APPLICATIONS OF WORLD PROJECTIONS

Ned Greene
Computer Graphics Lab
New York Institute of Technology
Old Westbury, NY 11568

Abstract

Various techniques have been developed which employ projections of the world as seen from a particular viewpoint. [Blinn and Newell] introduced reflection mapping for simulating mirror reflections on curved surfaces and their method can be extended to simulate refraction. [Miller and Hoffman] have presented a general illumination model based on world projections. [Greene] has used projections of the world to model distant objects, and [Greene and Heckbert] have used world projections to produce pictures with the fisheye distortion required for Omnimax® frames. World projections can also be used as a backdrop for ray tracing or beam tracing.

This paper proposes a uniform framework for representing and utilizing world projections and argues that the best general purpose representation is the projection onto a cube. Surface shading and texture filtering issues related to environment mapping are discussed including approximate methods for obtaining diffuse and specular shading values from prefiltered environment maps. It is noted that obtaining accurate diffuse reflection and antialiasing specular reflection, which are both problematical with ray tracing, can be effectively handled by environment mapping.

Keywords: Environment mapping, reflection mapping, surface shading, texture mapping, cube projection, Mercator projection.

Omnimax is a registered trademark of Imax Corporation, Toronto, Canada.

1. INTRODUCTION

Reflection mapping, introduced by Blinn and Newell in 1976, is a shading technique that uses a projection of the world (a "reflection map") as seen from a particular viewpoint (the "world center") to make rendered surfaces appear to be reflecting their environment. The mirror reflection of the environment at a surface point is taken to be the point in the world projection corresponding to the direction of a ray from the eye as reflected by the surface. Consequently, reflections are geometrically accurate only if the surface point is at the world center or if the reflected object is greatly distant. The geometric distortion of reflections increases as the distance from the surface point to the world center increases and as the distance from the reflected object to the world center decreases. To apply reflection mapping to a particular object, the most satisfactory results are usually obtained by centering the world projection at the object center.

This method for approximating reflections can be extended to encompass refraction. Obtaining accurate results, however, requires much more computation since the ray from the eye should be "ray traced" through the refractive object, and in this process the ray usually splits into reflected and refracted components at surface intersections. As with reflections, results are only approximate for geometric reasons. (Note: Simply bending the ray at the surface point and using this as the direction of the refracted ray is not accurate, but may convey the impression of refraction.)

As Miller and Hoffman have described, the concept of reflection mapping may be thought of in more general terms as an illumination model. Essentially, they treat a world projection as an area light source which produces sharp reflections in smooth glossy objects and diffuse reflections in low gloss objects. This is a good model of illumination in the real world, although shadows are not explicitly handled and, as with reflection mapping, results are only approximate for geometric reasons. In order to speak generically of this approach and conventional reflection mapping, the term "environment mapping" will be applied in this paper to techniques for shading and texturing surfaces which employ a world projection.

2. STRENGTHS OF ENVIRONMENT MAPPING

Environment mapping is often thought of as a poor man's ray tracing, a way of obtaining approximate reflection effects at a fraction of the computational expense. While ray tracing is unquestionably the more versatile and comprehensive technique, handling shadows and multiple levels of reflection and refraction [Whitted], it is interesting to note that environment mapping is superior in some ways, quite aside from the enormous advantage in speed. Obtaining accurate diffuse reflection and antialiasing specular reflection are both problematical with ray tracing since it point samples the three dimensional environment (Amanatide's approach of ray tracing with cones is an exception). Environment mapping can effectively handle these problems by filtering regions of the world projection. Under many circumstances, for example when a large area light source illuminates a low gloss object, the subjective quality of reality cues produced by environment mapping are superior to those produced by unadorned ray tracing. While it is noted that refinements to ray tracing proposed by [Cook-Porter-Carpenter] and [Amanatides] address the problems of aliasing and diffuse reflection, their methods increase several-fold the already extreme computational cost. (Incidentally, there is an interesting parallel between the way these refinements work and the use of texture filtering in environment mapping: both attempt to integrate over a region of the environment.) It should also be mentioned that ray tracing and environment mapping can be used in combination where foreground objects are ray traced and a world projection representing distant objects serves as a backdrop (Lucasfilm's "1984" serves as an example [Cook-Porter-Carpenter, Figure 8]).

3. AN EXAMPLE

The dimetrodon lizard of Plate 1 was rendered with reflection mapping (i.e. environment mapping with mirror reflections). Inspection of this image reveals an inherent limitation with reflection mapping: since the reflecting object is not normally in the world projection the reflecting object cannot reflect parts of itself, e.g. the legs are not reflected in the body. (Actually, this limitation can be partially overcome by using different world projections for different parts of an object.) But overall, reflection mapping performs well in this scene. The horizon and sky, which are the most prominent reflected features, are accurately reflected due to their distance from the reflecting object. The reflections of the tree and the foreground terrain are less accurate because of their close proximity, but surface curvature makes this difficult to recognize. As this example shows, reflections don't need to be accurate to look realistic, although attention should be paid to scene composition. Planar reflecting surfaces may cause problems, for example, since the distortion of reflections in them may be quite noticeable.

4. RENDERING A CUBE PROJECTION

Use of environment mapping presupposes the ability to obtain a projection of the complete environment. Suppose that we wish to obtain a world projection of a three dimensional synthetic environment as seen from a particular viewpoint. One method is to position the camera at the viewpoint and project the scene onto a cube by rendering six perspective views, each with a 90 degree view angle and each looking down an axis of a coordinate system with its origin at the viewpoint. The world projection used to shade the lizard in Plate 1 was created in this manner and the resulting cube is shown unfolded in Plate 2. This method of creating world projections has also been used by Miller and Hoffman as an intermediate step in creating a Mercator projection of the world, the format they prefer for environment mapping. Blinn and Newell also use a Mercator projection.

5. SURFACE SHADING AND TEXTURE FILTERING

Light reflected by a surface is assumed to have a diffuse component and a specular component. The diffuse component represents light that is scattered equally in all directions, the specular component represents light that is reflected at or near the mirror direction. This discussion of surface shading will be confined to obtaining diffuse and specular illumination from a world projection; the problem of combining this information along with surface properties (color, glossiness, etc.) to obtain the color reflected at a surface point will not be considered. See [Cook and Torrance] for a general discussion of surface shading, [Miller and Hoffman] for a discussion of shading in the context of environment mapping.

Diffuse illumination at a surface point comes from the hemisphere of "sky" in the world projection centered on the surface normal, and it can be found by filtering the region of the world projection corresponding to this hemisphere. Filtering should be done according to Lambert's law which states that the illumination coming from a point on the hemisphere should be weighted by the cosine of the angle between the direction of that point and the surface normal.

A region of the world projection should also be filtered to obtain the specular illumination component. Figure 1 shows the cone of "sky" reflected by a surface which corresponds to a pixel in the output image. An obvious approach to determining the specular illumination at a pixel is to find the region of the world projection subtended by the corresponding "reflection cone" and then average the pixels in this region. While this approach will produce reasonable results, they will not be optimal for several reasons. One problem is that the size of the region filtered should be influenced by surface roughness

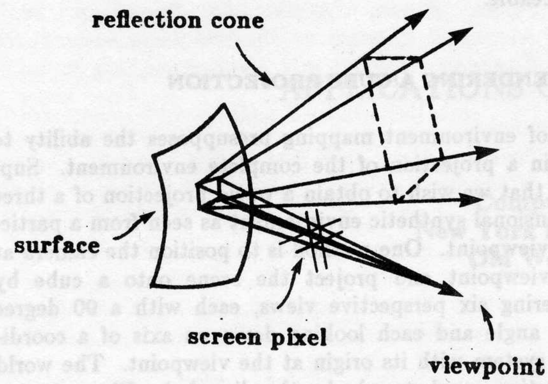

Figure 1. Rays from viewpoint through corners of a screen pixel are reflected by a surface, defining a reflection cone.

at the pixel level, since rough surfaces scatter specular illumination (thus, a larger area should be filtered if the surface is rough). Secondly, for theoretical reasons filtering should not be restricted to the region of texture corresponding to pixel bounds, and averaging of pixel values should be weighted by proximity to the center of the region being filtered. Subsequent references to reflection cones are made bearing in mind that they indicate only the approximate boundaries of regions to be filtered. See [Heckbert] for a general discussion of texture filtering.

While the details of shading formulas are beyond the scope of this paper, a rough rule for shading monochrome surfaces at a surface point is:

reflected color $=$ dc $*$ D $+$ sc $*$ S , where

dc is the coefficient of diffuse reflection
D is the diffuse illumination
sc is the coefficient of specular reflection
S is the specular illumination

(Note: dc and sc depend mainly on surface glossiness)

Plate 4 shows three monochrome spheres reflecting the environment of Plate 2. The relative weighting of the diffuse and specular components varies from completely diffuse on the left (a Lambertian reflector) to completely specular on the right (a perfect mirror).

6. CUBE PROJECTIONS VS. MERCATOR PROJECTIONS

When world projections are rendered from three dimensional models (rather than being photographed or painted), the cube representation of the world is preferred to a Mercator projection for reasons of computational efficiency and image quality. Rendering the cube projection is normally required in both cases. The further step of creating a Mercator projection from the cube projection requires additional computation, and the added generation of filtering can only degrade image clarity.

Moreover, Mercator projections are non-linear which complicates texture filtering since the region of texture subtended by a reflection cone does not have a polygonal boundary. (With cube projections, reflection cones always subtend polygonal regions of cube faces.) Figure 2 shows the regions subtended by the same reflection cone in Mercator and cube projections. Accurate filtering of the subtended region in the Mercator projection (the upper region) is problematical due to its irregular shape and the fact that pixels in the projection correspond to widely varying areas of "sky." Admittedly, filtering a cube projection presents a different problem: the multiplicity of regions to be filtered, five in the example of Figure 2. Fortunately, these problems occur only when surface curvature is high at the pixel level, which is not usually the case. Low surface curvature produces narrow reflection cones which map to approximately quadrilateral areas in a Mercator projection and which usually map to a quadrilateral on a single face of a cube projection.

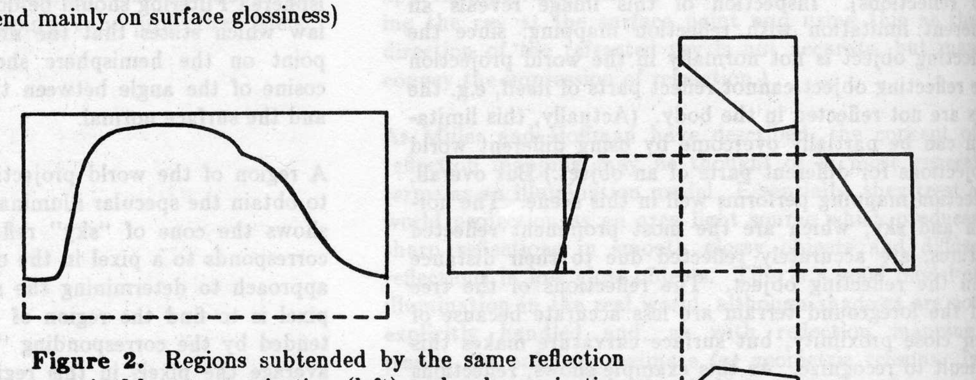

Figure 2. Regions subtended by the same reflection cone in Mercator projection (left) and cube projection (right). The reflection cone covers about 3/8 of the world.

7. PREFILTERED WORLD PROJECTIONS

As noted above, the width of a reflection cone depends on surface curvature so the area of "sky" reflected at a pixel can be arbitrarily large. A sphere covering a single pixel, for example, reflects nearly the entire world. Thus, environment mapping normally requires filtering large areas of the world projection at some pixels, even to produce reflections in mirror surfaces, so approximate filtering methods utilizing prefiltered texture greatly enhance efficiency.

As Miller and Hoffman describe, a diffuse illumination map can be created by convolving the world projection with a Lambert's law cosine function, a kernel covering one hemisphere of the world. This map, which is indexed by surface normal, may be thought of as indicating the colors reflected by a monochrome spherical Lambertian reflector placed at the world center. Since the diffuse illumination map has little high frequency content it may be computed and stored at low resolution and accessed with bilinear interpolation. Thus, prefiltering the world projection in this manner reduces the problem of finding the diffuse illumination at a surface point to a table lookup. Plate 3 is a magnified diffuse illumination map in Mercator projection format corresponding to the world projection of Plate 2. Incidentally, since diffuse illumination maps usually change only subtly from frame to frame, for animation applications it may suffice to create these maps at intervals (say every tenth frame) and interpolate between them.

Specular illumination can be obtained using fast, approximate filtering techniques developed for texture mapping surfaces. Our present implementation of environment mapping uses a prefiltered cube projection in the form of six "mip maps" [Williams], one for each cube face. Mip mapping is fast, but can only filter square regions of texture, so results are only approximate (in the example of Figure 2, the polygonal areas on the cube faces would need to be approximated by squares). [Crow] and [Perlin] have proposed similar approximate filtering techniques.

8. CHEAP CHROME EFFECTS

Environment mapping also has wide application where the objective is to produce a striking visual effect without particular regard for realism. Often the intent is to give objects a chrome plated look and the content of reflections is unimportant. In Robert Abel and Associate's "Sexy Robot" animation, for example, the reflection map was a smooth color gradation from earth colors at low elevations to sky colors at high elevations, and at a given elevation color was constant [Byles]. For applications of this sort, a colormap (or other one dimensional color table) suffices to specify the color gradation, and rendering computation can be greatly reduced by filtering this table instead of performing two dimensional texture filtering. Since an environment map isn't used, memory requirements are minimal and no setup time is required to prefilter texture. Highlights can be produced by adding point light sources independently of environment mapping.

9. MODELING BACKGROUND OBJECTS WITH WORLD PROJECTIONS

While the discussion thus far has been confined to using world projections for surface shading, they can also be used to model background objects. As described in [Greene], the sky component in frames of moving camera animation can be modeled as a half-world projection, for example the upper half of a cube. As the camera moves through the scene the appropriate region of the projection comes into view. Of course the advantage of this technique is speed: a scene element (in this case the sky) can be rendered from the world projection with texture mapping which, for complex scenes, is much faster than rendering from corresponding three dimensional models. This method assumes that objects in the sky are greatly distant from the camera, and results are only approximate when this is not the case. Plate 5 is a frame from animation in which the sky component was rendered from texture painted on a half-world projection.

Since half-world models cover the whole sky, they are particularly useful for creating world projections for environment mapping. The sky component of Plate 2 was modeled by projecting a 180 degree fisheye photograph of sky onto a half cube, shown in isolation in Plate 6. This model served two purposes in producing the picture of Plate 1, modeling the sky in the background and making the world projection which was used to shade the lizard. Incidentally, using photographic texture for sky models is an attractive option due to the difficulty of synthesizing complex sky scenes with realistic cloud forms and lighting effects, and using a 180 degree fisheye lens allows the whole sky to be photographed at once, thus avoiding problems associated with photo mosaics.

In scenes of animation where the camera rotates but doesn't change location, this approach to modeling sky can be extended to modeling the whole background. In this case, a single world projection centered at the camera is generated and the background can be rendered directly from this projection for the frames of animation. The geometry of the scene is faithfully reproduced regardless of the distance of objects in the scene from the camera; results are not just approximate. There are also limited applications of this technique to moving camera animation. A world projection of the distant environment centered at a typical camera position can be used to render distant objects in the scene while the near environment is rendered from models at each frame and then composited with the distant environment. This approach is analogous to using foreground and background levels in cel animation. As before, it is convenient to represent the world as a cube projection since it can be rendered directly from three dimensional models.

Rendering background objects from a world projection is a form of texture mapping since each pixel in the output image is rendered by determining the corresponding region in the world projection and then filtering a neighborhood of this region. Assuming that the world is represented as a cube projection, the cube faces should have substantially higher resolution than the output frames since only a small part of the world is subtended by the viewing pyramid for typical view angles. For example, a viewing pyramid with a 3:4 aspect ratio and a 45 degree vertical view angle subtends only about six percent of the world. Prefiltering the world projection is unnecessary since only small regions of texture are filtered.

This approach to rendering background objects suggests a method for performing motion blur. The region of texture traversed by the projection of an output pixel in the time interval between frames is determined, and then this region is filtered. The simplicity of performing accurate motion blur in this situation is due to the two dimensional nature of the model. The filter employed should be spatially variant because different output pixels may traverse differently shaped paths (this occurs, for example, if the camera rolls). Approximate results can be obtained by filtering elliptical regions in the world projection. Greene and Heckbert present an efficient method of filtering arbitrarily oriented elliptical areas.

10. NON-LINEAR PROJECTIONS

In addition to using world projections to generate perspective views of an environment, they can also be used to create non-linear projections such as the fisheye projection required for Omnimax frames. (The screen in an Omnimax theater is hemispherical and film frames are projected through a fisheye lens [Max].) Greene and Heckbert have obtained Omnimax projections of three dimensional scenes by projecting the scene onto a cube centered at the camera at each frame and then filtering regions of the cube faces to obtain pixels in the output image. This technique is very similar to the method described in the preceding section for rendering background objects from world projections. Plate 7 is an Omnimax projection made from the cube projection of Plate 2.

11. CONCLUSION

The projection of the environment onto a cube is a convenient and efficient format for world projections for the applications cited in this paper. In the context of a graphics system which employs world projections for multiple applications, the advantages of using a standard format are obvious: projection and filtering software is simplified, multiplicity of picture formats is avoided, etc. Generally, the methods described are ways of approximating a three dimensional problem with a two dimensional problem in order to reduce computational expense: environment mapping approximates ray tracing, rendering background objects from world projections with texture mapping approximates image rendering from three dimensional models. The subjective quality of reality cues in images produced with these approximate methods often compares favorably to results obtained by more expensive image generation techniques, and for complex environments approximate techniques may be the only practical way of producing animation having the desired features with moderate computing resources.

Acknowledgements

My thanks to Paul Heckbert who collaborated to incorporate environment mapping in his polygon renderer, who devised the "elliptical weighted average filter" used to create Omnimax projections, and who provided a critical reading of this paper. Thanks also to John Lewis, Tom Shermer, and John Schlag for their comments on the paper and to Dick Lundin who assembled and rendered the scene of Plate 1. Ariel Shaw and Rich Carter provided photographic assistance.

References

Amanatides, John, "Ray Tracing with Cones," *Computer Graphics,* vol. 18, no. 3, July 1984, pp. 129-135.

Blinn, James F., and Martin E. Newell, "Texture and Reflection in Computer Generated Images," *C.A.C.M.,* vol. 19, no. 10, Oct. 1976, pp. 542-547.

Byles, Torrey, "Displays on Display: Commercial Demonstrates State-of-the-Art Computer Animation," *IEEE Computer Graphics and Applications,* April 1985, pp. 9-12.

Cook, Robert L., Thomas Porter, and Loren Carpenter, "Distributed Ray Tracing," *Computer Graphics,* vol. 18, no. 3, July 1984, pp. 137-145.

Cook, Robert L., and Kenneth E. Torrance, "A Reflectance Model for Computer Graphics," *Computer Graphics,* vol. 15, no. 3, August 1981, pp. 307-316.

Crow, Franklin C., "Summed-Area Tables for Texture Mapping," *Computer Graphics,* vol. 18, no. 3, July 1984, pp. 207-212.

Greene, Ned, "A Method of Modeling Sky for Computer Animation," *Proceedings of the First International Conference on Engineering and Computer Graphics,* Beijing, China, Aug. 1984.

Greene, Ned, and Paul S. Heckbert, "Creating Raster Omnimax Images from Multiple Perspective Views using the Elliptical Weighted Average Filter," *IEEE Computer Graphics and Applications,* June 1986(?).

Heckbert, Paul S., "Survey of Texture Mapping," *Proceedings of Graphics Interface '86,* May 1986.

Max, Nelson, "Computer Graphics Distortion for Imax and Omnimax Projection," *Nicograph '83 Proceedings,* Dec. 1983, pp. 137-159.

Miller, Gene S., and C. Robert Hoffman, "Illumination and Reflection Maps: Simulated Objects in Simulated and Real Environments," *Siggraph '84 Advanced Computer Graphics Animation seminar notes,* July 1984.

Perlin, Kenneth, "Course Notes," *Siggraph '85 State of the Art in Image Synthesis seminar notes,* July 1985.

Whitted, Turner, "An Improved Illumination Model for Shaded Display," *C.A.C.M.,* vol. 23, no. 6, June 1980, pp. 343-349.

Williams, Lance, "Pyramidal Parametrics," *Computer Graphics,* vol. 17, no. 3, Aug. 1983, pp. 1-11.

Plate 1. Dimetrodon lizard model reflecting its environment. (Lizard modeled by Dick Lundin, tree modeled by Jules Bloomenthal)

Plate 5. Frame from animation utilizing a half-world sky model with painted texture. (Note: The reflection in the lake was *not* produced with reflection mapping.)

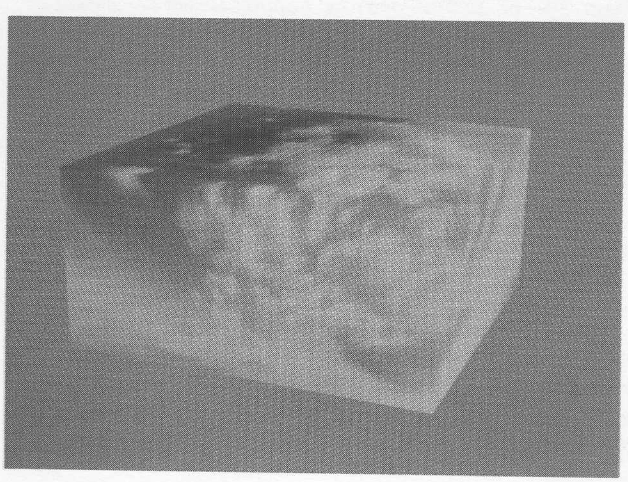

Plate 6. Half-world sky model made with photographic texture.

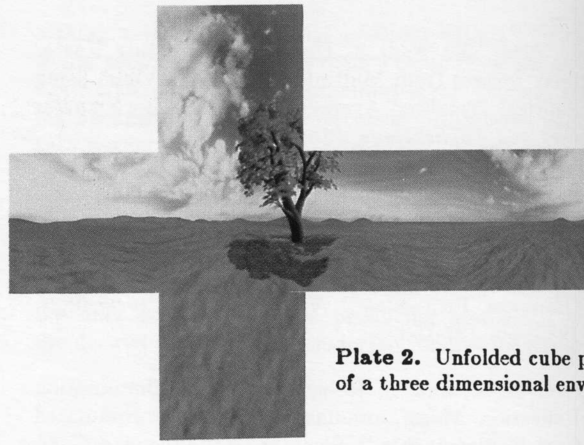

Plate 2. Unfolded cube projection of a three dimensional environment.

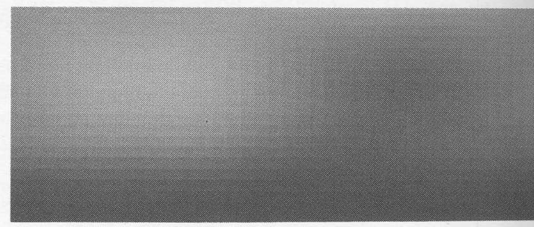

Plate 3. Diffuse illumination map of the environment of Plate 2 (Mercator format).

Plate 4. Spheres reflecting the environment of Plate 2. Lefthand sphere is a Lambertian reflector, righthand sphere is a perfect mirror.

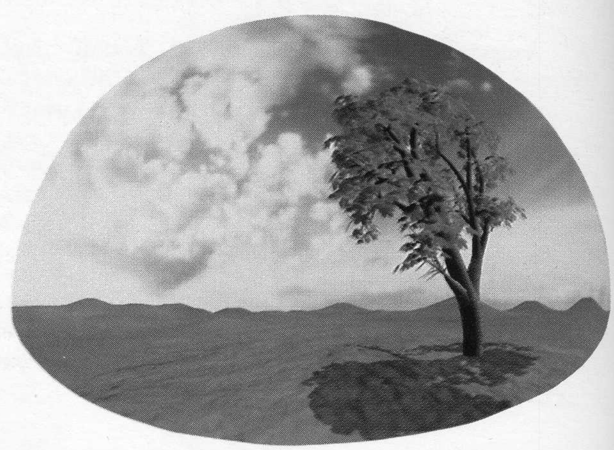

Plate 7. Omnimax projection of the environment of Plate 2.