Ian Ashdown, P. Eng., LC, FIES
620 Ballantree Road
West Vancouver, BC
Canada V7S 1W3

Tel: (604) 922-6148
Fax: (604) 987-7621
e-mail: ian@agi32.com

September 25th, 2004

Modeling Daylight for Interior Environments

by

Ian Ashdown, FIES
Senior Software Engineer
Lighting Analysts Inc.

Abstract

This paper[1] presents a fast and deterministic algorithm for daylight simulation and analysis of complex interior environments that models the transfer of luminous flux from exterior environments through arbitrarily-shaped and oriented windows and openings. These are modeled as arrays of independent point sources with luminous intensity distributions that are determined by the distribution of direct sunlight, diffuse daylight in accordance with analytic CIE and IESNA sky models, ground plane reflections, occlusions, and interreflections due to complex objects in the exterior environment. The algorithm includes Fresnel reflections from glass, and is applicable to both planar and non-planar windows such as barrel vault skylights.

Introduction

Daylight simulation and analysis with lighting design software generally involves an undesirable tradeoff between accuracy and calculation speed. Ray tracing programs such as *Radiance* (e.g., Larson and Shakespeare 1996) are capable of producing accurate and detailed photometric predictions for complex environments. However, they require a considerable amount of computation time. Radiative transfer programs such as *Lumen Micro* (e.g., Jongewaard 1993) produce useful predictions, but they are limited in their ability to model daylighting in complex environments.

In this paper we present a radiative transfer algorithm for daylight simulation and analysis of complex architectural environments, including the transfer of luminous flux from exterior environments through arbitrarily-shaped and oriented windows and openings (which are collectively referred to as "transition" surfaces). It combines the speed of radiative transfer calculations with the accuracy of ray tracing techniques. More important, it includes both direct sunlight and diffuse daylight in its calculations, and it can successfully model arbitrary sky conditions.

Previous Work

Daylight simulation algorithms have a long history in illumination engineering and related fields. Analytic solutions for simple environments include Moon and Spencer (1942) and Spencer and Stakutis (1951). Ray

---

[1] A preliminary version of this paper, called "Fast Daylight Simulation and Analysis," was presented at the Illuminating Engineering Society of North America Annual Conference in Salt Lake City, Utah in 2002. At that time the algorithm for transferring luminous flux through windows and openings had not been developed.

tracing solutions include Nishita and Nakamae (1986), McHugh (1995), Tsangrassoulis and Santaouris (1997), and Paule et al. (1998). Mardaljevic (1995) used *Radiance* for his daylighting validation studies.

Modest (1983) proposed a radiative transfer solution that became the basis for the SUPERLITE daylight simulation program. Similarly, the commercial product *Lumen Micro* (Jongewaard 1993) performs daylight simulation analysis using radiative transfer techniques.

Unfortunately, the development of radiative transfer algorithms specifically designed for daylight simulation has not kept pace with the development of more general algorithms for electric lighting simulation and analysis (e.g., Cohen and Wallace 1993, Sillion and Puech 1994). While there have been numerous papers on the topic, including Languenou and Tellier (1992), Tadamura et al. (1993), Dobashi et al. (1994), Nimeroff et al. (1994), Müller et al. (1995), Nakamae et al. (1995), and Dobashi et al. (1996), the proposed algorithms have not proved particularly useful for commercial daylight simulation and analysis programs where execution times are a primary concern.

Radiative Transfer

Radiative transfer techniques depend on the efficient and accurate calculation of occluded form factors. There are two popular approaches for these calculations: hemicubes and ray casting (e.g., Cohen and Wallace 1993).

The hemicube approach is conceptually simple and computationally efficient. As shown in Figure 1, an imaginary hemicube (half of a cube) is positioned over a planar polygonal element **A** and all other elements (e.g., **B**) in the environment are projected onto the five faces of the hemicube surfaces using perspective projection (e.g., Hill 2001). This is the geometric equivalent of Nusselt's analogy (Hermann 1900) for configuration factor determination[2], where the hemicube replaces the hemisphere.

To better understand the hemicube approach, Figure 2 shows three elements in an environment projected onto the top face as seen from the center of the element under the hemicube. The face is subdivided into an array of cells for which each form factor has been precalculated. (See for example Ashdown 1994 for calculation details.) Summing the form factors of the individual cells covered by each element provides its approximate (and possibly occluded) form factor.

Occlusion determination is an important component of the hemicube approach. All hemicube face cells are initially marked *empty*. As each polygonal element is projected onto a hemicube face, the distance from the center of the hemicube to the element as seen through each cell is determined. If the cell is empty, the element identifier and its distance are assigned to the cell and the cell is marked as *occupied*.

If the cell is already occupied, the two distances are compared. If the new distance is less, it means that the current element is closer and so occludes the previous element as seen through the cell. The current element identifier and its distance are therefore assigned to the cell. When all of the elements in the environment have been projected onto the hemicube faces, each cell records the identifier and distance of the closest element as seen from the center of the hemicube through the cell.

The hemicube is particularly elegant in that it can take advantage of scanline conversion techniques (e.g., Hill 2001) to perform the occluded form factor calculations of all projected elements with minimal computational effort. Even better, these calculations can be performed in graphics hardware using industry-standard display cards (Rushmeier et al. 1990).

The disadvantage of the hemicube approach is that it is prone to aliasing. As shown in Figure 2, small or distant elements may only project onto a few cells. Consequently, their form factors may be over- or underestimated. This can lead to both radiative transfer calculations errors and surface "mottling" effects when the results are used to visualize the environment.

Ray casting techniques (e.g., Wallace et al. 1989) avoid this problem by randomly casting rays from the center of the element to every other element in the environment. Counting the number of rays within a given solid angle that intersect each element without being occluded provides a Monte Carlo estimate of its form factor.

---

[2] The hemicube approach models the polygonal element under the hemicube as a differential element, and so it is really *configuration* factors that are calculated. However, because this term rarely appears in the computer graphics literature, the term "form factor" is used here instead.

The disadvantage of ray cast form factor determination is that it cannot take advantage of scanline conversion techniques or graphics hardware acceleration. It relies instead on ray tracing techniques, and is typically several orders of magnitude slower than hemicube form factor determination for complex environments with thousands to hundreds of thousands of elements.

Worse, "brute force" ray tracing techniques have a time complexity of $\mathbf{O}(n^3)$. That is, the execution time increases according to the cube of the number of elements n in the environment. By comparison, hemicube techniques have a time complexity of $\mathbf{O}(n^2)$, so that the execution time increases according to the square of the number of elements. (See for example Cohen and Wallace 1993 for a detailed explanation.)

A Ray Casting Approach

Müller et al. (1995) proposed a radiative approach for daylight simulation wherein the sky hemisphere (or "dome") surrounding an architectural environment is divided into a set of "sky patches" (Figure 3). The luminous flux due to the sun or sky patch luminance is then transferred (or "shot") to each element in the environment. Following Tregenza (1987), the sky dome is divided into 145 approximately equal area patches for sky luminance measurements, using 12-degree vertical increments[3], while the sky patch luminance values are determined in accordance with either IESNA (2000) or CIE (2003).

To ensure that the luminous flux arrives from the same direction for each element, the sky dome radius must be at least sixty times larger than the radius of a circle enclosing the architectural environment. This limits the divergence of the luminous flux direction to less than one degree.

A hemicube cell typically subtends on the order of 0.1 degrees when projected into an environment. The authors therefore concluded that a hemicube approach would produce unacceptable aliasing, and instead chose ray casting for their form factor determination. They indicated that it took about three hours and some 70 million rays to calculate the form factors of an environment consisting of 47,500 elements. (The machine was a Silicon Graphics Onyx Reality Engine 2 workstation, which is roughly comparable to commodity desktop computers today.)

This is the same problem that limits the performance of ray cast form factor determination in comparison with the hemicube method. To avoid this problem, we need to reconsider the use of hemicubes in terms of daylight simulation.

A Parallel Projection Approach

The hemicube was originally developed to model radiative transfer between surfaces with ideal diffuse reflection. As such, it is not the best choice to model direct sunlight or sky patches, which are essentially directional sources.

With this in mind, we can replace the hemicube with a single parallel projection plane, where the orientation of the plane is perpendicular to the daylight direction (Figure 4). Each element in the environment simply casts a shadow on the projection plane.

Single projection planes with both parallel and perspective projection have been previously proposed for form factor determination (e.g., Recker et al. 1990, Neumann 1995, and DiLaura et al. 1998), but only in the context of radiative transfer between surfaces with ideal diffuse reflectance properties and light sources at finite distances.

The single projection plane offers several advantages. First and foremost, the direct illuminance of a projection plane is constant across the plane. (The equations for determining illuminance due to solar illuminance or sky patch luminance are presented in Chapter 8 of IESNA 2000 and CIE 2003.)

Referring to Figure 4, the projection plane is placed immediately behind the environment in the direction of the sun or a sky patch. Every element is then projected onto the plane using parallel projection, with the plane being subdivided into an array of cells. Similar to perspective projection of an environment onto a hemicube face, this projection automatically performs occlusion determination. Knowing the direct

[3] Equal-area sky patches based on 12-degree vertical increments are also used in the BRE-IDMP data (e.g., Mardaljevic 2001). In terms of light sources, each sky patch complies with the IESNA five-times rule for far-field photometry in that this rule implies a maximum subtended angle of 0.2 radians, or 11.5 degrees for an extended area light source.

illuminance of each cell, the luminous flux received by each element can then be determined simply by summing the number of cells covered by the visible portions of the element.

This parallel projection approach works equally well for both direct sunlight and diffuse daylight. Even with 145 sky patches, it generally requires less than four seconds to compute the complete daylight component for a complex architectural environment.

A second advantage is that the parallel projection approach is scalable to environments of any size and complexity. Given for example an office building with both large and small elements (such as walls and window mullions), any number of projection planes can be tiled, with each plane covering a portion of the environment. This limits the amount of memory that must be allocated for the array of projection plane cells.

Aliasing Issues

The parallel projection approach has the same disadvantage as the hemicube approach: it is prone to aliasing. Fortunately, there is a simple modification that greatly alleviates the problem. It is a hybrid approach that provides both the accuracy of ray cast form factors and the $\mathbf{O}(n^2)$ time complexity of the parallel projection approach.

In this approach, the elements are first projected onto the projection plane. If an element's projected area (that is, the number of cells onto which it projects) exceeds a predetermined threshold, the direct illuminance of the element is calculated as outlined above. Otherwise, the element is marked for ray tracing.

Once all of the elements in the environment have been processed using parallel projection, a ray normal to the projection plane is traced through the center of each marked element. If the ray intersects an empty projection plane cell or if the occupied cell element is behind the marked element, the marked element is assumed to be fully visible. Its direct illuminance can then be calculated according to the cosine of the angle between its normal and the ray direction.

Despite using ray tracing techniques, this hybrid approach has $\mathbf{O}(n^2)$ time complexity. Moreover, the ray tracing process itself has only $\mathbf{O}(m)$ time complexity, where m is the number of marked elements. Even with complex environments consisting of tens of thousands of elements, the hybrid approach typically increases the execution time by only a few percent.

This technique can also be used to alleviate hemicube aliasing. Figure 5a illustrates the distribution of direct luminous flux on a horizontal surface from an IESNA Type II roadway luminaire, using conventional hemicube-based rendering techniques. The surface has a large number of elements, which results in form factor calculation errors and consequent surface mottling artifacts.

Figure 5b illustrates the same distribution using the proposed hybrid hemicube rendering technique. The surface mottling artifacts are almost completely eliminated, indicating greater accuracy in the form factor calculations. The difference in calculation times for the two renderings was less than 10 milliseconds.

Windows and Openings

When a preliminary version of this paper was presented in 2002, it was intended for the above algorithm to be applied to exterior environments only. It was known to be ill suited for transferring luminous flux through windows and openings (transition surfaces) because each sky patch projects a parallel beam of light. If these beams are projected for example through a narrow window into a room with no other lighting (and in particular no direct sunlight), the result is a pattern of light "spokes" (Figure 6A) that are projected on the floor and walls. (A similar problem was identified by DiLaura et al. [1998] for occluded form factor determination.)

What is clearly needed is some means of transferring the directional luminous flux through transition surfaces into interior environments without the coarse discretization due to the limited number of sky patches. The solution to this problem is once again the hemicube.

Ashdown (1998) proposed using the hemicube as a virtual photometer by reversing the flow of light. Whereas a hemicube in radiative transfer calculations is used to "shoot" luminous flux from one surface element to all other visible elements in an environment, it can also be used to "gather" luminous flux from all visible surfaces to a point in space. As such, it is analogous to an illuminance meter in a physical environment.

The virtual photometer integrates the luminous flux it receives and so discards any directional information. However, this information is still available. We can therefore position a hemicube on the surface of a window (or the imaginary plane of an opening) such that it gathers luminous flux from the visible surfaces of an exterior environment. Having done so, the hemicube can be reversed on its vertical axis and the flux it has gathered shot into the interior environment. The hemicube discretization is quite fine (two million cells is typical), and so there is essentially no possibility of visible spokes as with the previous approach.

An advantage of this technique is that transmission losses due to window absorption (including that due to colored glass) and Fresnel reflections (e.g., Schlick 1993) can be easily incorporated into the hemicube calculations on a per-cell basis as the gathered flux is shot into the environment. (Fresnel reflection losses are particularly important for accurate daylight insolation studies. By shifting from visible to infrared wavelengths, the technique can also be applied to thermal engineering studies.)

A hemicube models the gathering and shooting of luminous flux from a single point in space. To accurately model the transfer of luminous flux through a transition surface, it is necessary to discretize the (real or imaginary) surface into an array of surface elements and transfer flux through each element. This is equivalent to how opaque and translucent surfaces are modeled in radiative transfer programs.

This technique requires that all of the exterior environment surfaces visible to each hemicube have already received the direct luminous flux from the sun and sky patches, and also all interreflected flux between these surfaces. It is therefore necessary to calculate a radiosity solution for the exterior environment surfaces before transferring the luminous flux through the transition surfaces to the interior environment surfaces. In this sense, each hemicube is then equivalent to a directional light source when it shoots its gathered luminous flux to the interior surfaces.

The primary advantage of this technique is that, as shown in Figure 6B, it completely eliminates the spoking artifacts that are evident in Figure 6A. Another advantage relates to “light leaks” through surface joins. Computer-aided drafting (CAD) models represent surface vertices using floating point numbers. As these vertices are translated and rotated during radiative transfer calculations, it is possible to have very small cracks appear between adjacent surfaces (such as a wall and ceiling) due to floating-point roundoff.

The amount of light that can “leak” through such cracks is proportional to the crack width, which is typically equivalent to less than a millimeter. If however the light source is direct sunlight and the interior environment is illuminated by diffuse daylight through a small window, the light leak can become painfully evident in both renderings and isolux plots.

This problem typically requires the user to manually add additional dummy surfaces to prevent such light leaks. For complex environments with lengthy radiosity calculation times, this can be a time-consuming iterative process. With the above technique however, exterior surfaces are illuminated separately from interior surfaces. The only exception is direct sunlight, which is allowed to illuminate interior surfaces only after passing through a transition surface. There is no possibility of direct sunlight illuminating an interior surface through cracks due to floating-point roundoff errors, and so the light leak problem is automatically and completely eliminated.

Virtual Ground Plane

One more issue needs to be considered: reflections from the exterior ground plane. In many instances, it is inconvenient to model a ground plane to reflect direct sunlight and diffuse daylight into an interior environment through transition surfaces. To solve this problem, *Radiance* (Larson and Shakespeare 1996) introduced a virtual ground plane that was modeled as an inverted sky dome with uniform luminance.

The same approach can be applied here. The unobstructed horizontal illuminance can be analytically calculated in accordance with IESNA (2000), and the equivalent luminance value (based on an assumed ground plane reflectance of 18 percent) is applied to each patch of the inverted sky dome. All other calculations then proceed as above. If there are exterior surfaces present, they can be made “double-sided” to block luminous flux received from the ground plane as they reflect luminous flux received from direct sunlight, diffuse sky patches, and other exterior surface elements.

Conclusions

Daylight simulation and analysis has previously required ray-tracing techniques to obtain accurate predictions for complex architectural environments. Unfortunately, the amount of computation time required can be considerable and so interactive design and analysis is discouraged.

We have shown in this paper that daylight simulation and analysis of complex architectural environments, both interior and exterior, can be accomplished in a matter of seconds using radiative transfer techniques. These calculations include direct sunlight and diffuse daylight, and also the transfer of luminous flux to interior environments through arbitrarily-shaped and oriented windows and openings.

When the preliminary version of this paper was presented in 2002, we concluded that, "All that is needed are commercial lighting design and analysis programs to take full advantage of the capabilities [of the proposed algorithm] to provide fast and interactive daylight simulation and analysis." With this paper, we are pleased to say that the above algorithms and techniques have been successfully implemented in just such a program. It is only modesty that prevents us from saying which one.

References

1. Ashdown, I. 1994. *Radiosity: A Programmer's Perspective.* New York, NY: John Wiley & Sons.
2. Ashdown, I. 1998. "The Virtual Photometer: Modeling the Flow of Light," *IESNA 1998 Annual Conference Technical Papers*.
3. Cohen, M. F., and J. R. Wallace. 1994. *Radiosity and Realistic Image Synthesis*. San Francisco, CA: Morgan Kaufmann.
4. CIE. 2003. *Spatial Distribution of Daylight – CIE Standard General Sky*. CIE Publication S 011/E:2003. Vienna, Austria: CIE Central Bureau.
5. DiLaura, D. L., S. R. Santoro, T. F. Miller, and G. S. Terch. 1998. "Nondiffuse Radiative Transfer 5: Calculation of Occluded Form Factors," *Journal of the Illuminating Engineering Society* 27(2):77-82.
6. Dobashi, Y., K. Kaneda, H. Nakashima, H. Yamashita, T. Nishita, and K. Tadamura. 1994. "Skylight for Interior Lighting Design," *Computer Graphics Forum* 13(3):C85-C96.
7. Dobashi, Y., K. Kaneda, H. Yamashita, and T. Nishita. 1996. "Method of Calculation of Sky Light Luminance Aiming at an Interactive Architectural Design," *Computer Graphics Forum* 15(3):C109-C118.
8. Hermann, R. A. 1900. *Treatise of Geometric Optics*. Cambridge, MA: Cambridge University Press, pp. 209-211.
9. Hill. F. S., Jr. 2001. *Computer Graphics Using OpenGL*. Upper Saddle River, NJ: Prentice Hall, pp. 713-715.
10. IESNA. 2000. *The IESNA Lighting Handbook*, Ninth Edition. New York, NY: Illuminating Engineering Society of North America.
11. Jongewaard, M. P. 1993. "Daylight Calculations, Measurements and Visualization in Non-Empty Rooms," *Proceedings of Lux Europa 1993*, pp. 43–52.
12. Languenou, E., and P. Tellier.1992. "Including Physical Light Sources and Daylight in Global Illumination," *Proceedings of the Third Eurographics Workshop on Rendering*, pp. 217-225.
13. Larson, G. W., and R. Shakespeare. 1998. *Rendering with Radiance: The Art and Science of Lighting Visualization*. San Francisco, CA: Morgan Kaufmann.
14. Mardaljevic, J. 1995. "Validation of a Lighting Simulation Program under Real Sky Conditions," *International Journal of Lighting Research & Technology* 27(4):181-188.
15. Mardaljevic, J. 2001. "The BRE-IDMP Dataset: A New Benchmark for the Validation of Illuminance Prediction Techniques," *International Journal of Lighting Research & Technology* 33(2):117-136.
16. McHugh, J. R. 1995. *Daylighting Design via Monte Carlo*. Masters thesis, Colorado State University, Fort Collins, Colorado.
17. Modest, M. F. 1983. "Daylighting Calculations for Non-Rectangular Interior Spaces with Shading Devices," *Journal of the Illuminating Engineering Society* 12:226-241 (July).
18. Moon, P., and D. E. Spencer. 1942. "Illumination from a Non-uniform Sky," *Illuminating Engineering* 37(12):707-726.
19. Müller, S., W. Kresse, N. Gatenby, and F. Schöffel. 1995. "A Radiosity Approach for the Simulation of Daylight," *Rendering Techniques* '95 (Proceedings of the Sixth Eurographics Workshop on Rendering), pp. 137-146. New York, NY: Springer-Verlag.

20. Nakamae, E. G. Jiao, K. Tadamura, and F. Kato. 1995. "A Model of Skylight and Calculation of its Illuminance," *Lecture Notes in Computer Science* Vol. 1024, pp. 304-312. Berlin, Germany: Springer-Verlag.
21. Neumann, L. 1995. "Monte Carlo Radiosity," *Computing* 55:23-42.
22. Nimeroff, J. S., E. Simoncelli, and J. Dorsey. 1994. "Efficient Re-rendering of Naturally Illuminated Environments," *Proceedings of the Fifth Eurographics Workshop on Rendering*, pp. 359-373.
23. Nishita, T., and E. Nakamae. 1986. "Continuous Tone Representation of Three-Dimensional Objects Illuminated by Sky Light," *Computer Graphics* 20(4):125-132.
24. Paule, B., J.-L. Scartezzini, and S. Citherlet. 1998. "LESO-DIAL Daylighting Design Software," *Proceedings of the International Daylighting Conference* (Daylighting '98), pp. 29-36.
25. Recker, R. J., D. W. George, and D. P. Greenberg. 1990. "Acceleration Techniques for Progressive Refinement Radiosity," *Computer Graphics* 24(2):59-66 (1990 Symposium on Interactive 3D Graphics).
26. Rushmeier, H. E., D. R. Baum, and D. E. Hall. 1990. "Accelerating the Hemi-Cube Algorithm for Calculating Radiation Form Factors," *ASME Journal of Heat Transfer* 113:1044-1047.
27. Schlick, C. 1993. "A Customizable Reflectance Model for Everyday Rendering," *Proceedings of the Fourth Eurographics Workshop on Rendering*, pp. 73-84.
28. Sillion, F. X., and C. Puech. 1994. *Radiosity & Global Illumination*. San Francisco, CA: Morgan Kaufmann.
29. Spencer, D. E., and V. J. Stakutis. 1951. "The Integral-Equation Solution of the Daylighting Problem," *Journal of the Franklin Institute* 252(9):225-237.
30. Tadamura, K., E. Nakamae, K. Kaneda, M. Baba, H. Yamashita, and T. Nishita. 1993. "Modeling of Skylight and Rendering of Outdoor Scenes," *Computer Graphics Forum* 12(3):C189-C200.
31. Tregenza, P. R. 1987. "Subdivision of the Sky Hemisphere for Luminance Measurements," *International Journal of Lighting Research & Technology* 19(1):13-14.
32. Tsangrassoulis, A., and M. Santaouris. 1997. "Daylight Modeling with Passport-Light," *Proceedings of Building Simulation '97*, Volume I, pp. 73-78.
33. Wallace, J. R., K. A. Elmquist, and E. A. Haines. 1989. "A Ray Tracing Algorithm for Progressive Radiosity," *Computer Graphics* 23(3):315-324.

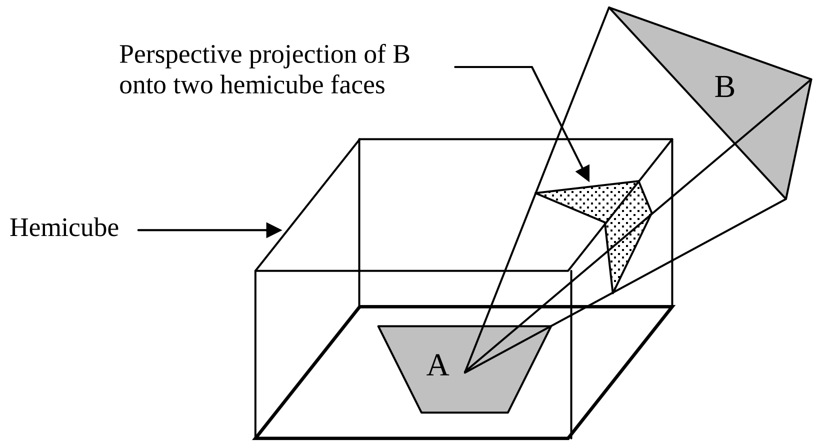

Figure 1 – Hemicube projection

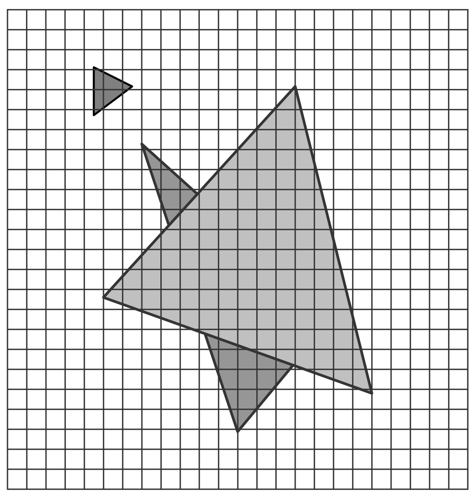

Figure 2 – Element projection onto a hemicube face

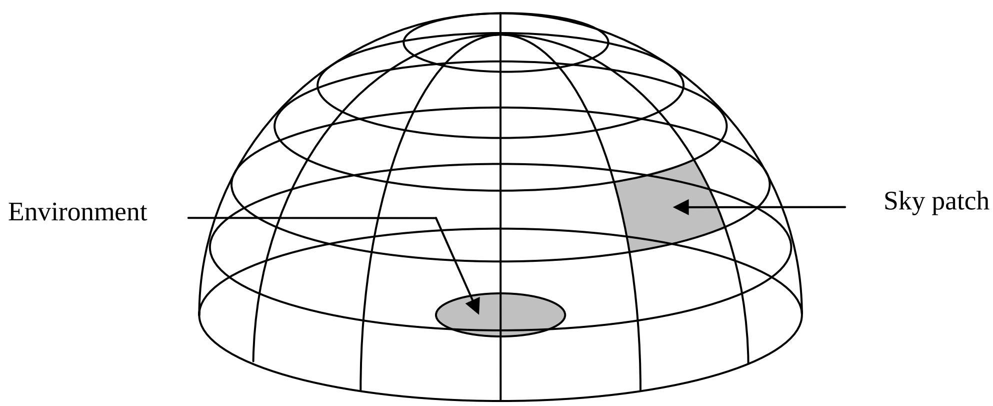

Figure 3 – Sky dome (Müller et al. 1995)

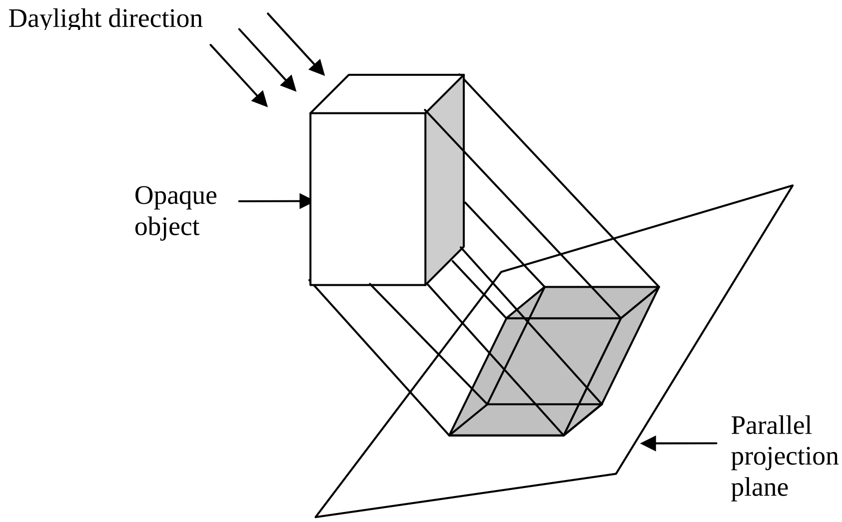

Figure 4 – Single parallel projection plane

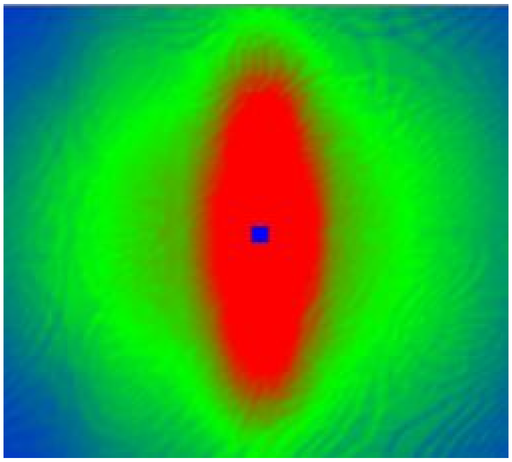

Figure 5A – Hemicube rendering

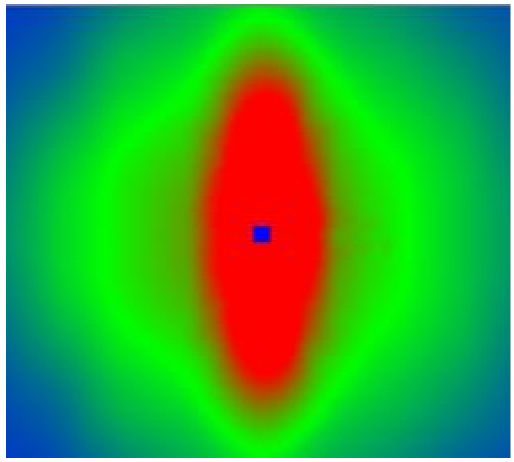

Figure 5B – Hybrid hemicube rendering

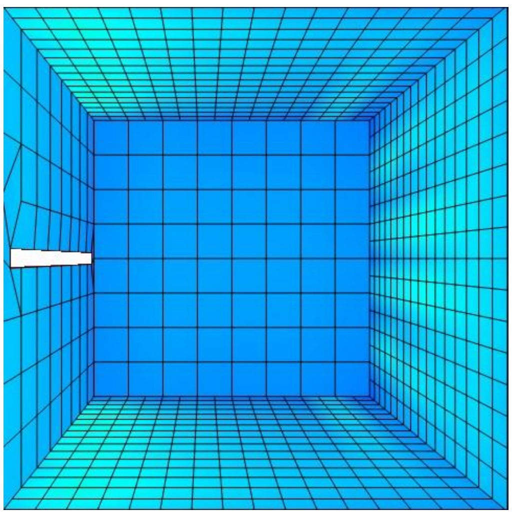

Figure 6A – Spoked patterns on walls

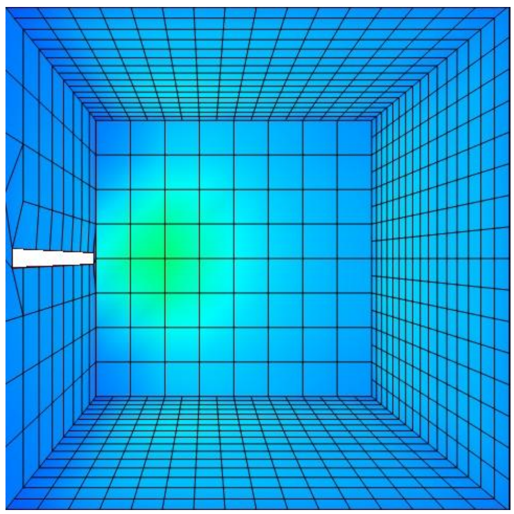

Figure 6B – Elimination of spoked patterns