Visualizing Objects with Mirrors

F. Hurtado and M. Noy and J.-M. Robert and V. Sacristán and S. Skiena

I. Introduction

Computer graphics techniques such as ray tracing are commonly 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.

However, realizing these advantages requires carefully positioning the mirrors in the scene. Consider the deficiencies in the image shown in Figure 1(a), which consists of the Earth and three axis-parallel mirrors. Secondary reflections off the mirrors create a total of six reflected images, three of which spurious and confusing. In addition, the direct view of the model obscures parts of one primary and two secondary reflections. A third, more subtle, deficiency is that certain parts of the world (such as Europe) remain completely and unnecessarily occluded.

Contrast this with the image shown in Figure 1(b), which consists of the Earth amid three carefully positioned, non-axis oriented mirrors. These flat but elliptical mirrors are cut so as to eliminate the possibility of secondary reflections. There is no occlusion, in fact the projection of the reflected images is tangent to the primary image. Finally, this configuration has the property that (1) every point on the surface of the earth is visible in either primary or reflected image, and further (2) the mirror images are large and clear – in fact provably the highest-resolution possible under any such configuration.

In this paper, we provide algorithms and geometries for optimal object-independent 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 II, 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.

To evaluate the practical merits of our derived mirror configuration, we present ray-traced images showing simulated views from both optimal and non-optimal mirrors configurations from a variety of viewpoints. We literally provide a new way to see the world. The color images in this paper are are best viewed at http://www.cs.sunysb.edu/~skiena/mirrors.

Our proposed mirror configuration has potential applications beyond simple visualization:

- *Fly-throughs* – One can easily imagine using our configuration as the basis of a fly-through, where the entire surface of the object is visible in every frame of the visualization. In this paper, we show how to calculate the optimal mirror positions from any possible view point.
- *3D Scanning* – Three-dimensional objects such as sculptures (e.g. the Digital Michaelangelo project [10]), can be digitized using either laser/sonic range data or interpolating between multiple images. This require getting direct access to all sides of the object, which may not be possible due to occlusion or difficulties in physically positioning the object and/or sensor system. By instead embedding physical mirrors in configurations such as presented in this paper, we can completely view the object from the sensor without moving either.
- *Dental Imaging* – Mirrors are commonly employed by dentists to see the back surfaces of teeth. Properly interpreting the orientation of the resulting views is non-trivial, and our visualizations can provide excellent examples for training. Further, illustrations of simulated dental conditions can be provided by ray-traced graphics, with our techniques being applicable to identifying the best placement of the mirrors.

[1]Partially supported by Acción CICYT no. TIC-747/93-C02, Acción CICYT no TIC95-0957-E, Proyecto MEC-DGES-SEUID PB96-0005-C02-02, Projecte UPC PR-9410, CUR Gen. Cat. 1999SGR00356, Proyecto DGES-MEC PB98-0933, and the Joint Commission USA-Spain for Scientifical and Technological Cooperation Project 98191. The stay of Vera Sacristán at SUNY Stony Brook was supported by 1995BEAI200036.

Dep. Matemàtica Aplicada II, Universitat Politècnica de Catalunya.
Dep. Matemàtica Aplicada II, Universitat Politècnica de Catalunya.
Dép. d'Informatique et de Mathématique. Université du Québec à Chicoutimi.
Dep. Matemàtica Aplicada II, Universitat Politècnica de Catalunya.
Dept. of Computer Science, State University of New York at Stony Brook

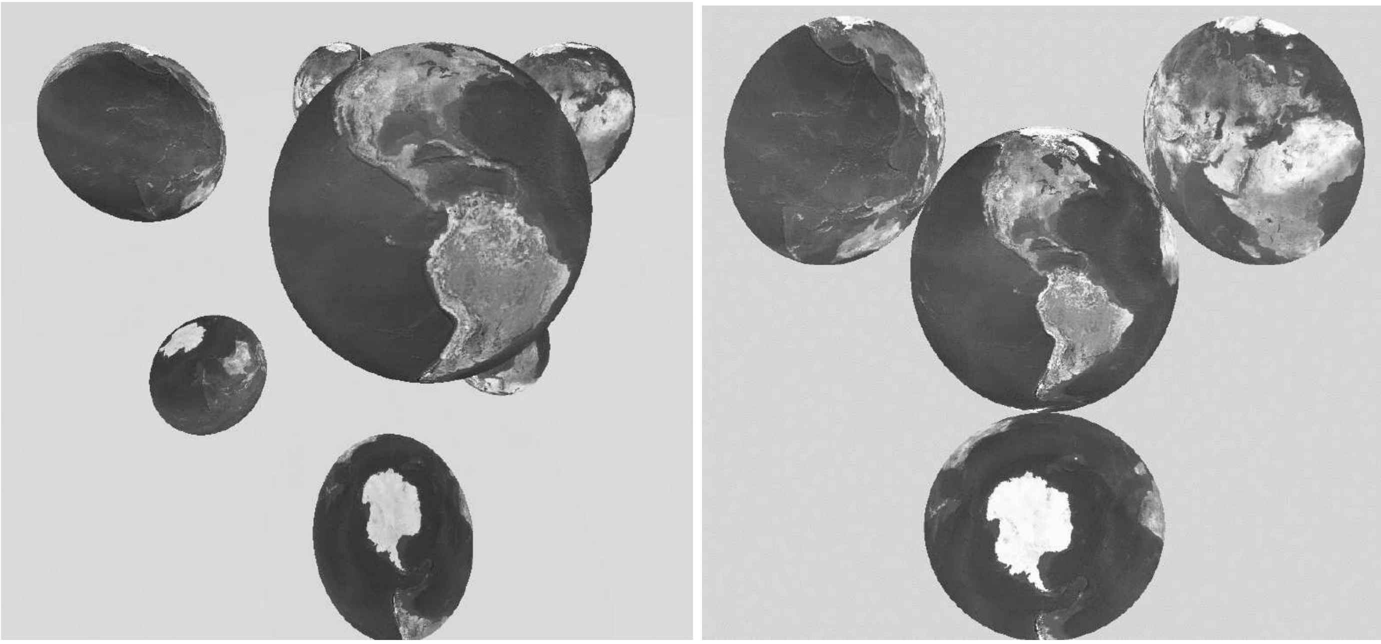

Fig. 1. Examples of bad and good mirror placement in a scene. (a) contains three orthogonal mirrors, illustrating the problems of secondary reflections and occlusion, while (b) shows the complete surface of the earth using our optimal mirror configuration.

In Section IV, we briefly 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. These problems prove to be very difficult for non-convex objects. We illustrate some the complexities, and point to algorithmic results (of a primarily theoretical nature) to overcome them.

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 [8] posed the question of whether all polygons with mirror reflecting edges can be illuminated from a point source, which generated a substantial literature [9] before being answered in the negative by Tokarsky [13]. 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.

II. 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 position of 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.

A. The two dimensional problem

Theorem II.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. 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. Thus, p can be replaced by any of these p'.

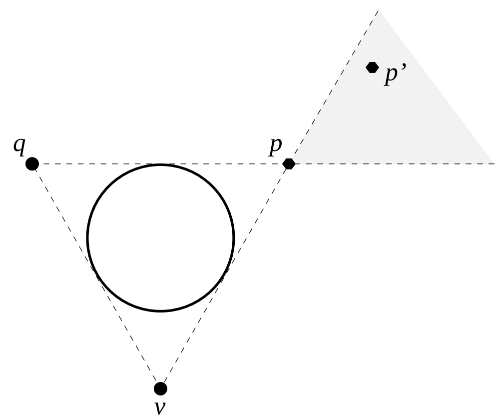

Fig. 2. Viewpoints v, p, and q together see the entire circle. Further, viewpoint p can be replaced by p'.

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.

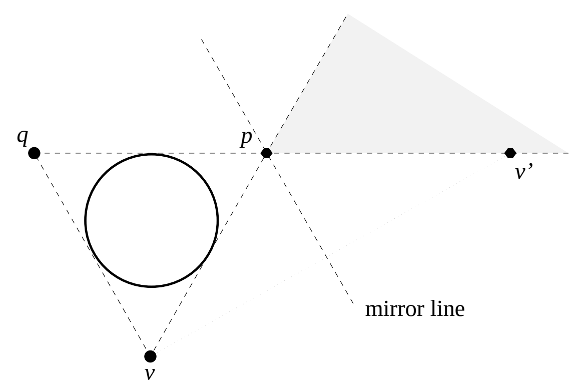

Fig. 3. Where to place the mirror line.

Finally, notice that the size of the mirror can be found by intersecting the mirror line with the cone determined by the reflected image of the arc seen from p and the viewpoint v, for there is no need to see with the mirror the entire arc that v' directly sees, but only the arc that p sees. Observe that the mirrors can be made arbitrarily small by translating v' away from the scene along any line contained in the shaded wedge.

■

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 viewpoint to (1) the visible portion of the 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 II.2: The construction of Theorem II.1 minimizes the maximum distance between the viewpoint 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 4). It should be obvious that a circumscribed one will always be better, as shown in the second step of Figure 4. 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 4). For all these triangles, the maximum distance from the viewpoint to the circle is equal to the semi-perimeter of the triangle. Among all of them, the equilateral triangle minimizes the perimeter and hence is the better one for our purposes. ■

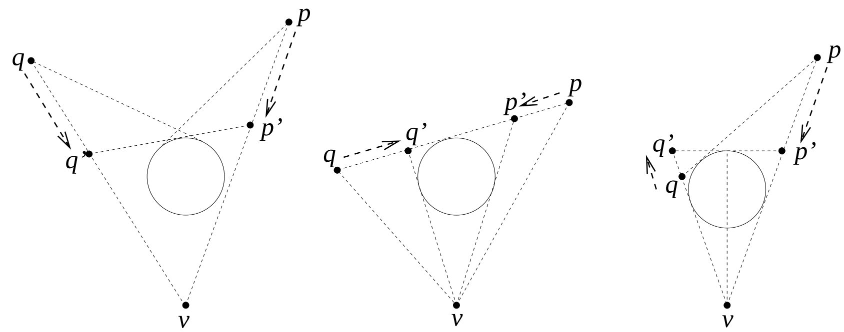

Fig. 4. How to decrease the maximum distance in three steps.

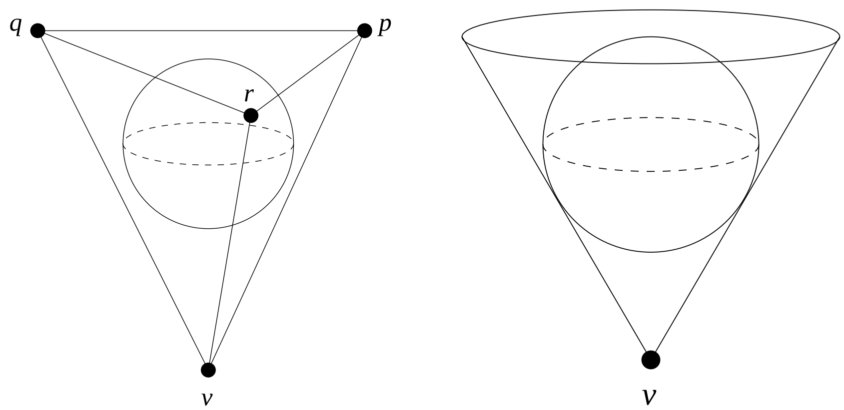

Fig. 5. Placement of the four viewpoints, and the obstruction cone

B. 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 II.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.

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.

The points whose view is obstructed by the sphere lie in a cone circumscribed to the sphere having apex v. 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 $60°$.

Figure 6 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 major axis lies in the planar section of Figure 6, and the minor one is perpendicular to it.

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

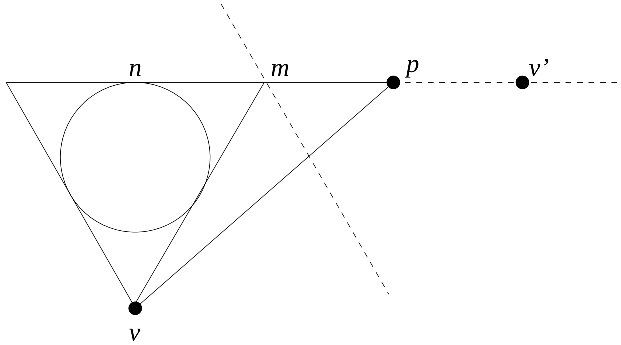

Fig. 6. Planar section of the solution.

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). ■

Theorem II.4: The construction of Theorem II.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 6), which is again the semi-perimeter of the equilateral triangle circumscribed to a circle. The two dimensional argument holds. ■

III. Implementing the Optimal Mirror Configuration

Without loss of generality, we will explain how to position only one of the mirrors, since once it has been determined the other two are obtained by rotating the first one with respect to the z-axis by 120 and -120 degrees. We break the computation into two phases, (1) computing the mirror plane, and (2) computing the mirror center and axis.

1. *Computing the mirror plane* – Without loss of generality, assume That is, we are assuming that the sphere has center $(0, 0, 0)$ and radius 1, and that the viewpoint is $v = (0, 0, -d)$, for some distance d. The mirror plane will be determined by a point m and a vector nn where m belongs to the mirror plane and nn is normal to the plane. In fact, m is the point from which the mirror will hang.

The mirror plane may be computed as follows:

(a) Compute r, the radius of the circle basis of the cone that is circumscribed to the sphere and has vertex in v. The result is $r = \sqrt{\frac{d+1}{d-1}}$. The point m is then $m = (r, 0, 1)$.

(b) The vector nn is the direction of the bisector of the angle $\widehat{mnv}$, Where $n = (0, 0, 1)$ is the north pole of the sphere (the point diametrically opposite to the viewpoint). The result is $nn = (\sqrt{d+1}, 0, \sqrt{d-1})$.

2. *Computing the center and the two axis of the mirror* – Recall that the needed mirror is elliptical (the paper says why) as illustrated in Figure 7. Given the viewpoint distance d (as before), we seek the center and two axes (axis-A and axis-B) of the ellipse. axis-A lies in the plane $y = 0$, while axis-B is perpendicular to it. These may be computed as follows:

(a) Compute p the vertex of the circumscribed tetrahedron which will be replaced with our mirror. The point p sees a cap of the sphere. We will first determine the diameter of that cap that lies in the plane $y = 0$. One of its endpoints is the $n = (0, 0, 1)$. The other one will be called q. The point $p = (2r, 0, 1)$, where r is computed above.

(b) Compute q. This can be done in two ways: either (1) finding the point belonging to the sphere such that $\widehat{npo} = \widehat{opq}$, where o is the origin, and center of the sphere or (2) finding the point belonging to the sphere such that the vectors $\overrightarrow{nq}$ and $\overrightarrow{op}$ are orthogonal.

Now we use the diameter nq to compute the axis-A. This can be done by reflecting q through the mirror and obtaining the reflected point q'. Then one can compute the portion of the mirror which is necessary to see q' from v.

(c) Compute q' and let a be the intersection of the mirror plane and the line vq'

(d) Thus the center of the mirror is $c = \frac{m+a}{2}$.

(e) The axis lying in the plane $y = 0$ is then $axisA = \frac{1}{2}d(m, a) = d(c, m) = d(c, a)$.

(f) Computing the axis-B is complicated because it does not correspond to a diameter of the cap.

i. Let c' be the intersection of the line $v'c$ and the line nq (That is, c' is the projection of the center of

ii. Compute b', one of the two points on the sphere that has the same x and z coordinates as c' (That is, project c' onto the basis circle of the cap, perpendicularly to nq).

iii. Compute b, the intersection of $v'b'$ with the mirror plane (That is, b is the endpoint of the axis-B).

iv. Compute $axisB = d(b, c)$.

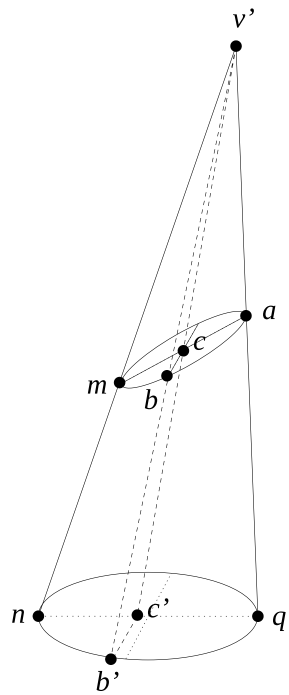

Fig. 7. Computing the center and the two axis of the mirror

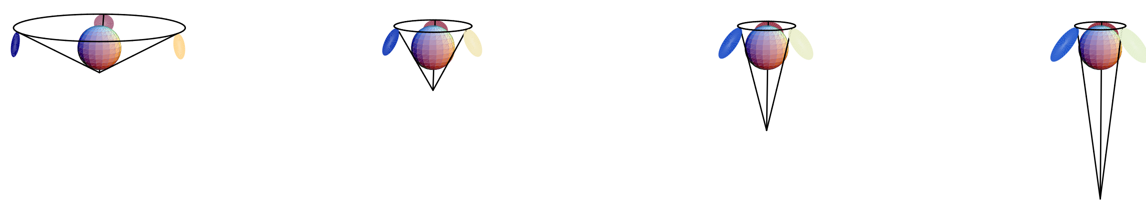

Fig. 8. Illustrating the mirror positions as the view point increases, for viewpoint distances 1.15, 2, 4, and 7.7 from the center of the sphere.

Here are the coordinates for the optimal position assuming the sphere is located at $(0,0,0)$ and of radius 1. The viewpoint then is at distance 2 from the origin and center of the sphere. Assuming that $v = (0,0,-2)$, this implies that the calculations give the following results:
- Mirror hanging point $m = (\sqrt{3}, 0, 1)$
- Mirror normal vector $nn = (\sqrt{3}, 0, 1)$
- Mirror center: $c = (\frac{51\sqrt{3}}{43}, 0, \frac{19}{43}) \approx 0.6445$
- Mirror major axis: $axisA = \frac{4\sqrt{43}}{43} \approx 0.6101$

Figure 8 illustrate the relative positions of the mirror and as we move the viewpoint away from the center of the sphere. Observe that as the viewpoint gets farther from the sphere, the mirrors get closer to the sphere (as the obstruction code thins) and the mirrors fan out to wider angles.

There are two limit positions, one as $d \to 1$ and the other one when $d \to +\infty$. When the $d \to +\infty$:
- Mirror hanging point: $m \to (1, 0, 1)$
- Mirror normal vector: $nn \to (\frac{\sqrt{2}}{2}, 0, \frac{\sqrt{2}}{2})$, that is, tends to form 45 degrees with vertical.
- Mirror center: $c \to (3/2, 0, 1/2)$
- Mirror major axis: $axisA \to \sqrt{2}/2$
- Mirror minor axis: $axisB \to 1/2$

When $d \to 1$,
- Mirror hanging point: $m \to (+\infty, 0, 1)$
- Mirror normal vector: $nn \to (1, 0, 0)$, that is, tends to be vertical.
- Mirror center: $c \to (+\infty, 0, 1/2)$
- Mirror major axis: $axisA \to 1/2$

- Mirror minor axis: $axisB \to 1/2$

The simulations of the scenes with our mirror configurations were created with rayshade, a ray-tracer created at Princeton University by C. Kolb and successively updated (http://www-graphics.stanford.edu/˜cek/rayshade). In all these images, the sphere has center (0,0,0) and radius 1, and only the viewpoint and mirrors change.

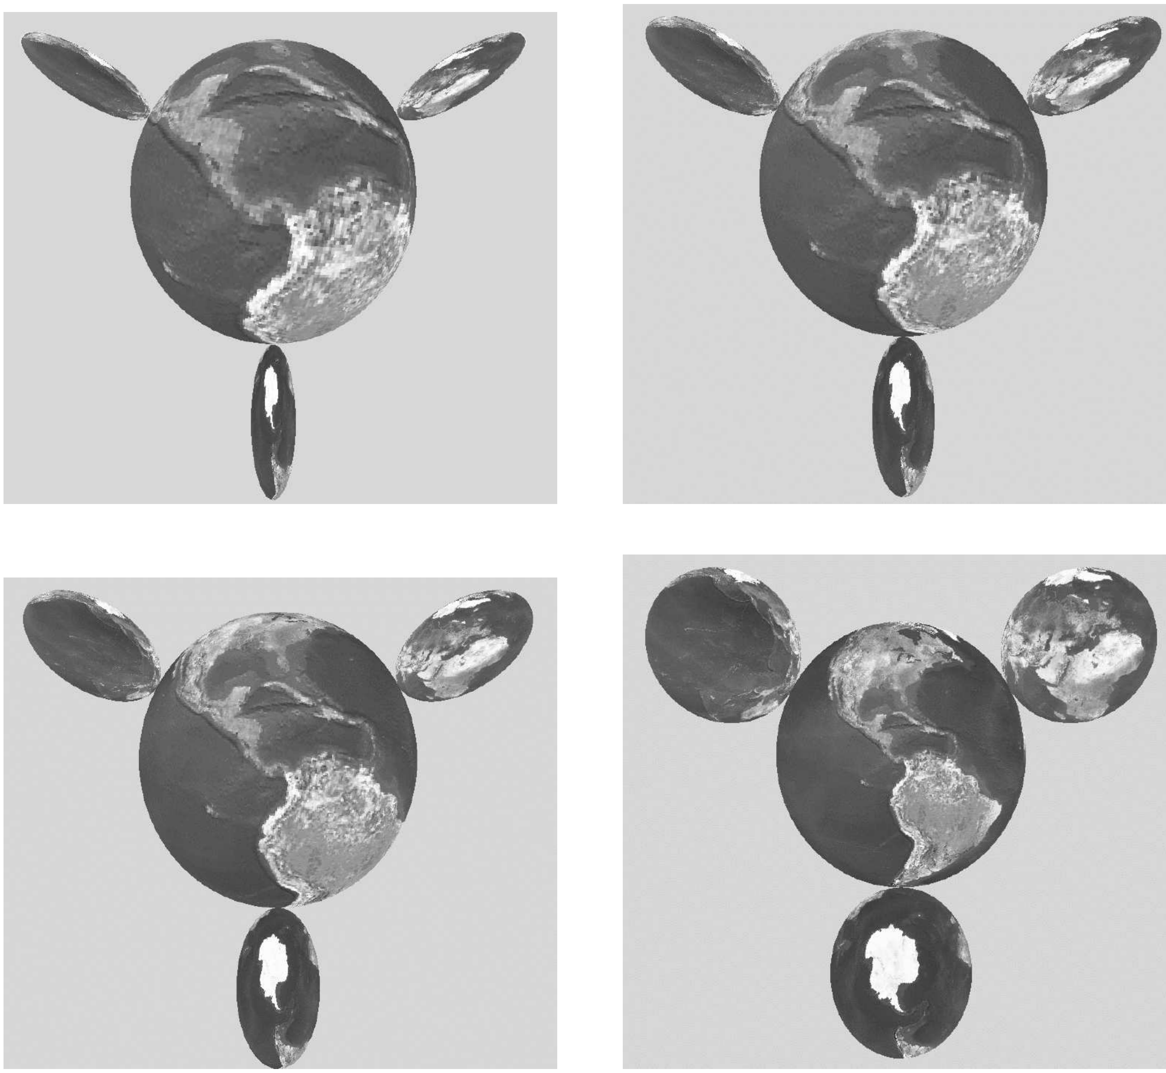

Fig. 9. Perspectives of the Earth from different viewpoints in our mirror construction, top row, $d = 1.1$ and $d = 1.2$. Bottom row $d = 1.4$ and $d = 3.0$.

Figure 9 shows the Earth in our mirror configuration as the viewpoint moves back after starting just above its surface. Observe that the focussed region expands and the mirrors gain increasing perspective as we move farther back. Note that the raytracer produces strange optical illusions due to its "camera like" behavior. For example, when the viewpoint approaches the sphere in Figure 9, it shows the mirrors as being more and more elliptic, while in fact they get closer to circular. When the viewpoint is very far away from the sphere, it shows the mirrors as being as big as the sphere, while they should be seen as half of it. When the viewpoint is very far away from the sphere, and the angle is small, it produces some precision errors.

Other effects are more reasonable. For example when the viewpoint is too far from the scene, the direct view looks very unnatural, because it is like a cylindrical projection. When the viewpoint is too close to the scene, the direct view looks very unnatural, because it is like a conic projection, and the deformation is big.

It is instructive to compare our optimal mirror configuration with the optimal axis-oriented configuration. The two are illustrated side-by-side in Figure 10, with the axis-oriented view taken up the main diagonal at a viewpoint distance such that the reflected images are tangent to the primary image. Note that the mirrors on Figure 10(b) appear significantly larger than those of the axis-oriented construction, yielding higher resolution in the reflected images.

These illustrate the position of the optimal viewpoint and mirror position.

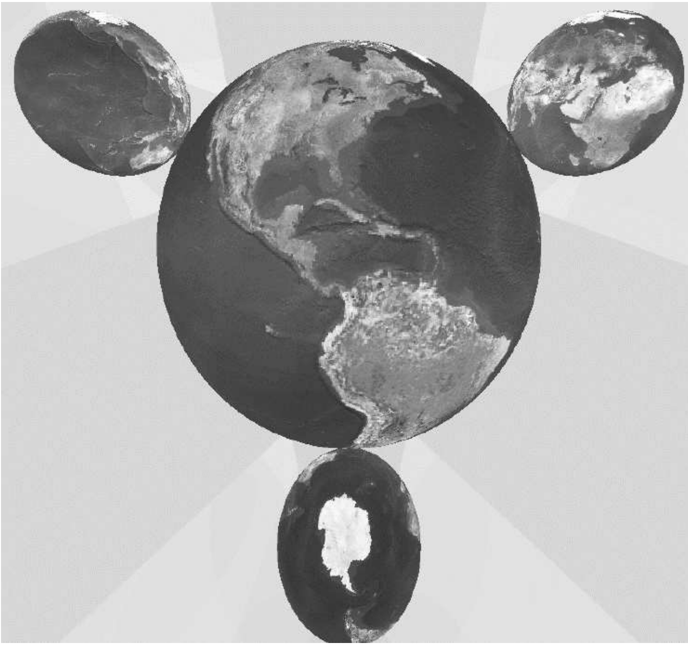

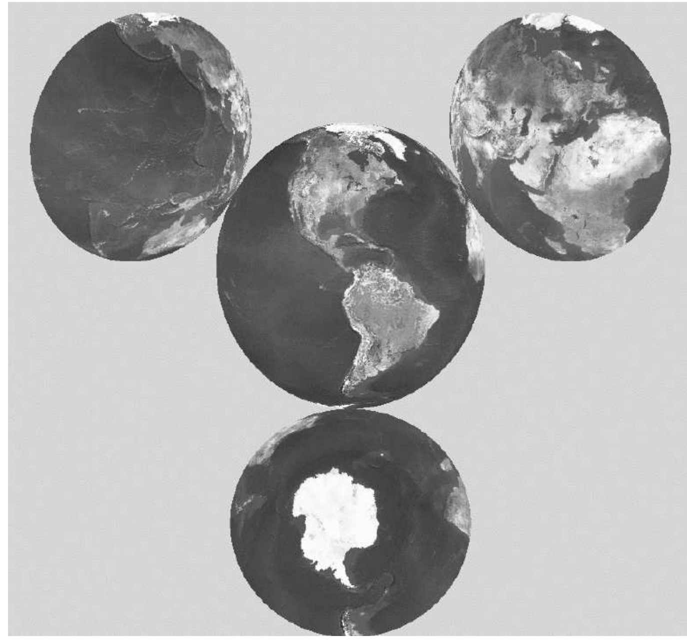

Fig. 10. Comparing the best axis-oriented view (a) against the proposed mirror configuration (b). Observe the increase in the size and resolution of the mirrors in the right view.

IV. Object-Dependent Problems

In 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.

The algorithmics of mirror position prove to be very complicated, even in two dimensions, and are clearly too messy to justify implementation. It is important to understand some of these complications, as they impact the ad hoc selection of mirror configurations.

Note that the fixed mirror configuration of the previous section can be profitably employed even to view objects which are not convex. Consider the arrangement of eggs in a shell shown in Figure 11. As configuration is not convex, there is no guarantee that all surfaces are visible. However, the mirror reflections clearly reveal different sides of the eggs with the shell. This effect is particularly clear in the full color rendition of the scene.

- *Multiple concavities* – Every convex polygon can be viewed with one viewpoint and one mirror, if one positions the viewpoint and mirror properly. However, tiny concavities can require up to $n/3$ viewpoints to see all faces (see Figure 12(a)), and 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 12(b).
- *Disconnected visibility regions* – For visual coherence, it is important that adjacent points of the surface of the object appear within the same (reflected) image, except perhaps at the boundary of the mirror. However, this coherence property cannot be maintained while simultaneously minimizing the number of viewpoints. Specifically, 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 13, 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 $lmnopa$.
- *Division of surface edges* – Even if we restrict ourselves to consider only connected vision, individual surface features may need to be divided in unsatisfactory ways. Figure 14 shows 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. Note that a cutting point is needed in the thick part of the polygon in order to see the polygon as three connected chains. The evaluation of such a cutting point involves non-trivial numerical issues.

For completeness, we cite our algorithmic results [7] on mirror placement, which overcome these difficulties in restricted cases.

- 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.

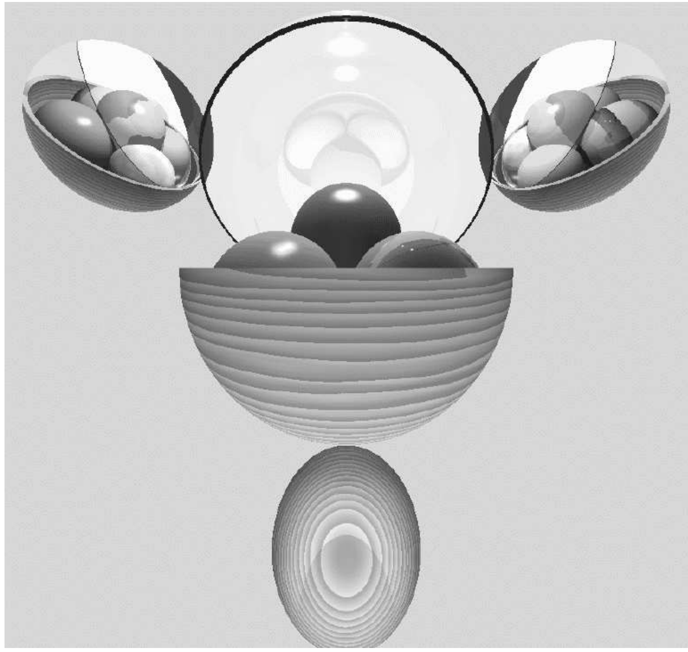

Fig. 11. Visualization of hidden surfaces in a complex scene using the optimal mirror configuration.

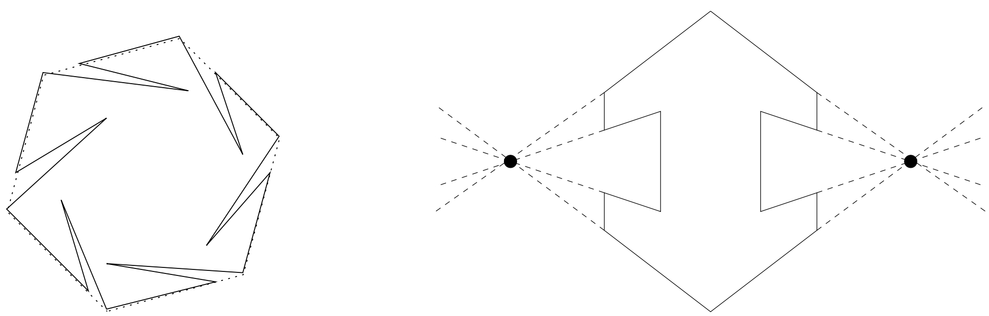

Fig. 12. (a) A polygon requiring $n/3$ viewpoints to completely see, and (b) a polygon that cannot be seen with a viewpoint and a mirror.

- An $O(n \log n)$ algorithm to determine all placements of pairs of viewpoints sufficient to view the complete surface of P.
- An $O(n^2 \log^2 n)$ algorithm to determine the minimum number of viewpoints sufficient to view the complete surface of P in connected chains without inter-edge cutting points.
- An $O(n)$ algorithm to position a viewpoint and mirror around a convex n-gon so as to minimize the maximum viewing distance.

A. Acknowledgments

The authors thank Belén Palop and Martí Sánchez, of the Dep. Matemàtica Aplicada II of the Universitat Politècnica de Catalunya, for their essential help in generating the Earth images.

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] F. Hurtado, M. Noy, J.-M. Robert, V. Sacristan, and S. Skiena, "Visualizacio'n de Objectos Mediante Espejos" *CEIG'96* (in Spanish), Valencia, Spain, June 1996. Also in English as Visualizing Objects with Mirrors report 96-01, Department of Computer Science, State University of New York, Stony Brook, March 1996.

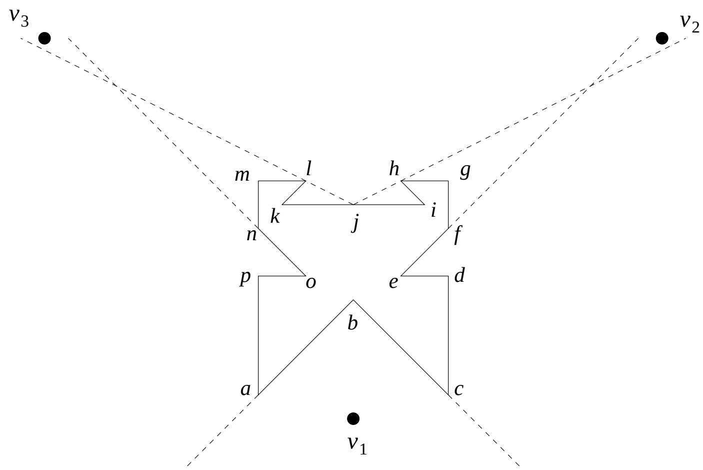

Fig. 13. A polygon that cannot be seen with three viewpoints in a connected manner.

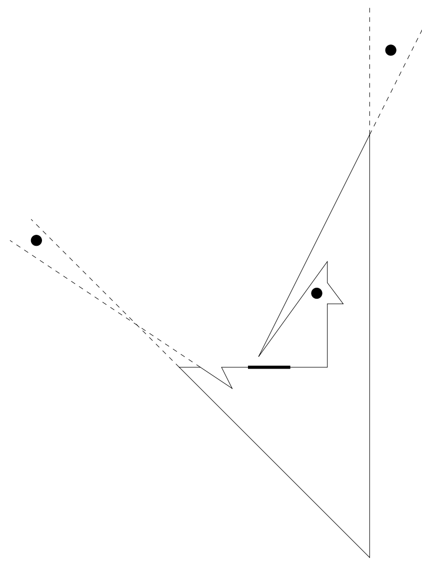

Fig. 14. A cutting point is needed.

[8] V. Klee, Is every polygonal region illuminable from some point? *Amer. Math. Monthly*, 76 (1969) p. 80.
[9] V. Klee and S. Wagon, *Old and New Unsolved Problems*, Mathematical Association of America, 1991.
[10] M. Levoy, K. Pully, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk. The Digital Michelangelo Project: 3D Scanning of Large Statues, *Proc. SIGGRAPH 2000*. To appear.
[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.
[13] G. M. Tokarsky, Polygonal rooms not illuminable from every point. *Amer. Math. Monthly* 102 (1995) pp. 867-879.