

Cem Yuksel

Gradient Space Projection

Projections onto the Hemisphere

Abstract In this paper we present a simple method for handling projections on the hemisphere. The proposed method defines a coordinate system on the hemisphere surface that effectively converts the hemisphere to an infinite plane tangent to the top of the hemisphere, which is called the gradient space. Using this coordinate system, projections onto the hemisphere are replaced by simple perspective projections onto the gradient space. This approach totally eliminates the non-linearities caused by the spherical surface and permits exact hemisphere projection computations using only a single projection operation. We present different sampling techniques on the gradient space and provide qualitative comparisons.

Keywords Hemisphere Projection · Gradient Space

1 Introduction

Projecting polygons onto the hemisphere is a fundamental operation performed by many graphics algorithms such as visibility computations and radiosity. Different approaches have been proposed to make this operation accurate or efficient both by directly using the hemisphere or approximating its shape by simpler structures.

In this paper we present a different approach to handle projections onto the hemisphere by effectively converting it to an infinite plane (*gradient space*). The main difference of our approach is the use of an infinite plane as opposed to a finite plane that represents the full hemisphere or a part of it. Using projections onto this infinite plane instead of the hemisphere, the linearity of polygons are preserved avoiding non-linearities caused by the spherical shape of the hemisphere. Furthermore, the entire hemisphere surface can be addressed without dividing it into pieces and handling them separately. Thus, a single perspective projection operation is enough to cover the full hemisphere.

Cem Yuksel
 Department of Computer Science, Texas A&M University
 E-mail: cem@cemyuksel.com

Because of this infinite plane, we cannot use a simple grid based sampling method. Therefore, we also propose an efficient sampling scheme that works on the infinite plane. Our results show that our method is efficient and produces less noise.

In the next section we briefly overview the previous work. Section 3 explains our method in detail and Section 4 discusses sampling on the infinite plane. We provide our results in Section 5 and conclude in Section 6.

2 Previous Work

The most common way of computing projections onto the hemisphere is using a hemicube (half cube) structure [1]. Even though the shape of the hemicube can only roughly approximate the shape of the hemisphere, in practice good results can be achieved by computing appropriate weights for each sample position on the hemicube. On the other hand, a separate projection operation is required for each face of the hemicube.

Sillion and Puech [11] used a single finite plane that is larger than the top of the hemicube to compute form factors for radiosity computation. They reported that the single plane produced better results with higher performance. However, their finite plane cannot address the full hemisphere and the illumination from near horizontal directions is totally ignored. Single pass projections are also used to compute environment maps [5] and radiosity [2] on graphics hardware.

Spencer [12] proposed a method for directly using the hemisphere surface, which is accelerated by Doi and Itoh [3]. Strzlinger [13] used exact projections directly on the hemisphere surface.

3 Gradient Space Projection

When lines of polygons are projected onto the hemisphere surface, they form circular arcs. To simplify this

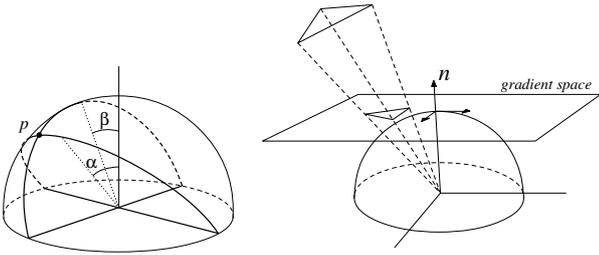


Fig. 1 α - β coordinates (left), gradient space (right).

and use simple linear equations instead, we define a coordinate system on the hemisphere surface that converts all these arcs back to line segments.

We start by picking an arbitrary cartesian coordinate system such that z-axis is aligned with the surface normal. We represent the coordinate of a point p on the hemisphere that is aligned with the z-axis (surface normal) using two angles α and β as shown in Figure 1. Here, α is the angle between the z-axis and the plane that contains the point p and the y-axis, and β is the angle between the z-axis and the plane that contains the point p and the x-axis. Using α - β coordinates we can uniquely represent all points on the hemisphere surface.

Then we define a u - v coordinate system such that $u = \tan(\alpha)$ and $v = \tan(\beta)$. This u - v coordinate system corresponds to the projection of the hemisphere on the infinite plane tangent to the top of the hemisphere. This infinite plane is called the gradient space of the hemisphere [6].

Using this coordinate system, projections onto the hemisphere become simple perspective projections onto the infinite gradient space (Figure 1), where the center of projection is the center of the hemisphere. This simple perspective projection allows us to use low-complexity z-buffer projection algorithms like [4, 8, 14, 10, 16] for reducing the number of target polygons to be projected.

In practice, it is not possible to use an infinite plane on a computer system, and the size of this plane is bounded by the limits of the floating point representation. However, using the exponential representation of floating point numbers of IEEE standard 754 that is used in most computers today, we can address the part of the hemisphere down to the horizontal angles of 10^{-37} degrees with single precision (32-bit). Using double precision (64-bit), this limit falls below 10^{-307} degrees. Therefore, for any practical purpose, we can safely assume that the full hemisphere is addressed.

On the other hand, in this exponent form, as we get closer to the limits of the floating point representation, the accuracy of the numbers drops significantly. However, this does not affect the accuracy of our computation, which is defined by the accuracy on the hemisphere, not on the gradient space. Even when the difference between two consecutive floating point numbers becomes large towards the limits of the gradient space, the corresponding difference on the hemisphere is still very small.

4 Sampling

Since we are using an infinite plane, we cannot sample it using a simple grid that distributes the sampling points evenly over the plane, as the number of sampling points would be infinite regardless of the resolution.

A simple alternative is to place sampling points randomly or using a quasi-random sequence; thus, the result of the sampling corresponds to Monte Carlo ray tracing. In this case, sample positions are first computed on the hemisphere and then projected onto the gradient space. Even though random sample distributions are known to produce good results without aliasing artifacts, an ordered distribution of sampling points is preferred for projection based methods, since this order can be used to compute intersections of a projected polygon with multiple sampling points at once.

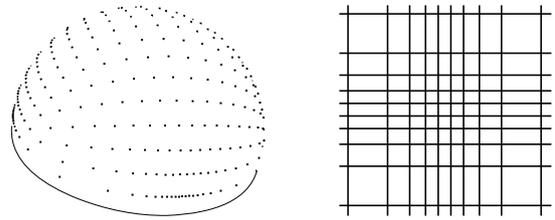


Fig. 2 Sampling points of Sillion and Puech on the hemisphere, and on the gradient space as intersections of lines.

One way to achieve ordered sample point distribution is to extend the finite plane sampling scheme introduced by Sillion and Puech [11] to an infinite plane. Sillion and Puech used the intersections of horizontal and vertical lines as shown in Figure 2. Lines are placed such that $u_i = \tan(i\Delta\alpha)$ and $v_j = \tan(j\Delta\beta)$, where u_i and v_j are the positions of i^{th} horizontal and j^{th} vertical lines respectively, and $\Delta\alpha$ and $\Delta\beta$ are constants. Even though the sample positions on the gradient space are well ordered in this scheme, sample positions on the hemisphere form over-sampled and under-sampled areas towards the horizontal plane as can be seen in Figure 2.

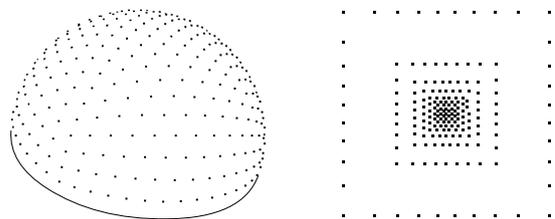


Fig. 3 Sampling points of a geodesic dome on the hemisphere, and on the gradient space.

We propose a different sampling scheme, which is used in ray tracing for uniform sampling of the hemisphere. We use the nodes of a geodesic dome as our sam-

pling points, so that we make sure the sampling points are almost evenly distributed over the hemisphere surface and our sampling is unbiased. We used a geodesic dome that is produced by subdividing a square pyramid, because it forms an ordered distribution of samples on the gradient space (Figure 3). Note that nodes that are on the horizon of the hemisphere are discarded.

To determine the sampling points underlying a projected polygon efficiently, we separate the sampling points into two groups: points that are on horizontal lines and points that are on vertical lines as shown in Figure 4. While sampling a projected polygon, we first determine the lines that the polygon intersects. Then, we find the two intersection points of each line with the polygon edges. All sampling points on the line within the two intersection points intersect with the polygon.

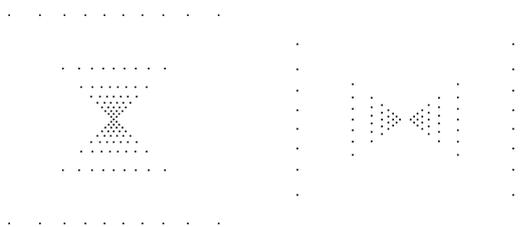


Fig. 4 Samples on horizontal (left) and vertical (right) lines.

5 Results

We implemented our gradient space projection method for gathering illumination from polygonal meshes. We project each polygon onto the gradient space and find intersections with predefined sampling points. Occlusion test is handled by comparing the distance of the polygon to the previously stored value on the sampling point similar to a z-buffer algorithm. Once all the target polygons are projected and sampled (their intersections with the sampling points are computed), we gather the incoming radiance from all the sampling point directions. Note that for all sampling methods we use an infinite plane (gradient space). To avoid aliasing artifacts, we apply a random rotation to the hemisphere along the surface normal prior to the projection operation as in [15] and [9].

Figure 5 shows a simple Cornell box scene for a visual comparison of different sampling methods. For random samples we used a Poisson disc distribution, for it is known to produce the best sampling distribution. As can be seen from these images, noise decreases with the increasing number of sampling points for all methods. However, noise on the objects does not decrease much for the sampling scheme of Sillion and Puech, because the sampling points are not evenly distributed on the surface of the hemisphere. Here, the quality of geodesic

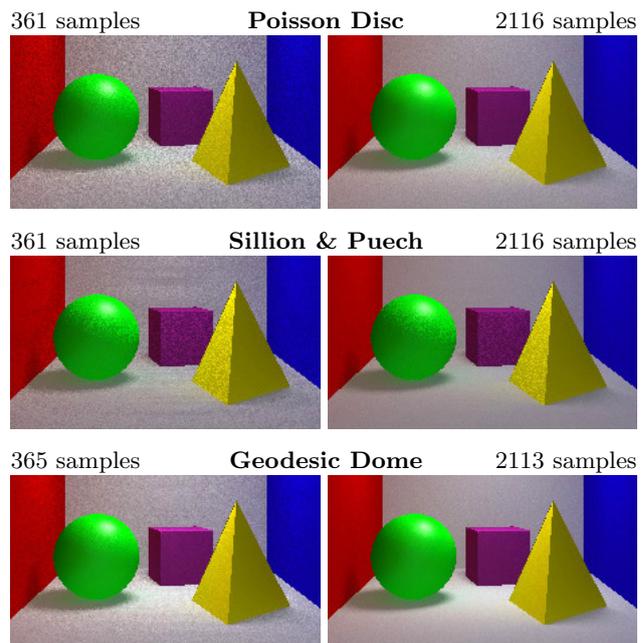


Fig. 5 Cornell box scene comparing sampling techniques.

dome sampling is similar to poisson disc sampling. Note that replacing the projection operation with ray tracing using the same samples as ray directions would produce exactly the same results.

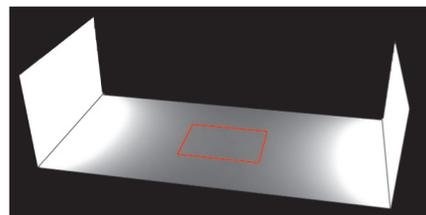


Fig. 6 A simple scene for testing sampling accuracy.

To quantify the differences, we prepared the scene in Figure 6, in which the two planes on either side of the horizontal plane are assigned a constant illumination. The camera is placed such that the frame line is the red rectangle in the middle of the horizontal plane. Table 1 shows the standard deviation of intensities in the rendered images. These results show that Poisson disc and geodesic dome sampling produce less noise.

Table 1 Standard deviation of intensities

	145	365	1301	2113	samples
Random	22.4	14.2	7.4	5.9	
Poisson Disc	13.8	8.3	4.2	3.3	
Sillion & Puech	34.3	25.1	25.9	25.5	
Geodesic Dome	9.8	6.0	4.9	4.0	

Table 2 Sampling times in milliseconds

	145	685	1301	2113	samples
Poisson Disc	0.16	0.95	1.82	3.26	
Sillion & Puech	0.09	0.10	0.11	0.27	
Geodesic Dome	0.10	0.11	0.13	0.25	

The average sampling times for the scene in Figure 5 are shown in Table 2. Here we can see that well-ordered sampling point distribution in geodesic dome and Sillion and Puech sampling methods helps reduce the sampling time significantly. These results show that we can get fast results with low noise using geodesic dome sampling.

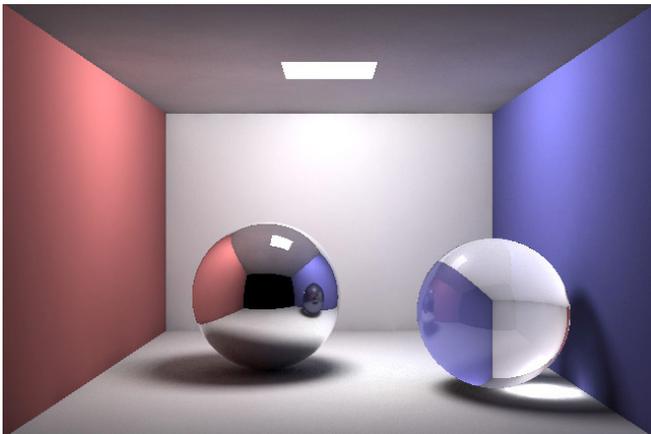
**Fig. 7** Gradient space projection with photon mapping.

Figure 7 shows another Cornell box scene rendered using photon mapping [7] and gradient space projection with geodesic dome sampling for final gathering. Figure 8 shows a room scene illuminated by a single light source and indirect illumination with single bounce using the same method without photon mapping.

**Fig. 8** A room scene illuminated by a single light source.

6 Conclusion

We introduced the gradient space projection method for handling projections onto the hemisphere. Using this approach exact projections can be computed with a single perspective projection operation that preserves the linearity of polygons. We also show sampling techniques on the infinite gradient space and provide qualitative and performance comparisons. As future work, we would like to explore possible GPU-based implementations (note that the geodesic dome sampling may not be suitable for such applications).

References

1. Cohen, M.F., Greenberg, D.P.: The hemi-cube: A radiosity solution for complex environments. In: Proceedings of SIGGRAPH 1985, pp. 31–40 (1985)
2. Coombe, G., Harris, M.J., Lastra, A.: Radiosity on graphics hardware. In: GI '04: Proceedings of Graphics Interface 2004, pp. 161–168 (2004)
3. Doi, A., Itoh, T.: Acceleration radiosity solutions through the use of hemisphere-base formfactor calculation. *Journal of Visualization and Computer Animation* **9**(1), 3–15 (1998)
4. Greene, N., Kass, M., Miller, G.: Hierarchical Z-buffer visibility. In: Proceedings of SIGGRAPH 1993, vol. 27, pp. 231–238 (1993)
5. Heidrich, W., Seidel, H.P.: View-independent environment maps. In: HWS '98: Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware, pp. 39–ff. (1998)
6. Horn, B.K.P.: *Robot Vision*. The MIT Press (1986)
7. Jensen, H.W.: *Global Illumination Using Photon Maps*. In: *Rendering Techniques '96*, pp. 21–30 (1996)
8. Lafortune, E.P., Willems, Y.D.: A 5D Tree to Reduce the Variance of Monte Carlo Ray Tracing. In: *Rendering Techniques '95*, pp. 11–20 (1995)
9. Max, N.: Optimal sampling for hemicubes. *IEEE Transactions on Visualization and Computer Graphics* **1**(1), 60–76 (1995)
10. Rusinkiewicz, S., Levoy, M.: Qsplat: A multiresolution point rendering system for large meshes. In: Proceedings of SIGGRAPH 2000, pp. 343–352 (2000)
11. Sillion, F., Puech, C.: A general two-pass method integrating specular and diffuse reflection. In: Proceedings of SIGGRAPH 1989, pp. 85–92 (1988)
12. Spencer, S.: The hemisphere radiosity method: a tale of two algorithms. *Photorealism in Computer Graphics* pp. 127–135 (1991)
13. Strzlinger, W.: Exact projections onto the hemisphere. In: *Spring Conference on Computer Graphics '95*, pp. 31–38
14. Sudarsky, O., Gotsman, C.: Output-sensitive visibility algorithms for dynamic scenes with applications to virtual reality. *Computer Graphics Forum* **15**(3), 249–258 (1996)
15. Wallace, J.R., Cohen, M.F., Greenberg, D.P.: A two-pass solution to the rendering equation: A synthesis of ray tracing and radiosity methods. In: Proceedings of SIGGRAPH '87, pp. 311–320 (1987)
16. Wand, M., Fischer, M., Peter, I., auf der Heide, F.M., Straßer, W.: The randomized z-buffer algorithm: Interactive rendering of highly complex scenes. In: Proceedings of SIGGRAPH 2001, pp. 361–370 (2001)