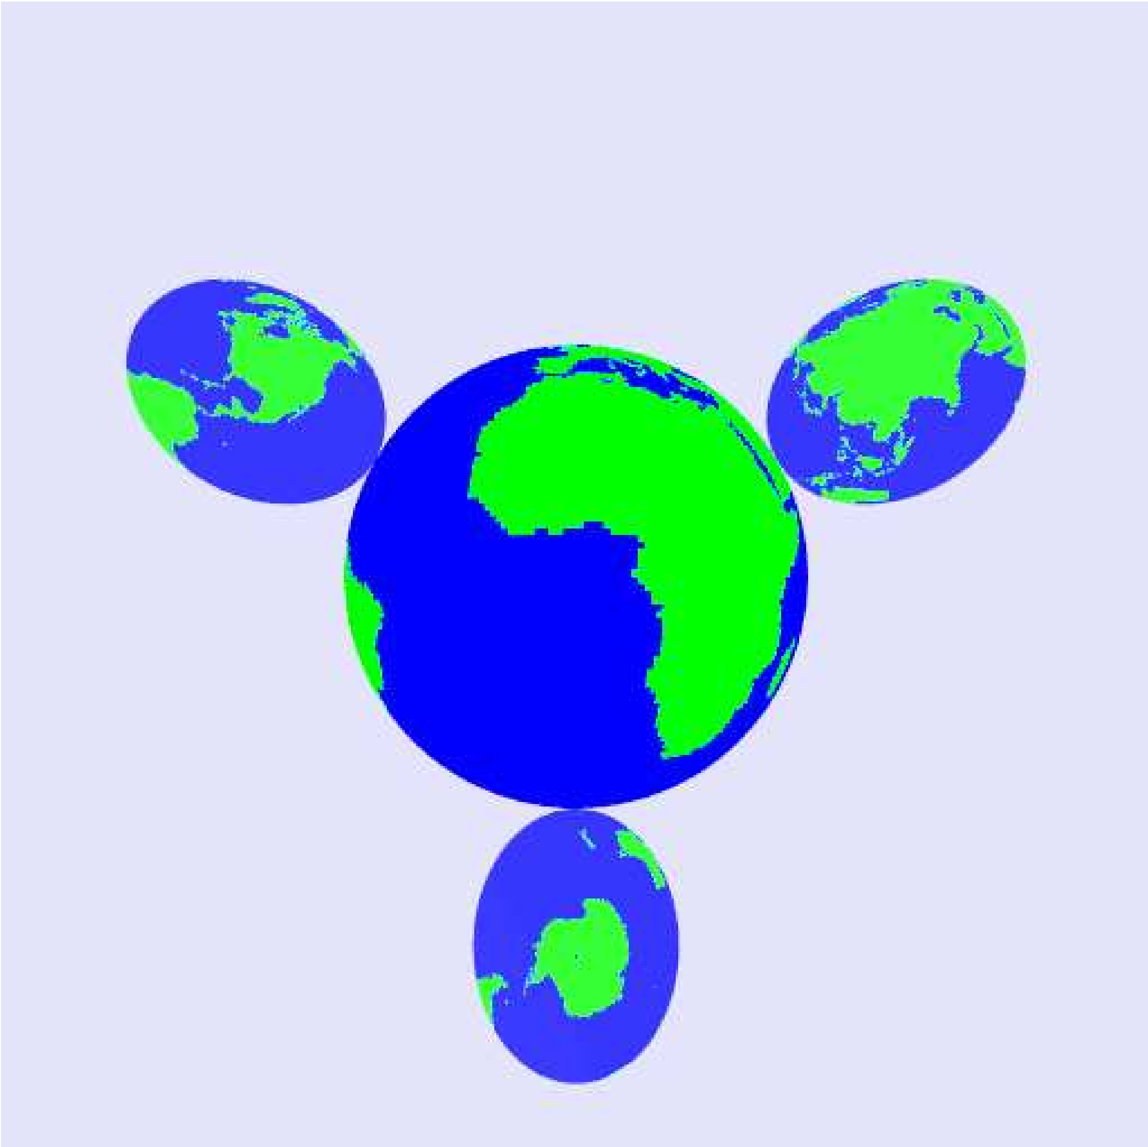

Figure 11: Visualizing the Earth using mirrors. In this image, all shadows and light reflections have been eliminated. The result is a very "flat" image, like a map. The central disk shows the part of the earth that can be directly seen from the viewpoint, while the three elliptic disks are the mirrors that show the inverted images of three caps. Notice that the Earth is completely seen.

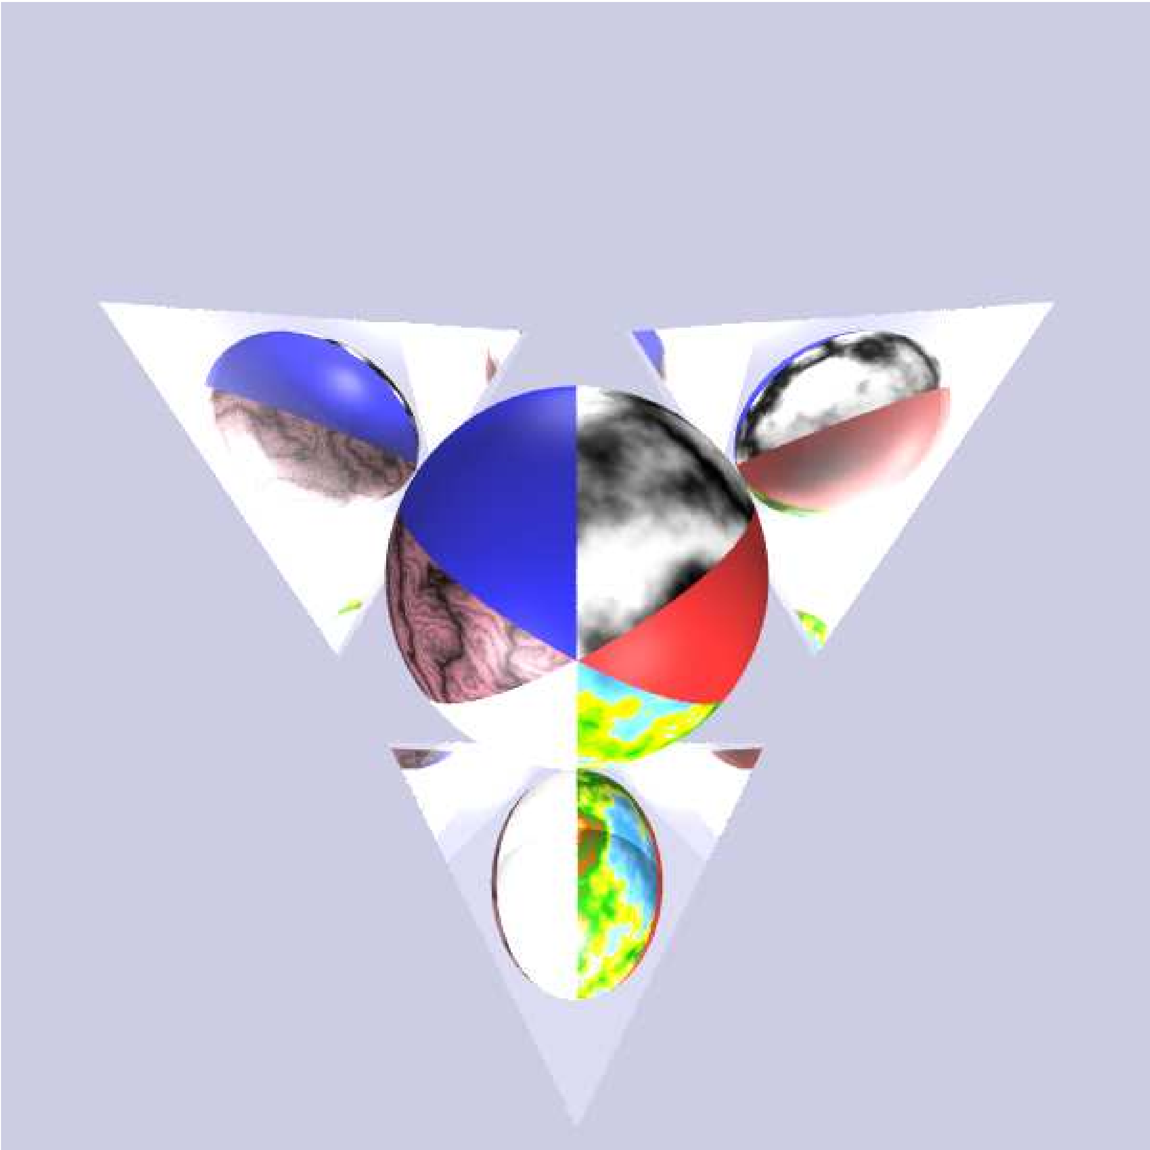

Figure 12: Visualizing a beach balloon using mirrors. In this image, the mirrors have been made greater that necessary, so that the scene can be better understood. The small images that appear in the angles of the mirrors are reflections from a mirror to another.

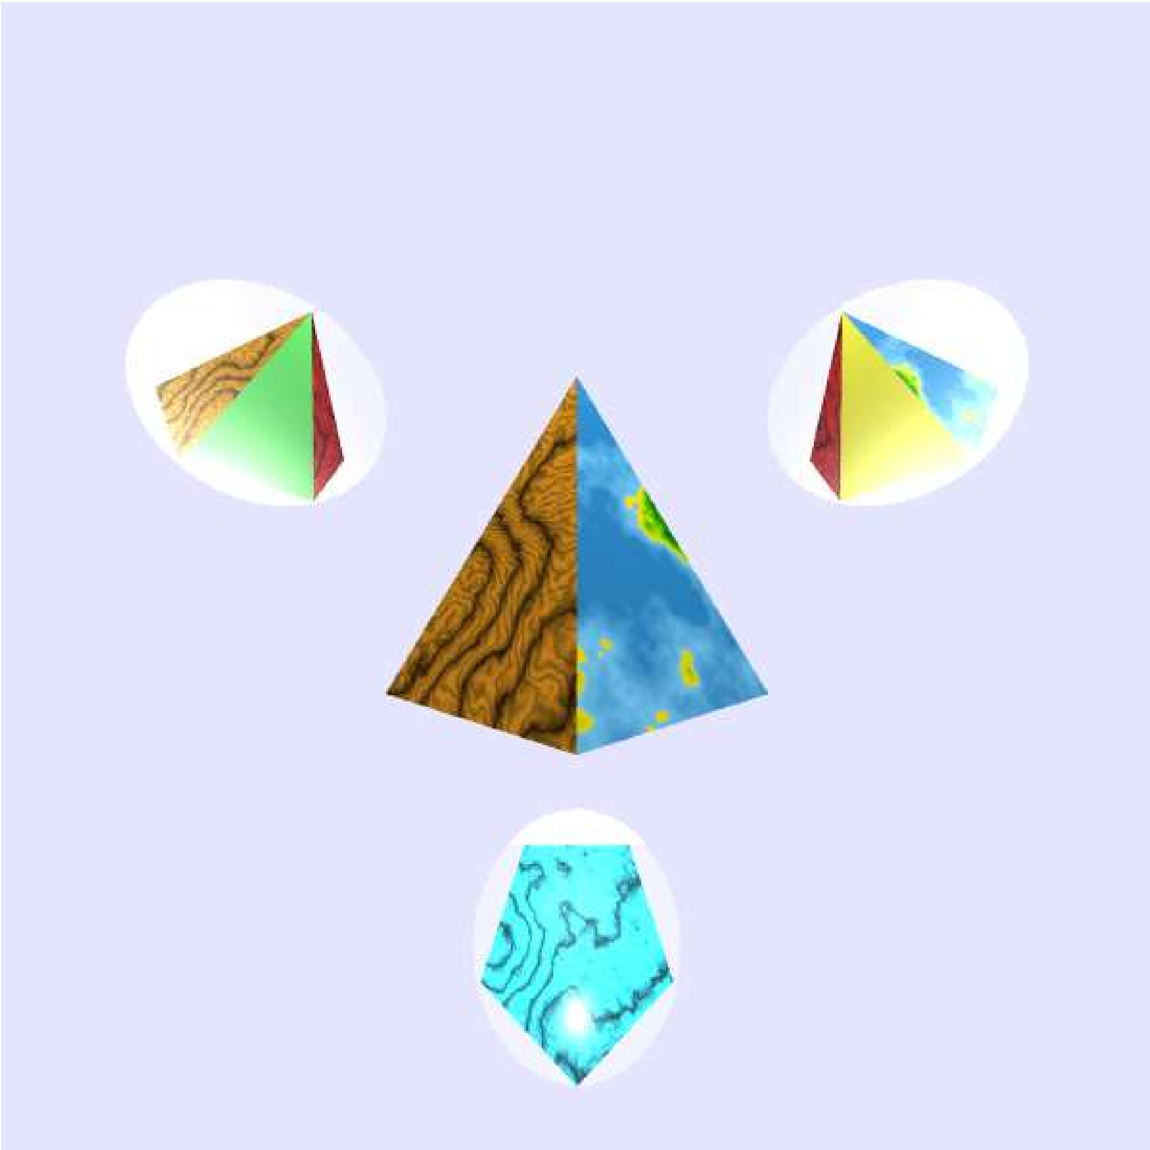

Figure 13: Visualizing a convex body using mirrors. As said, our "camera" allows to see any convex object positioned in the interior of the sphere in a complete way, using the three mirrors.

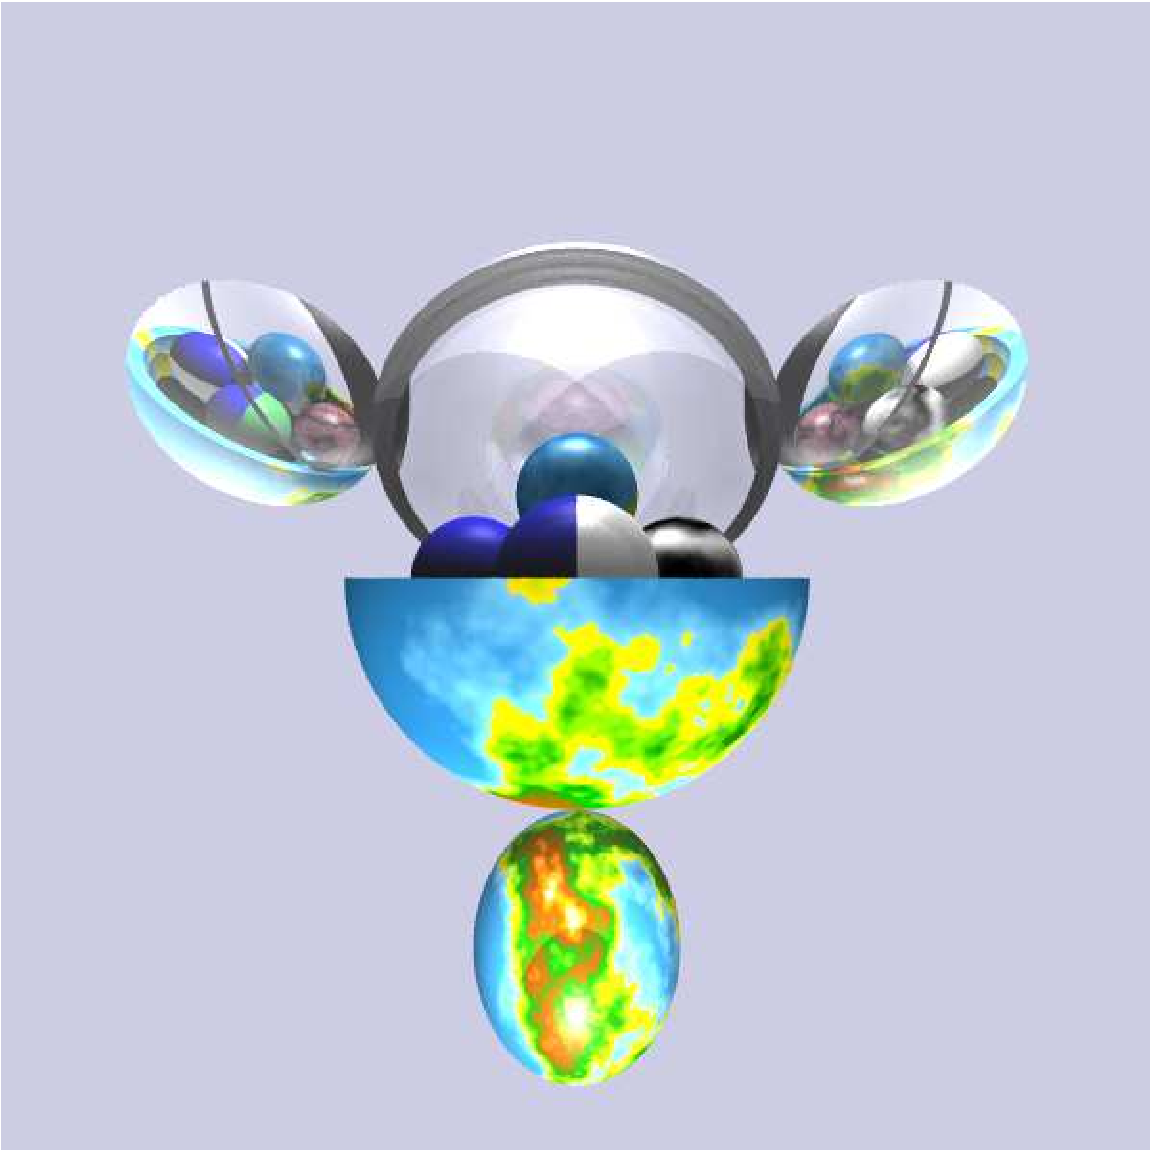

Figure 14: An image to show the three different visions through the three mirrors. Even if the object to be seen is not convex, we think that this image shows the potentiality of our "camera": it allows to clearly see and understand the scene in a very complete way. (The cover is not completely seen because it is placed outside the "camera".)

3 Object-Dependent Problems

In these object-dependent problems we have the freedom to position the viewpoint and the mirrors as a function of the object. As in the object-independent problem, the question of whether or not an object can be seen from a viewpoint with a given number k of mirrors is related to the question of whether it can be seen from $k+1$ viewpoints: any object that can be seen with a viewpoint and k mirrors can also be seen from the viewpoint and its k reflected (view)points.

In Section 3.1 we proof that convex polygons can always be seen from a viewpoint with a mirror, and we solve the associate optimization problem. When dealing with non-convex polygons, some problems can arise, because it is not always possible to see a non-convex polygon with two viewpoints and, even when it is, some obstruction problems can make it impossible to replace one of the viewpoints with a mirror. In Section 3.2 we study the problem of seeing a simple polygon from two viewpoints, and then replace one of the viewpoints with a mirror. Finally, in Section 3.3 we study the three-or-more viewpoints problem, which turns out to be a different one.

3.1 Seeing a convex polygon with one viewpoint and one mirror

We seek to position a viewpoint and a single mirror that allow us to entirely see a given convex polygon. As said, the first step is to position two viewpoints, then replace one of them by a mirror.

Lemma 3.1 *Any convex n-gon can be entirely seen from 2 viewpoints. One pair of such viewpoints can be found in $O(\log n)$ time. All pairs of regions of such viewpoints can be found in $O(n)$ time.*

Proof: It is always possible to find two parallel supporting lines that touch the convex polygon in only one point each (Figure 15). The two tangent points decompose the boundary of the polygon into two chains.

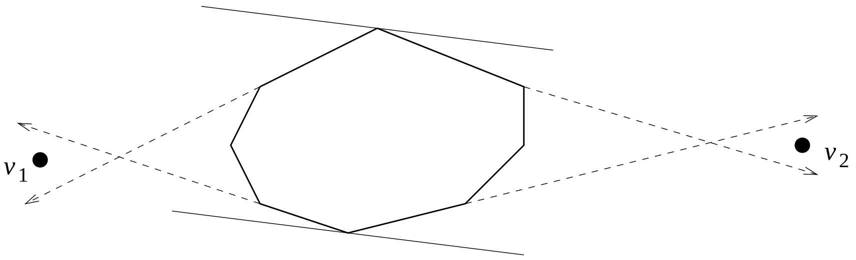

Figure 15: Two viewpoints can together see a convex polygon.

The intersection of the halfplanes determined by the edges belonging to one of these chains is certainly non empty. Actually, because of the convexity, it is just the intersection of two halfplanes, corresponding to the first and the last edges of the chain. Any point in the interior of this intersection sees the entire chain. So, two points, one in each intersection, see the entire polygon. Given a convex polygon, such two points can be found in $O(\log n)$ time, n being the number of edges of the polygon. In fact, for a given direction of the supporting lines, the two regions of possible viewpoints can be found in $O(\log n)$ time.

On the other hand, observe that if two viewpoints, v and v', can see the entire convex polygon P, then ∂P can be decomposed into two subchains, C and C', such that v sees C, v' sees C', and the endpoints of C and C' form an antipodal pair of P. As all antipodal pairs can be found in linear time, all pairs of viewpoint regions can be found in linear time. □

Theorem 3.2 *Any convex polygon can be entirely seen from one viewpoint with one mirror.*

Proof: We start with two regions, one for each viewpoint, and try to replace one of the viewpoints with a mirror. In other words, we want to find a viewpoint in one of the two regions and a position for a mirror such that the reflected viewpoint lies in the other region and no obstruction problems arise.

Let us call the two endpoints od the two chains n and s, the lines determining the left visibility region l_n and l_s, and the lines determining the right visibility region r_n and r_s, as in Figure 16.

Place a viewpoint v in the wedge determined by l_n and l_s, and trace the "southern" supporting line obs_s from v to the polygon. Any view line from v seeking to see n through a mirror must lie below the obstruction line obs_s and then reflect in a mirror above line r_s. So, taking any point m in the wedge determined by obs_s and r_s (notice that v can always be placed such that obs_s and r_s intersect), we can place a mirror through m in the line orthogonal to the bisector of vm and nm, as in Figure‘16. This way, we can be sure that (1) the reflected viewpoint lies in the adequate visibility region and (2) there is no obstruction.

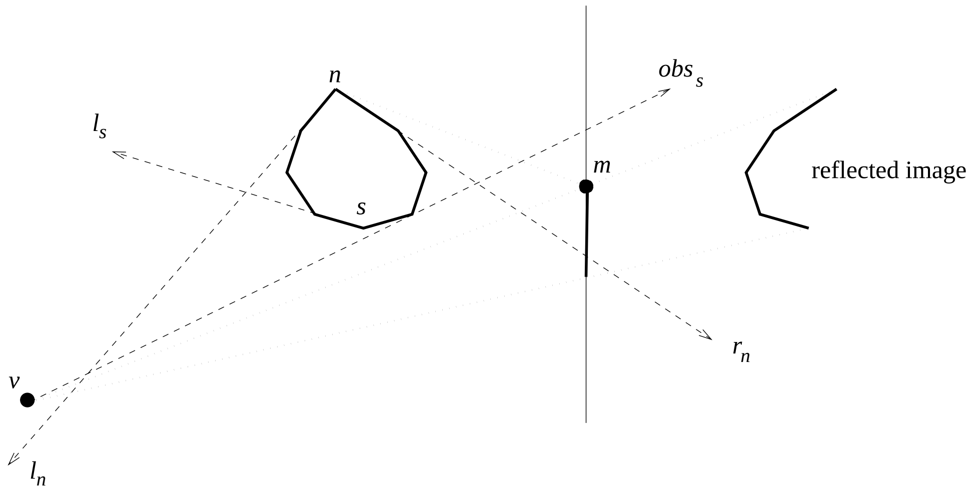

Figure 16: Placing a viewpoint and a mirror.

This proves that any convex polygon can be seen with an adequate viewpoint and a mirror. Figure 16 also shows the minimum required size of the mirror for a given v and m. □

Theorem 3.3 *Given a convex n-gon, in $O(n)$ time it is possible to place a viewpoint and a mirror such to minimize the maximum distance between the viewpoint and the image over all possible placements.*

Proof: The algorithm considers each antipodal pair (n, s). As said, any feasible configuration of a viewpoint and a mirror must place the viewpoint in one of the visibility regions and the reflected viewpoint in the other, such that the view line from the viewpoint v to the opposite chain must reflect in the mirror avoiding the obstruction. In Figure 16 this was done by placing the point m in the wedge determined by obs_s and r_n. There are three other possible constructions, which are analogous to this one: (2) place v in the left region and m in the wedge determined by obs_n and r_s; (3) place v in the right region and m in the wedge determined by obs_n and l_s; (4) place v in the right region and m in the wedge determined by obs_s and l_n.

Consider again the construction of Figure 16, where v lies in the left region and m is in the wedge determined by obs_s and r_n. In order to minimize the maximum distance from v to the points of the polygon, it should be obvious that v must be placed in l_n and m in the apex of the wedge, that is, in the intersection of obs_s and r_n. While moving v along l_n, two things can happen: (1) obs_s rotates around its supporting vertex (and hence m moves along r_n) or (2) obs_s changes its supporting point from a vertex of the polygon to the next one (which, in such case, must also be an antipodal pair of n). Between two events of type (2), it is possible to locate the position of m that minimizes the distance $vm + mn$ in constant time. So, the viewpoint v that minimizes the distance $vm + mn$ over all possible points in l_n can then be found in $O(k_i)$ time, where k_i is the number of antipodal pairs of vertex $v_i = n$. Hence, the entire algorithm runs in time $T(n) = \sum_{i=1}^{\#\text{ant.pairs}} O(k_i) = O(n)$.

The only problem comes from the fact that the maximum distance from a viewpoint v to the polygon is not always achieved in the “north” pole n, which means that the distance to be minimized is not always $vm + mn$. Because of the convexity, the maximum distance from v to the polygon can be achieved either in the “north” pole n, either in “south” pole s and no other possibility is left. This means that the distance to be minimized can be either $vm + mn$, or $vm' + m's$ (see Figure 17).

A simple modification of the previous algorithm solves this inconvenience. Take the output of the algorithm we have so far: it gives the position of the viewpoint v that minimizes the distance $vm + mn$. If $vm' + m'n$ is smaller than $vm + mn$, the problem is solved. Otherwise, move v towards the polygon:

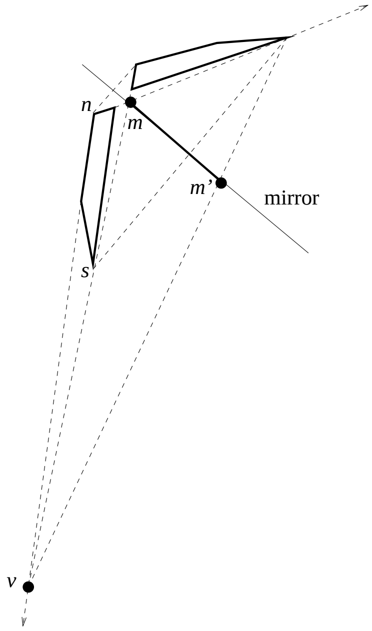

Figure 17: The maximum distance can be to the "south" pole.

this will increase the distance to n and decrease the distance to s. It is easy to prove that the position that makes the two distances equal is the optimal one. □

3.2 Seeing a simple polygon with one viewpoint and one mirror

In the non-convex case, things become more complicated. First of all, some polygons can not be seen with two viewpoints, as shown in Figure 18.

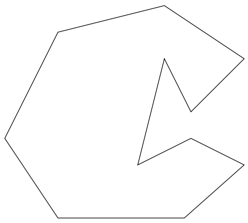

Figure 18: This polygon cannot be seen with two viewpoints.

But even when a polygon can be seen with two viewpoints, it is not always true that it can be seen with a viewpoint and a mirror, as shown in Figure 19.

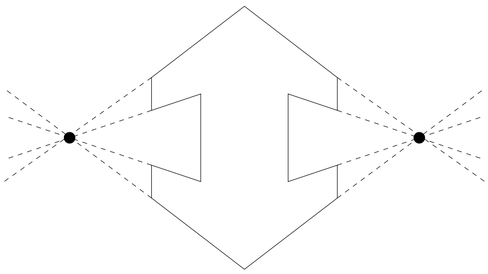

Figure 19: This polygon cannot be seen with a viewpoint and a mirror.

Nevertheless, we start studying whether or not a polygon can be seen with two viewpoints. Later, we will show how to find if one of the viewpoints can be replaced by a mirror.

First of all, remember that if a polygon can be entirely seen from two viewpoints, then each of them must be responsible for seeing a connected subchain of its boundary, such that there must exist a pair of parallel supporting lines passing through the chains endpoints, as can be seen in Figure 20.

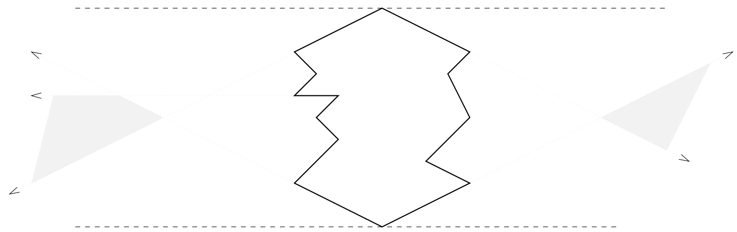

Figure 20: A polygon that can be seen from two viewpoints.

So, the first thing to be done is to understand how the visibility of such polygonal chains behaves. Let C be one side of a simple and connected polygonal chain of n edges. We are interested in determining whether or not there exists a point in the plane that can see all the points of C. This problem is analogous to deciding whether or not a simple polygon is star-shaped, but here is posed for an open polygonal chain. For the case of simple polygons, the decision problem can be solved in linear time by linear programming, because it is a matter of deciding whether or not the kernel of the polygon (the intersection of all the halfplanes determined by its edges) is empty. It is not the same for a polygonal chain.

Fact 3.4 *A point can lie in the intersection of all the halfplanes determined by the edges of a polygonal chain without seeing the whole chain.*

Proof: For a proof, see Figure 21, where none of the points lying in the intersection of all the halfplanes determined by the edges of the chain (the shaded region) sees the whole chain. □

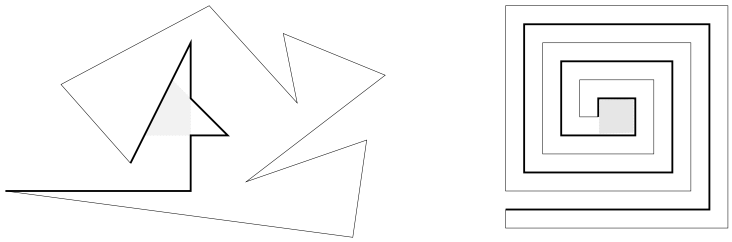

Figure 21: The visibility kernel of a polygonal chain is not the intersection of the halfplanes determined by its edges.

But in the case we are dealing with, the chains have the property that their endpoints always form an antipodal pain, and hence lie in the convex hull of the polygon.

Lemma 3.5 *Let C be a connected subchain of a polygon P, such that the two endpoints of C belong to the convex hull of P. Then any point lying in the intersection of the halfplanes determined by the edges of C sees the entire chain C.*

Proof: First, decompose C into subchains such that the endpoints of each subchain are consecutive vertices of the convex hull of the polygon, $CH(P)$. We will prove that the property holds for each such subchain. The result for the whole chain C will then be trivial.

Consider a chain C whose endpoints x and y are consecutive vertices of $CH(P)$. Let p be a point in the intersection of all the halfplanes determined by the edges of C, $p \in \bigcap H_i$. We want to prove that p sees all the points of C.

If p lies inside the simple polygon $\overline{C}$ obtained by adding the edge $\overline{xy}$ to C, then p belongs to the kernel of $\overline{C}$ and hence sees all C. The fact that $\overline{C}$ is simple and that p must belong to the appropriate

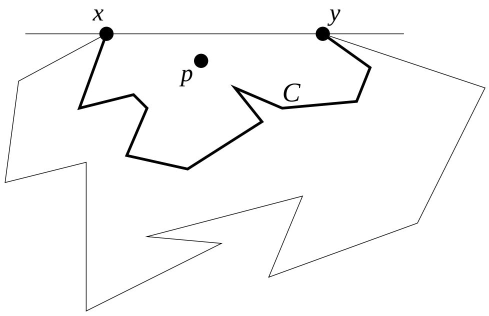

Figure 22: The entire polygon lies this side of $\overline{xy}$.

halfplane determined by $\overline{xy}$ both come from the assumption that x and y are consecutive in $CH(P)$ (see Figure 22).

If p does not lie inside $\overline{C}$, then we can remove from $\overline{C}$ an ear with tip $\neq x, y$ and proceed recursively. Such an ear always exists (as long as $\overline{C}$ has as least 4 edges) because of the two non overlapping ears theorem [11]. When removing an ear, edges a, b of C are replaced by a new one c (refer to Figure 23). Before the removal, $p \in H_a \cap H_b \cap (\bigcap H_i)$. After the removal, $p \in H_c \cap (\bigcap H_i)$. Otherwise, p would have

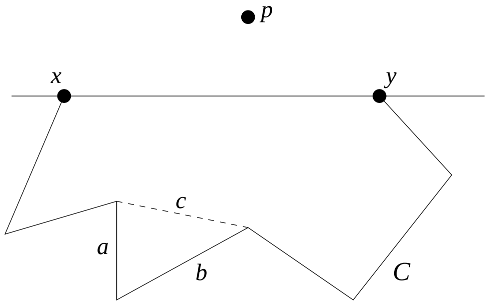

Figure 23: Proceeding recursively.

been interior to the triangle abc, but p is exterior to $\overline{C}$. So, by induction, p sees the new chain, that is, p sees all the edges of the old chain that have not been removed and the newly introduced edge c. But then p must also see a and b, because for all $q \in a \cup b$, the segment line $\overline{pq}$ must intersect c (remember that p lies in the appropriate halfplanes). As p sees all the points of c, p must also see q (Figure 24). The

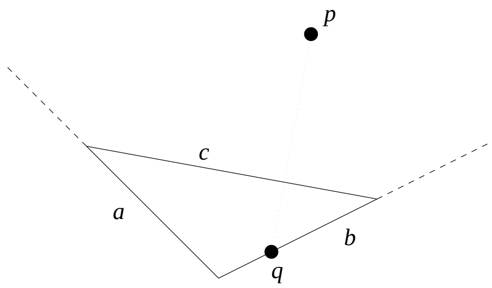

Figure 24: If p sees c, then it must see a and b.

induction goes on until no more ears can be removed, that is, the pocket $\overline{C}$ reduces to a triangle. In this case, the proof is immediate.

So far, we have proved that any point lying in the intersection of the halfplanes sees the entire chain. It should be quite obvious that any point that sees the whole chain must lie in the intersection of the halfplanes. □

Theorem 3.6 *Let P be a polygon with n edges. The existence of two viewpoints that can together see P*

can be determined in $O(n \log n)$ time. Further: all pairs of regions of such pairs of points can be computed in $O(n \log n)$ time.

Proof: As we said yet, if a polygon can be entirely seen from two viewpoints, then each of them must be able to see a connected subchain of its boundary, such that the endpoints of the two chains must form an antipodal pair of vertices of the polygon (remember Figure 20).

So, in order to decide whether or not a polygon can be seen from two viewpoints, the first step is to compute all antipodal pairs of vertices. Then, start with one pair. As it breaks the polygon into two chains, whose endpoints belong to the convex hull of the polygon, Lemma 3.5 applies, and the feasibility region for the viewpoints of each chain can be found as the intersection of the halfplanes determined by its edges. Now, rotate the direction of the supporting lines, and actualize the two visibility regions. After a complete rotation of 180°, you will know if there is any possibility of seeing the entire polygon from two viewpoints, and you will also have, if any, the information about the regions where the viewpoints can lie. This can be done in $O(n \log n)$ time by an off line maintenance of the intersection of the halfplanes [8]. $\square$

Given a polygon that can be simultaneously seen from two viewpoints, the algorithm that decides whether or not one of the viewpoints can be simulated with a mirror and locates a possible position for it is based in the following observation (refer to Figure 26). If a polygon can be seen from a viewpoint and a mirror, then the viewpoint can be translated towards the polygon, along the line to the reflected viewpoint, as long as it remains in its feasibility region, while the reflected viewpoint can be rotated away from the polygon, centering the rotation in the viewpoint. This way, the following result is obtained:

Lemma 3.7 *If a polygon can be seen from a viewpoint with a mirror, then the mirror can always be placed in such a way that both the viewpoint and its reflected image lie in the boundary of their feasibility regions.*

Proof: Suppose that the supporting lines decompose the boundary of the polygon in a left and a right chains. The left chain is visible from the left feasibility region and the right chain is visible from the right feasibility region. Suppose also that it is possible to see the polygon from a viewpoint in the left region and a mirror in the appropriate position, so that both the viewpoint v and its reflection v' lie below the polygon, as in Figure 25.

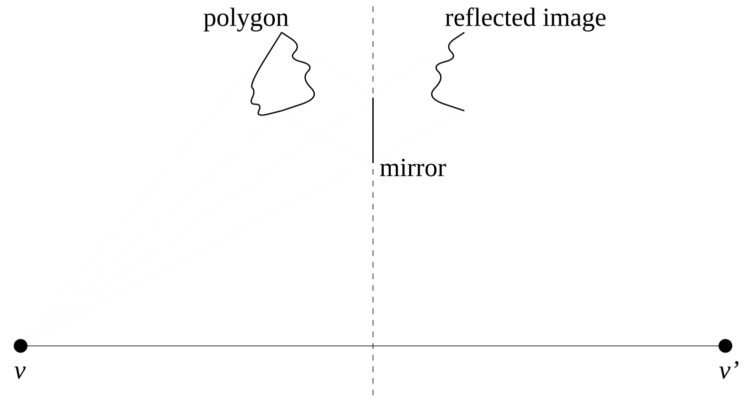

Figure 25: A viewpoint and a mirror.

The viewpoint v can be translated towards the polygon along the line vv' as long as it remains in its feasibility region (see Figure 26). Such a translation of vector $\vec{t}$ implies a translation of the mirror line (away of the polygon) of vector $\frac{\vec{t}}{2}$, and thus produces a translation of the reflected image of vector $\vec{t}$. With these movements, no obstruction can appear, because the lines of view are translated away from the polygon.

The reflected point v' can be rotated away from the polygon centering the rotation in v, as long as it remains in the corresponding feasibility region (see Figure 26). Such a rotation of angle α implies a rotation of the mirror line of angle α, which produces a movement of the reflected image. With this movement, no obstruction can appear, because the lines of view rotate away from the polygon.

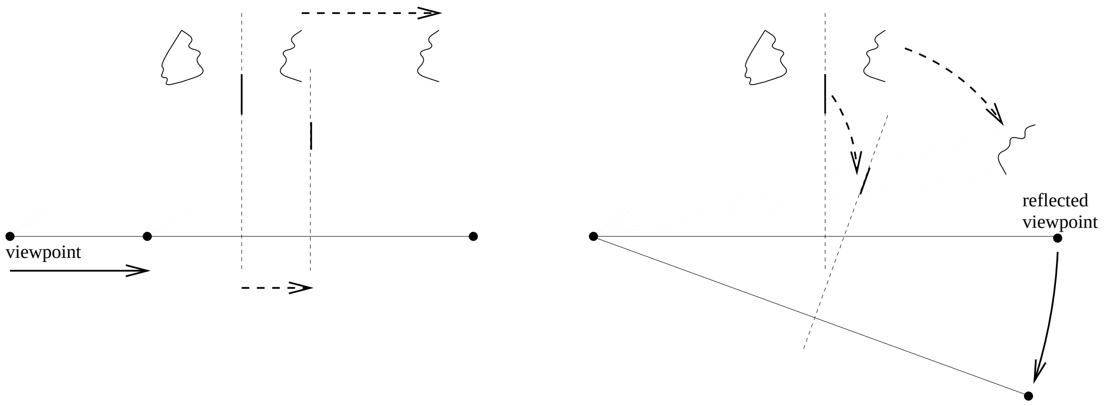

Figure 26: How to move the viewpoint and its reflected image without generating obstructions.

Both movements, the translation of v and the rotation of v', always take these points to the boundary of their visibility regions, because both take them to the line joining the antipodal pair, and such a line cannot be crossed without leaving the visibility regions. □

This nice result, anyway, does not mean that the viewpoint and its reflected image can always be found in the vertices of their feasibility regions. Figure 27 shows an example in which the viewpoint cannot be moved to a vertex: moving the viewpoint upwards makes the polygon obstruct the vision line from the viewpoint to the reflected image of the polygon, while moving the viewpoint downwards makes the mirror intersect the polygon. Simple modifications of this example can be used to prove the analogous

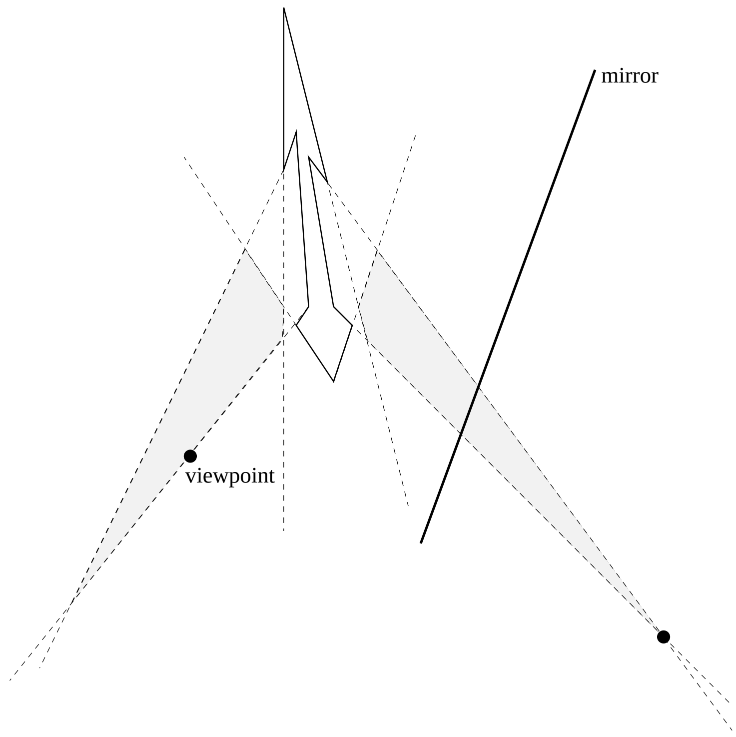

Figure 27: The viewpoint v cannot be moved to a vertex.

fact for the reflected viewpoint, as in Figure 28.

Now we are ready to solve the one-viewpoint-one-mirror problem.

Theorem 3.8 *In $O(n^3)$ time it is possible to decide whether or not a given n-gon can be seen with a viewpoint and a mirror, and find possible locations for them.*

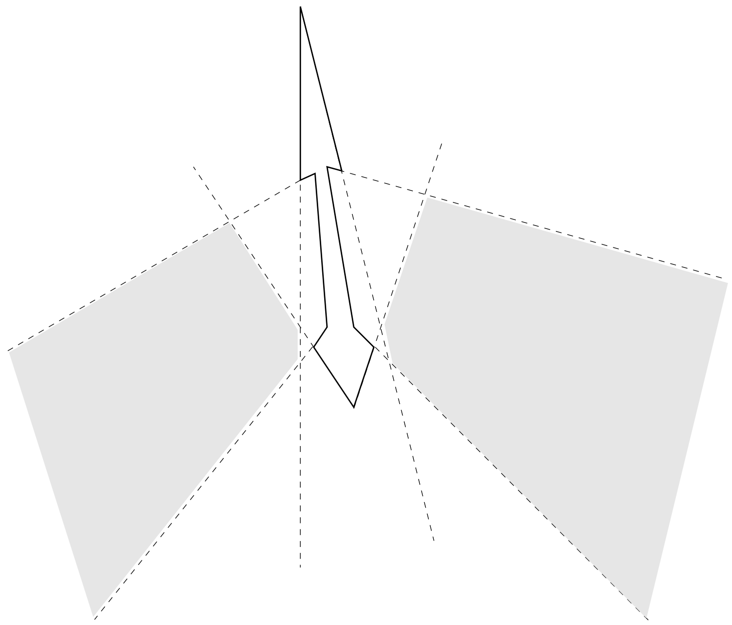

Figure 28: The reflected viewpoint v' cannot be moved to a vertex.

Proof: For every antipodal pair of vertices of the polygon, the algorithm essentially checks the boundary of a feasibility region against the boundary of its complementary one. Some subtleties must be taken into account, because given a viewpoint in one of the boundary chains, the obstruction to the vision of the reflected other chain can be determined by all antipodal pairs containing one of the two selected vertices.

Concretely, when checking an edge of the boundary of the say "left" region against an edge of the "right" region, looking for a feasible position of the viewpoint in the left one and for the reflected viewpoint in the right one, the problem to be considered, which is the obstruction, depends on the supporting line from the viewpoint to the polygon. This line, as can be seen in Figure 29, can touch the polygon in different vertices: namely the antipodal pairs of vertex n. The left edges must hence be decomposed into

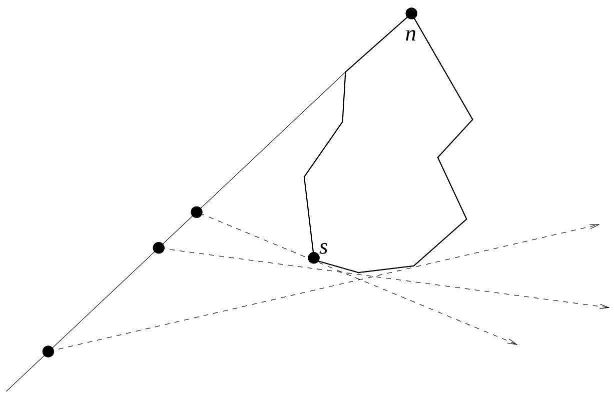

Figure 29: Different supporting points.

subsegments, each one giving the same vertex of the polygon for the supporting line.

The algorithm checks each left subsegment against each right edge for a feasible position of the viewpoint and its reflected viewpoint. This can be done in constant time, by deciding whether or not the critical reflection point r can be made to lie below the obstruction line *obs* (refer to Figure 30).

The entire algorithm runs in $O(n^3)$ time, because it considers each antipodal pair of vertices of the polygon (there are $O(n)$) and, for each one, checks each subsegment of the boundary of one region (there are $O(nk_i)$, where k_i = #antipodal pairs of vertex v_i) against each edge of the boundary of the other

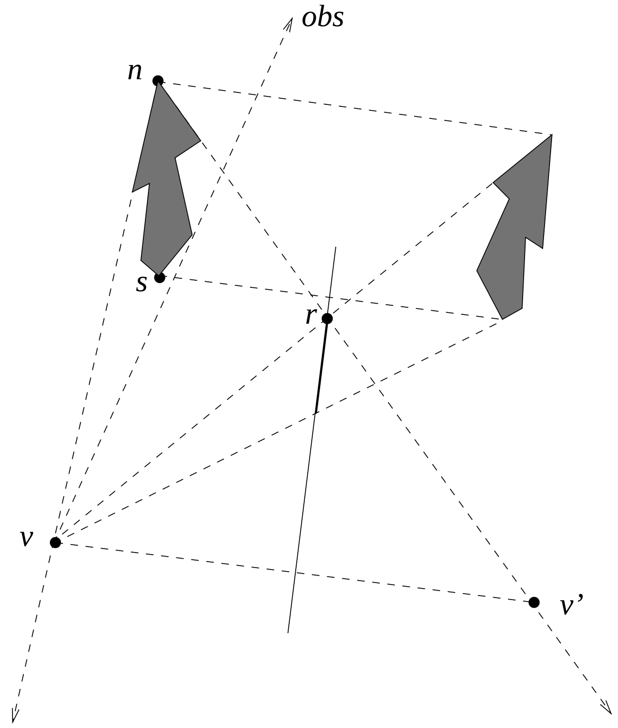

Figure 30: How to solve the obstruction problem.

region (there are $O(n)$) in constant time. The overall cost is, then

$$T(n) = \sum_{i=1}^{\#\text{ant.pairs}} O(nk_i)O(n) = O(n^2) \sum_{i=1}^{\#\text{ant.pairs}} O(k_i) = O(n^2)O(n) = O(n^3).$$

□

3.3 Seeing a polygon with one viewpoint and two mirrors

When one viewpoint and a mirror do not suffice to visualize a polygon, additional mirrors are needed. However, the algorithmic problems become considerably more difficult with two mirrors, or even with the less restricted problem of three viewpoints. First, observe that there exist polygons that can be seen with three viewpoints, but cannot be partitioned into three contiguous connected chains so that each viewpoint sees one of the chains. An example is provided in Figure 31, in which the only way to place three viewpoints makes two of them see unconnected portions of the polygon. Concretely, v_1 sees abc, v_2 sees $cdefgh$ and $jklm$, and v_3 sees $ghij$ and $lmnopa$.

Things do not get easier if we restrict ourselves to consider only connected vision, for there exist polygons that can be seen with three viewpoints so that each viewpoint sees a connected chain, but some endpoints of the chains are not vertices. See the example of Figure 32, where a cutting point is needed in the thick part of the polygon in order to see the polygon as three connected chains.

Finally, we could restrict ourselves to only consider decompositions of the polygon in three connected chains with endpoints in the vertices. Even in that case, results cannot be very encouraging, because it may be possible to decompose the boundary of a polygon into three of these chains, each one seen from a viewpoint, in such a way that the endpoints of the chains do not belong to the convex hull of the polygon, making Lemma 3.5 impossible to apply.

Theorem 3.9 *In $O(n^2 \log^2 n)$ time it is possible to decide how many viewpoints are needed to completely see a given n-gon in a connected way without cutting points, and find their feasibility regions.*

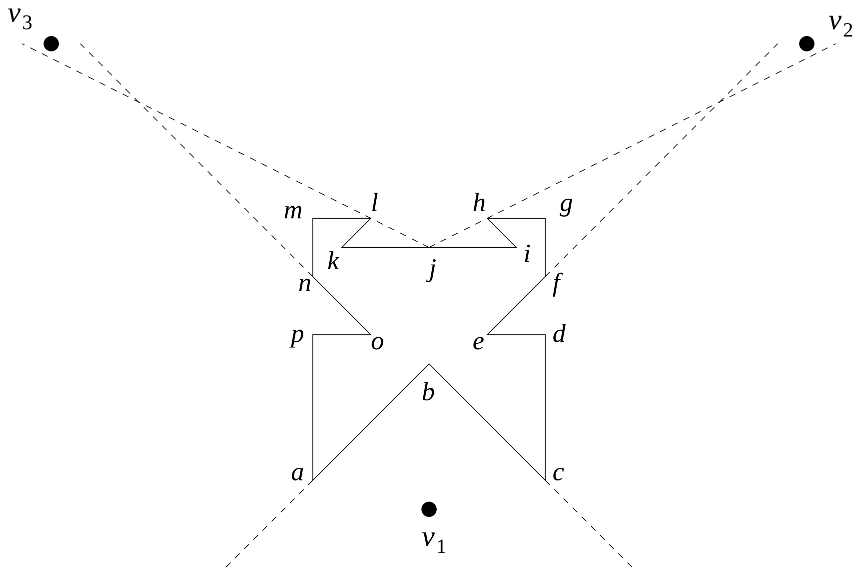

Figure 31: A polygon that cannot be seen in a connected way.

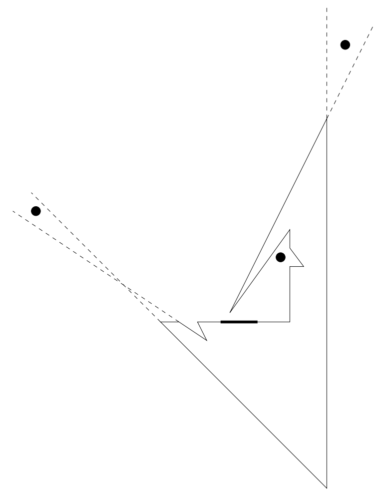

Figure 32: A cutting point is needed.

Proof: The algorithm essentially computes the visibility region of each vertex of the polygon, then finds all maximal chains that can be seen with a single viewpoint and finally walks along the polygon to find a minimum cover.

The visibility regions of all vertices can be computed in $O(n)$ time [6].

In order to compute the maximal chains efficiently, observe that the chains having two (or more) vertices belonging to different pockets of the polygon are easy to deal with, because their visibility regions are intersections of halfplanes external to the convex hull of the polygon. This would take us to an algorithm based on a on-line maintenance of the intersection of halfplanes (that can be done in $O(n \log^2 n)$ time [12]) if it was the same for the chains completely lying in a pocket, but it is not: the region from where one of such chains can be seen is not necessarily convex (see Figure 33.

Nevertheless, the visibility region of any subchain of a pocket can be decomposed into two parts: the external and the internal one (with respect to the lead of the pocket). The external one is again an intersection of halfplanes, while the internal one is not, and must be analyzed in more detail.

Consider one pocket of the polygon. In order to compute the maximal chains of the pocket that can be seen from a single viewpoint internal to it, we could intersect the visibility regions of contiguous vertices until the result was not empty, and redo it again and again, starting with each vertex. This

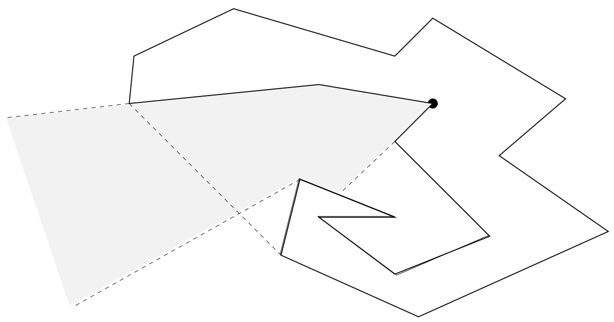

Figure 33: The visibility region of this chain is not convex.

process would solve the problem, but with a high cost. Instead of doing that, we triangulate the polygon and then maintain, for each triangle of the triangulation, the information of its intersections with the visibility region of each vertex of the polygon (that is, at most two halfplanes for each vertex). All right most maximal chains of the pocket that can be seen from a single viewpoint can then be computed in $O(n^2 \log^2 n)$ time by an on-line maintenance of the intersection on the halfplanes [12] in *each* triangle (they are $n-3$).

The external visibility region of a chain in a pocket is, as said, an intersection of halfplanes. Concretely, each vertex in a pocket can be externally seen from a region formed by the intersection of three halfplanes (remember Figure 33). One of them is determined by the lead of the pocket, while the other two can be found in $O(\log n)$ time per vertex with a linear preprocess of the pocket [7].

All together, a combination of both pocket internal and external vision allows to compute all maximal chains in the polygon in $O(n^2 \log^2 n)$ time.

Finally, we use all maximal chains to find a minimum cover of the polygon by walking along its boundary. Start at any vertex v_1. One single viewpoint allows to see the entire maximal chain starting at v_1. Go to its endpoint v_2 and use the maximal chain starting at that point v_2 to continue covering the polygon. Keep walking around the polygon until you find an endpoint that was yet visited (this will happen after at most n steps). At that point, you can be sure that you found a cover. Now you only need to check whether or not it wraps only once around the polygon. If so, we have finished. If not, take any subcover of it that wraps around the polygon once. The maximality of the chains allows to prove, in both cases, that the cover uses a minimum number of viewpoints. □

3.4 Acknowledgments

The authors whish to thank Belén Palop and Martí Sánchez, of the Dep. Matemàtica Aplicada II of the Universitat Politècnica de Catalunya, for her essential help in generating the Earth model for Figure 11 and in creating the images of Figures 12, 13 and 14.

References

[1] B. Aronov, A. R. Davis, T. K. Dey, S. P. Pal and D. C. Prasad, Visibility with Reflection, *Proc. 11th ACM Symp. on Computational Geometry*, 1995, pp. 316-325.

[2] J. Czyzowicz, E. Rivera-Campo and J. Urrutia, Illuminating rectangles and triangles in the plane *J. Combin. Theory Ser. B*, 57 (1993) pp. 1-17.

[3] J. Czyzowicz, E. Rivera-Campo and J. Urrutia, Optimal floodlight illumination in polynomial time, *Proc. 5th Canad. Conf. Comput. Geom.*, 1993, pp. 393–398.

[4] G. Das and G. Narasimhan, Optimal Linear-Time Algorithm for the Shortest Illuminating Line Segment in a Polygon, *Proc. 10th Annu. ACM Sympos. Comput. Geom.* 1994, pp. 259-266.

[5] H. Edelsbrunner and J. W. Jaromczyk, How often can you see yourself in a convex configuration of mirrors?, Report, UIUCDCS-R-86-1256, Dept. Comput. Sci., Univ. Illinois, Urbana, IL, 1986.

[6] H. ElGindy and D. Avis, A linear algorithm for computing tha visibility polygon from a point, *J. Algorithms*, 2 (1981) pp. 186-197.

[7] L. J. Guibas and J. Hershberger, Optimal Shortest Path Queries in a Simple Polygon, *J. Comp. Sys. Sciences*, 39 (1989) pp. 126-152.

[8] J. Hershberger and S. Suri, Offline Maintenance of Planar Configurations, *Proc. 2nd ACM/SIAM Symp. on Discrete Algorithms (SODA)*, 1991, pp. 32-41.

[9] V. Klee, Is every polygonal region illuminable from some point? *Amer. Math. Monthly*, 76 (1969) p. 80.

[10] V. Klee and S. Wagon, *Old and New Unsolved Problems*, Mathematical Association of America, 1991.

[11] G. H. Meisters, Polygons have ears, *Amer. Math. Monthly*, 82 (1975) pp. 648-651.

[12] F. P. Preparata, An optimal real time algorithm for planar convex hulls, *Comm. ACM*, 22 (1979) pp. 402-405.