



A TWO-PASS SOLUTION TO THE RENDERING EQUATION:
A SYNTHESIS OF RAY TRACING AND RADIOSITY METHODS

John R. Wallace, Michael F. Cohen,
Donald P. Greenberg
Cornell University
Ithaca, N. Y. 14853

ABSTRACT

View-independent and view-dependent image synthesis techniques, represented by radiosity and ray tracing, respectively, are discussed. View-dependent techniques are found to have advantages for calculating the specular component of illumination and view-independent techniques for the diffuse component. Based on these observations a methodology is presented for simulating global illumination within complex environments using a two-pass approach. The first pass is view-independent and is based on the hemi-cube radiosity algorithm, with extensions to include the effects of diffuse transmission, and specular to diffuse reflection and transmission. The second pass is view-dependent and is based on an alternative to distributed ray tracing in which a z-buffer algorithm is used to sample the intensities contributing to the specularly reflected or transmitted intensity.

CR Categories and Subject Descriptors: I.3.3 [Computer Graphics]: Picture/Image Generation; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism

General Terms: Algorithms

Additional Key Words and Phrases: radiosity, distributed ray tracing, z-buffer, global illumination, view-dependence, view-independence.

1.0 INTRODUCTION

The creation of realistic images for a raster display requires the calculation of the intensity of light leaving visible surfaces in the direction of the observer through each pixel of the image plane. Early methods for calculating this intensity accounted only for the direct illumination of surfaces by light sources [14][20]. The subsequent evolution of illumination models has

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

provided increasingly sophisticated methods of accounting for global illumination, which includes the effect of all objects in the environment.

Algorithms to compute global illumination may be characterized by their approach to selecting the sample points within the environment and the directions for which the final intensity is calculated. For view-dependent algorithms these sample points and directions are determined by the view position and by the discretization of the image plane. View-independent algorithms, on the other hand, calculate the illumination of all surfaces for a set of discrete environment locations and directions determined by criteria that are independent of view. A two-pass method is presented in this paper that overcomes disadvantages of previous methods by computing the global diffuse component of illumination during a view-independent preprocess and the global specular component during a view-dependent postprocess.

Recently, in separate papers, Kajiya [18] and Immel et al [17] presented essentially the same equation to completely describe the intensity of light leaving a surface in a given direction in terms of the global illumination within an environment. However, each paper described radically different approaches to the equation's solution. Kajiya extended view-dependent ray tracing [24][8] to include global diffuse illumination, which ray tracing had previously not accounted for correctly. Immel extended the view-independent approach taken by the radiosity method [13][5][19] to include directional reflection, eliminating the restriction of purely diffuse environments. Although both approaches ultimately converge to the same solution, they illustrate the complementary advantages and disadvantages of solving the problem entirely on a view-dependent or view-independent basis.

Kajiya and Immel each described a complete illumination equation based on earlier work in radiative heat-transfer [22]:

$$I_{out}(\theta_{out}) = E(\theta_{out}) + \int_{\Omega} \rho''(\theta_{out}, \theta_{in}) I_{in}(\theta_{in}) \cos(\theta) d\omega \quad (1)$$

where

I_{out} = the outgoing intensity for the surface



I_{in} = an intensity arriving at the surface from the rest of the environment
 E = outgoing intensity due to emission by the surface
 θ_{out} = the outgoing direction
 θ_{in} = the incoming direction
 Ω = the sphere of incoming directions
 θ = the angle between the incoming direction θ_{in} and the surface normal
 $d\omega$ = the differential solid angle through which the incoming intensity arrives
 ρ'' = the bidirectional reflectance/transmittance of the surface

The outgoing intensity, I_{out} , leaving a differential area of a surface in a given direction θ_{out} is the sum of the light emitted by the surface in that direction, $E(\theta_{out})$, plus any light arriving at the surface which is then reflected or transmitted in that direction. The reflected intensity depends on light arriving at the surface from all directions above the surface and the transmitted intensity depends on light arriving from all directions below the surface. Hence the integration is over the entire sphere of incoming directions. The bidirectional reflectance - transmittance $\rho''(\theta_{out}, \theta_{in})$ represents the physical reflection properties of the surface and is expressed as the ratio of the reflected or transmitted intensity in the outgoing direction θ_{out} to energy arriving from the incoming direction θ_{in} . All quantities are also dependent on wavelength; assuming no energy exchange between wavelength bands, independent equations of the same form can be written for each wavelength band.

2.0 DIFFUSE RAY TRACING AND SPECULAR RADIOSITY

Kajiya describes an algorithm that solves equation (1) using an efficient approach to stochastic ray tracing [8] in which many rays are shot per pixel with each ray generating only one path through the environment. Since light sources seen by a surface typically make a very large contribution to the reflected intensity, one ray is also sent towards a light source for every ray shot stochastically at the rest of the environment. These modifications, along with other variance reduction techniques described in the same paper, allow sufficient incoming directions to be sampled by ray tracing to evaluate the outgoing intensity for diffuse reflection, thus successfully reproducing phenomena previously only modeled by radiosity methods.

Like its ray tracing antecedents, Kajiya's method is view-dependent, being restricted to those rays that ultimately reach the eye. This is enforced by tracing rays backwards from the eye, through the pixels of the image plane and into the environment. Thus, for a given point on a surface visible to the viewer, the intensity need only be determined for one outgoing direction. The directional nature of the bidirectional reflectivity function then provides an efficient basis for selecting the important directions in which to sample the incoming intensities. In other words, the incoming intensities may be sampled at a higher frequency in the direction of the specular peak (Fig. 1a).

For ideal diffuse reflection, however, the bidirectional reflectivity is independent of the outgoing direction; thus the view direction does

not provide a basis for "importance sampling". To reduce the amount of sampling required, techniques may be used that take advantage of knowledge gained as sampling progresses. However a large number of directions may still have to be sampled, particularly when the assumption that lights are the most significant sources of illumination does not hold. Such cases may occur, for example, when a diffuse surface is entirely in shadow or when light reflected from a specular surface contributes significantly to the illumination of a diffuse surface (Fig. 1b). In addition, the pixel by pixel determination of intensity imposed by ray tracing from the eye may result in the performance of more work than necessary since the illumination of a diffuse surface as perceived by the viewer typically changes relatively slowly from one pixel to the next.

By contrast, in the radiosity approach all calculations are view-independent. The standard radiosity method treats only diffuse surfaces. The set of sample points for which the intensities are calculated depends on the discretization of the environment surfaces rather than the viewpoint and image resolution (Fig. 1d). This greatly reduces the number of sample points compared to the number of pixels in the final image. For surfaces with high intensity gradients, which occur at shadow boundaries, for example, a uniform surface discretization does not suffice. However, the distribution of these sample points for a given surface may be adaptively determined by the gradient of intensity over the surface [6].

Immel expanded the view-independent radiosity approach to include specular reflection. The relationship between a surface patch and all the other patches in the environment becomes, in Immel's approach, a relationship between a given outgoing reflection direction for a patch and all outgoing directions for all other patches. A simultaneous solution of the resulting system of equations gives an intensity in each direction for each patch.

Unfortunately, for specular surfaces the intensity as seen by the viewer typically changes very quickly from pixel to pixel. Therefore a view-independent solution may require that specular surfaces be subdivided to the point where each patch covers approximately one pixel in the final view. Although Immel's approach would eventually converge to an accurate solution as the discretization of the environment increased, the computation and space required to solve the resulting system of simultaneous equations precludes subdivision to the required level. As a result, artifacts appear when specular surfaces are rendered, since the correct intensity is known only at relatively widely spaced sample points (Fig. 1c).

2.1 A TWO COMPONENT MODEL OF THE TRANSFER OF LIGHT ENERGY

A reasonable approach to solving the global illumination problem is to divide it in a manner that allows the solution to take advantage of the strengths of both view-dependent and view-independent sampling.




Figure 1a. View dependent calculation of specular component: The intensity calculation is limited to the outgoing directions that reach the eye. Incoming intensities may then be sampled at a higher frequency in the direction of the specular peak. The outgoing intensity on the floor will be calculated at points corresponding to each pixel, thus capturing the reflection of the table's edge to the accuracy required by the view.




Figure 1c. View independent calculation of specular component: All outgoing directions must be accounted for, hence there is no preferred incoming direction. The outgoing intensity used when rendering from a specific view is the result of weighting the equally spaced incoming samples. The intensity at point A will be determined by interpolating from the widely spaced sample points. Thus the reflection of the edge of the table will not be accurately captured.




Figure 1b. View dependent calculation of diffuse component: For diffuse reflection the contribution of incoming intensities to the outgoing intensity is independent of the outgoing direction. If the light source is not visible to the sample point, many incoming directions may have to be sampled, since the significant sources of illumination may be difficult to find. This will be repeated for each pixel in which the floor is visible.




Figure 1d. View independent calculation of diffuse component: Many incoming directions are sampled for each sample point with their contributions weighted according to Lambert's cosine law. The sample points are spaced according to the gradient of illumination. The intensity at point A will be determined by interpolating from the surrounding sample points.

Local light reflection models for computer graphics traditionally separate the scattering of light from a surface into two components, a non-directional or diffuse term and a directional or specular term [20][4][7]. This corresponds to approximating the bidirectional reflectance function in equation (1) as the sum of a diffuse portion, ρ_d , which is independent of viewpoint, and a specular portion, ρ_s , which depends on the view direction. Thus,

$$\rho''(\theta_{out}, \theta_{in}) = k_s \rho_s(\theta_{out}, \theta_{in}) + k_d \rho_d$$

where

- k_s = fraction of reflectance that is specular
- k_d = fraction of reflectance that is diffuse
- $k_s + k_d = 1$

Equation (1) can then be rewritten as

$$I_{out}(\theta_{out}) = E(\theta_{out}) + I_{d,out} + I_{s,out}(\theta_{out}) \quad (2a)$$

where

$$I_{d,out} = k_d \rho_d \int I_{in}(\theta_{in}) \cos(\theta) d\omega \quad (2b)$$

$$I_{s,out}(\theta_{out}) = k_s \int \rho_s(\theta_{out}, \theta_{in}) I_{in}(\theta_{in}) \cos(\theta) d\omega \quad (2c)$$

The outgoing diffuse and specular terms each depend on all incoming intensities I_{in} . Since these incoming intensities are, in fact, just outgoing intensities from other surfaces, they in turn contain both diffuse and specular components. Thus, it is not theoretically correct to solve for the diffuse and specular components independently, precluding a simple superposition of ray tracing and standard radiosity solutions, (although such an approach may be acceptable in many cases). Since the standard radiosity solution considers only diffuse inter-reflection, the illumination of a diffuse surface by light reflected specularly from another surface will not be accounted for and the subsequent effect of this light on the global illumination will be lost. The superposition of ray tracing will not recover it, since ray tracing




Figure 2. The four "mechanisms" of light transport: (a) diffuse to diffuse, (b) specular to diffuse, (c) diffuse to specular and (d) specular to specular.

does not consider the illumination of diffuse surfaces by other surfaces at all.

Given the separation of ρ into ρ_d and ρ_s , the transfer of light from one surface to another can be thought of as occurring by way of four "mechanisms". These are illustrated in Figure 2. Light reflected diffusely by one surface may have arrived at that surface via diffuse reflection (Fig. 2a) or specular reflection (Fig. 2b) from another surface. Similarly, light reflected specularly may also have arrived via diffuse reflection (Fig. 2c) or specular reflection (Fig. 2d) from the other surface. The first two mechanisms are expressed by equation (2b) and the latter two by equation (2c). The diffuse to diffuse transfer is handled by the standard radiosity algorithms. Specular to specular and diffuse to specular transfer are handled by standard ray tracing algorithms, although the diffuse component is not correctly modeled. Neither approach in its conventional form has handled the specular to diffuse transfer, although other approaches, such as backwards ray tracing [2], beam tracing [16] and the methods of Kajiyama and Immel described above, have been presented to account for this transfer mechanism. The backwards ray tracing method described by Arvo, in particular, anticipates the two-pass nature of the approach described below.

3.0 TWO-PASS IMAGE SYNTHESIS

The two-pass approach for computing global illumination described in the following sections accounts for all four mechanisms of light transfer. In a view-independent stage of the algorithm, called the preprocess, the complete global propagation of light is approximated in order to determine the diffuse component of intensity for all surfaces (Eq. 2b). All four transfer mechanisms are included in the preprocess. Thus specular reflection must be accounted for, but only to the extent necessary to accurately calculate the diffuse component. A view-dependent stage of the algorithm, called the postprocess, then uses the results of the preprocess as the basis for calculating the specular component (Eq. 2c) to the accuracy required by the view. The postprocess accounts for the specular to specular and diffuse to specular transfer mechanisms, in the latter case using the diffuse component calculated during the preprocess as the source. For each pixel in the final image, the resulting specular component is

added to the diffuse component, interpolated from the preprocess sample points, providing a complete solution to equation (2a).

3.1 THE PREPROCESS

The standard radiosity solution using the hemi-cube algorithm provides the basis for the preprocess [5]. (Details of the radiosity method are not included here. The hemi-cube algorithm is a geometrically based numerical integration technique for calculating form-factors, which are purely geometric terms describing the transfer of energy from one surface to another.) Since the standard radiosity method accounts only for diffuse reflection, extensions to the hemi-cube algorithm must be made to add diffuse transmission and to approximate specular reflection and specular transmission in so far as it effects the diffuse component for all surfaces.

3.1.1 Translucency

Ideal diffuse transmission (translucency) is readily included by placing a hemi-cube on the back, as well as the front, of the transmitting surface and calculating backward in addition to the usual forward form-factors [11][21]. The backward form-factors represent the effect that light from surfaces seen by the back of the translucent surface has on the intensity of the front of the translucent surface.

3.1.2 Specular to Diffuse Transport

Specular reflection or transmission may also be accounted for in the standard radiosity solution by performing additional work during the calculation of form-factors [12]. The form-factor, it will be recalled, represents the fraction of the total energy leaving a given diffuse surface patch that arrives at a second diffuse surface patch. Specular surfaces are treated as additional routes by which light energy may reach one diffuse surface from others (Fig. 3). Thus, two diffuse surfaces that "see" each other via specular reflection or transmission by intermediate surfaces have an extra form-factor representing the additional fraction of energy that will be transferred over this route. The additional form-factors may be calculated by any method that can determine the path of specular reflection or transmission. The radiosity solution proceeds as before by solving the system of equations describing the interrelationships between




Figure 3. Calculation of extra form-factors to account for mirror reflection. Patch A receives light directly from patch B and indirectly through reflection by the mirror. The mirror is treated as a window into a virtual "mirror world." Projecting patch B' onto the hemicube is then equivalent to following the actual path of reflection back to patch B.

diffuse surfaces, with specular reflection included implicitly to the extent that it affects these interrelationships. This accounts for the intensity of the diffuse component due to the diffuse to specular, specular to specular, and specular to diffuse transfer mechanisms without explicitly calculating the specular component for any surface.

In the current implementation, this technique is restricted to perfect, planar mirrors or filters. This allows the additional form-factors to be calculated with the current z-buffer hemi-cube algorithm by creating a "virtual world" which is seen through the mirror or filter [21] (Fig. 3). The approach may be generalized to handle curved surfaces and refraction by using ray tracing to determine the form-factors.

3.13 Result of the Preprocess

The final result of the preprocess is an accurate determination of the diffuse intensity at selected sample points (the element vertices) within the environment. Although some of these sample points may end up being hidden from a particular viewpoint, the number of sample points required is typically much less than the number that would have been imposed by a pixel to pixel calculation (although this would not necessarily be the case for environments with a very large number of surfaces). The hemi-cube algorithm, in effect, determines the incoming intensities at each sample point for a number of directions determined by the hemi-cube resolution. The effective number of directions is in the tens of thousands for typical hemi-cube resolutions. In a sense, the effort that would have been required to determine the diffuse component at every pixel in a view-dependent approach is applied instead to more accurately determining the illumination at a smaller number of points.

3.2 THE POSTPROCESS

The postprocess takes the results of the preprocess and, for a given view, completes the solution of the illumination equation for the surface visible at each pixel. The view-dependence of the postprocess permits the efficient calculation of the specular component $I_{s,out}(\theta_{out})$ (Eq. 2b) by limiting the direction for which it must be calculated to the view direction. As in previous radiosity methods, the diffuse component $I_{d,out}$ is calculated by interpolation from the element vertex intensities determined during the preprocess. The sum $I_{d,out} + I_{s,out}(\theta_{out})$ completes the solution to the integral in equation (1) and, with the addition of intensity due to emission, provides the final intensity at each pixel.

The specular component $I_{s,out}(\theta_{out})$ depends on the intensities arriving at the surface from the entire hemisphere of directions, weighted by the specular bidirectional reflectance $\rho_s(\theta_{in}, \theta_{out})$. Distributed ray tracing [8] and "cone tracing" [1] were the first algorithms for sampling these incoming intensities. The method presented in this section is based on the observation that in most specular reflection, the bidirectional reflectance is such that only a fraction of the incoming intensities over the hemisphere make a significant contribution to the outgoing intensity in a particular direction. Thus, with little loss of accuracy, the limits of integration may be reduced from the entire hemisphere to the smaller solid angle over which the weighted intensity is significant. The determination of the incoming intensities requires that the visible surfaces and intensities be found over the solid angle of interest, a problem that is precisely equivalent to that of computing an image from the surface intersection point with a view in the mirror direction.




Figure 4. (a) A specular bidirectional reflectivity function for a particular outgoing direction showing the reflection frustum and the solid angle through which incoming intensities will be sampled.

(b) A two-dimensional slice through the reflectivity function of figure 4a plotted in rectangular coordinates with incoming angle on the x axis and reflectivity on the y axis. The weight that will be given to incoming intensities is proportional to the value of this function at the incoming angle.




Figure 5. Sampling of the intensities arriving through the reflection frustum using the z-buffer algorithm. The incoming diffuse component due to element B at each "pixel" is obtained by interpolating from the diffuse intensities at the element vertices. The incoming specular component is obtained recursively.

3.21 The Reflection Frustum

A view frustum encompassing this solid angle is constructed (Fig. 4). The visible surfaces for this reflection frustum are determined using a standard low resolution z-buffer algorithm (typically on the order of 10 by 10 "pixels"). The incoming intensity "seen" through each "pixel" of this frustum is simply the intensity of the surface visible at that pixel (Fig. 5). This intensity may contain both diffuse and specular components. The incoming diffuse component at each reflection frustum pixel is determined during the z-buffer operation by Gouraud shading from the element vertex intensities calculated during the preprocess. Where a surface visible in a reflection frustum pixel has a specular component, its intensity is calculated by applying the entire postprocess algorithm recursively, analogously to the similar case in traditional ray-tracing. The work required for this recursive step may be reduced by adaptively limiting the depth of recursion [15] and by reducing the resolution of the frustum for successive bounces.

Once the incoming intensities have been determined, the integral in equation (2c) can be numerically evaluated. The subdivision of the reflection frustum into pixels acts as a discretization of the domain of integration. The integral thus becomes the summation:

$$I_{s,out} = k_s \sum_{i=0}^n \sum_{j=0}^n W_{i,j} I_{in}(\theta_{i,j}) \cos(\theta_{i,j}) \Delta\omega_{i,j}$$

where the bidirectional reflectance function is represented by $W_{i,j}$, an array of weights, and n is the resolution of the reflection frustum. Each weight in the array corresponds to the incoming direction $\theta_{i,j}$, determined by the "pixel" (i,j) of

the reflection frustum and is proportional to the bidirectional reflectance for the view direction and the incoming direction determined by that "pixel".

The array of weights can be pre-computed as a look-up table for a simple Phong-like reflectance function at a given reflection frustum resolution, making the approximation that the variation of $\cos(\theta)\Delta\omega$ over the frustum is small and can be ignored. Various surface finishes can be simulated by simply varying the size of the frustum; the smaller the frustum, the more mirror-like the reflection. More complex physically based reflection models may require more elaborate look-up tables or require the weights to be computed on the fly based on reflection angle, frustum size and resolution.

3.22 Transparency

Specular transmission, including refraction, is accounted for analogously to reflection by defining a bidirectional transmittance and constructing a transmission frustum in the transmitted direction.

3.23 Antialiasing

Inevitably, a uniform point sampling method will produce aliasing in one form or another. Figure 6a provides an example of aliasing produced by the regular point sampling of the incoming intensities through the reflection frustum. The uniform sampling produces a striped or plaid effect due to an alignment of the sampling pattern with the projection of object edges in the environment. A small shift in the reflection frustum may result in a sudden jump in the number of reflection frustum "pixels" covered by the object. Selecting sample points stochastically rather than uniformly, as imposed by the regular grid of "pixels", may eliminate aliasing by converting it to noise [9][10]. However, this also eliminates the speed




Figure 6. (a) Aliasing due to uniform point sampling of the incoming intensities.

(b) Result of rotating the reflection frustum around its central axis by a random amount for each image pixel (enlargement of lower left corner of image in Figure 11b).




Figure 7. Stainless steel spheres with finishes ranging from completely diffuse to highly specular.

hemi-cube resolution = 50 by 50
 reflection frustum resolution:
 first bounce = 15 by 15
 second bounce = 10 by 10
 reflection frustum size:
 middle sphere = 0.79 radians
 right-hand sphere = 0.16 radians

advantages of an incremental algorithm like the z-buffer. To preserve this advantage, instead of randomly distributing the samples themselves, the reflection frustum used to calculate the specular component at each pixel in the final image is rotated around its central axis by a random amount before the z-buffering is performed. This rotation alters the alignment between the reflection frustum sampling pattern and the environment from pixel to pixel in the final image, thus reducing the aliasing artifacts considerably (Fig. 6b).

4.0 RESULTS

Figure 7 shows a series of spheres that range from completely diffuse to highly specular. The spheres were subdivided into 96 patches for the preprocess. Adaptive subdivision [6] was found to be

particularly important for curved surfaces, because of the relatively rapid change in diffuse intensity produced as the surface normal turns away from a light source. The depth of recursion during the postprocess was limited to two by treating specular surfaces as entirely diffuse on the third bounce.

Figure 8 shows a series of thin glass sheets demonstrating a range from ideal diffuse to ideal specular transmission. The left hand glass panel acts as a purely diffuse transmitter (like opal glass), the center panel is more specular (like frosted glass), while the panel on the right is an almost-ideal specular transmitter. The etched glass in Figure 9 was produced using a procedurally defined "translucency map" to determine the size of the transmission frustum at each point on the glass.




Figure 8. Three panes of translucent glass. The left-hand panel demonstrates ideal diffuse transmission.

hemi-cube resolution = 200 by 200
 transmission frustum resolution = 15 by 15
 transmission frustum size:
 middle panel = 0.79 radians
 right-hand panel = 0.16 radians




Figure 9. Etched glass.

Figures 10a, b and c show an environment for which the illumination of diffuse surfaces by specular reflection (via the specular to diffuse transfer mechanism) is significant. Several versions of the image have been calculated to emphasize the effects of various levels of sophistication provided by the preprocess. Overall illumination of the room is provided by a dim area source on the ceiling. The reflection in the mirror in each case was calculated during the postprocess. In Figure 10a only the direct illumination of surfaces by the light sources is accounted for, hence surfaces facing away from lights or in shadow are black. In Figure 10b the full radiosity solution has been applied but the effect of the specular reflection from the mirror on the illumination of diffuse surfaces has been ignored. In Figure 10c this specular to diffuse transfer has been accounted for using the mirror form-factor technique. The illumination of the top of the vanity along with the objects on it results primarily from light emitted by the lamp on the vanity and subsequently reflected by the mirror. The environment is subdivided into 4224 elements, providing only 5379 sample points for which the diffuse component was calculated. This is a significantly smaller number than the number of pixels at which an entirely view-dependent calculation would have calculated the diffuse component. The hemi-cube resolution during the preprocess was 150 by 150. The specular component of the mirror was calculated with a frustum of 0.032 radians.

Figure 11 shows an environment inspired by a well-known painting, "Lady and Gentleman at the Virginals", by the 17th century Dutch painter Jan Vermeer. Vermeer is particularly known for his use of light to define space, based, at least in part, on a sensitivity to the effects of what we would now call "global illumination", effects such as penumbra and "color bleeding", for example. For the image in Figure 11a the floor was treated as an ideal diffuse reflector. As far as is known, the floor of his studio never attained the high sheen demonstrated in Figure 11b. The hemi-cube resolution for the preprocess was 50 by 50 for both images. The resolution of the reflection frustum was 10 by 10 for the floor in Figure 11b, with a reflection frustum size of 0.079 radians.

All images were calculated on DEC VAX 11/750, 11/780, 8300 or 8700 computers and displayed using 1280 by 1024 24-bit frame buffers from Raster Technologies or Hewlett-Packard.

5.0 CONCLUSION AND FUTURE DIRECTIONS

A hybrid, two-pass methodology for simulating global illumination within general complex environments has been presented. The approach takes advantage of the complementary strengths of view-dependent and view-independent methods. The view independent preprocess, based on the radiosity method, provides an efficient computation of the diffuse component. The view-dependent postprocess, based on ray tracing, efficiently calculates the

specular component. By taking advantage of the best features of each a complete simulation is achieved.


The standard radiosity algorithms are also extended to handle curved surfaces and diffuse transmission and to account for the contribution of specular reflection to the illumination of diffuse surfaces. An alternative to distributed ray tracing is described in which the specular component is determined by a regular sampling of intensities contributing to the reflected intensity using a z-buffer algorithm.

Both the preprocess and the postprocess require the performance of a very large number of independent visible surface calculations. These are necessary during the preprocess to calculate the form-factors and during the postprocess to determine incoming intensities contributing to the specular component. In fact, the total computation time in both processes is overwhelmingly dominated by these calculations. The choice of the standard z-buffer algorithm was motivated in part by recent advances in the implementation of such algorithms in special-purpose hardware [23]. The use of such hardware to perform the z-buffer portions of the hemi-cube algorithm and the distributed ray tracing algorithm described above promises dramatic increases in the speed of both the preprocess and the postprocess. In addition, with increased processing speed, aliasing problems may be reduced by increasing the resolution of the reflection frusta.

Future implementations should increase the generality of the preprocess by including ray tracing to calculate reflected and refracted form-factors where necessary. Future implementations should also provide for cases where the gradient of the diffuse intensity is extremely high, such as along sharp shadow boundaries. Just as with the specular component, the attempt to solve for these cases in an entirely view independent manner is inefficient. Regions of high diffuse gradient may instead be merely identified in the preprocess and solved for in the postprocess on a pixel by pixel basis.

The two-pass approach presented here is more general than the particular implementation described. Many alternative pre and post processes may be imagined. Conventional distributed ray tracing, for example, could be substituted for the z-buffer algorithm in the postprocess. Path tracing, as Kajiya named his approach, might use the approximate description of the propagation of light gained during the preprocess as a basis for importance sampling.

The two-pass approach also provides a framework for the progressive refinement of images, in the spirit of the approach described by Bergman, et al [3]. Using currently available hardware the pre-computed diffuse component provided by the preprocess may be rendered at interactive rates, enabling "walk-throughs" of static environments. When the viewer lingers at a certain view, refinement of the shading of visible surfaces by the postprocess may commence.



(a)



(c)



(b)

Figure 10. (a) Direct illumination by light sources only.

(b) Diffuse to diffuse transfer included. Specular to diffuse ignored.

(c) Full solution.



Figure 11a. A 17th-century Dutch interior.



Figure 11b. A 17th-century Dutch interior after polishing the floor.

ACKNOWLEDGMENTS

The z-buffer distributed ray tracing algorithm is based on ideas developed by Michael Cohen and Lisa Maynes. Particular thanks go to Holly Rushmeier for many valuable discussions, to Stewart Feldman and Wayne Lytle for the diagrams, to Emil Ghinger for photography, and to the reviewers for their helpful comments. This research was conducted under grant No. DCR-8203979 from the National Science Foundation and was supported by generous equipment grants from the Digital Equipment Corporation and Hewlett-Packard.

REFERENCES

1. Amanatides, John, "Ray Tracing with Cones," Proceedings of SIGGRAPH'84, In Computer Graphics, Vol. 18, No. 3, July 1984, pp. 129-136.
2. Arvo, James, "Backward Ray Tracing," Developments in Ray Tracing, SIGGRAPH Course Notes, Vol. 12, 1986.
3. Bergmann, Larry, Henry Fuchs, Eric Grant, Susan Spach, "Image Rendering by Adaptive Refinement," Proceedings of SIGGRAPH'86, In Computer Graphics, Vol. 20, No. 4, Aug. 1986, pp. 29-38.
4. Blinn, James F., "Models of Light Reflection for Computer Synthesized pictures," Proceedings of SIGGRAPH'77, In Computer Graphics, Vol. 11, No. 2, 1977, pp. 192-198.
5. Cohen, Michael F. and Donald P. Greenberg, "A Radiosity Solution for Complex Environments," Proceedings of SIGGRAPH'85, In Computer Graphics, Vol. 19, No. 3, July 1985, pp. 31-40.
6. Cohen, Michael F., Donald P. Greenberg, David S. Immel, Philip J. Brock, "An Efficient Radiosity Approach for Realistic Image Synthesis," IEEE Computer Graphics and Applications, Vol. 6, No. 2, March 1986, pp. 26-35.
7. Cook, Robert L., and Kenneth E. Torrance, "A Reflection Model for Computer Graphics," ACM Transactions on Graphics, Vol. 1, No. 1, January 1982, pp. 7-24.
8. Cook, Robert L., Thomas Porter and Loren Carpenter, "Distributed Ray Tracing," Proceedings of SIGGRAPH'84, In Computer Graphics, Vol. 18, No. 3, July 1984, pp. 137-145.
9. Cook, Robert L., "Stochastic Sampling in Computer Graphics," ACM Transactions on Graphics, Vol. 5, No. 1, January 1986, pp. 51-72.
10. Dippe, Mark A. Z., Erling Henry Wold, "Antialiasing Through Stochastic Sampling," Proceedings of SIGGRAPH'85, In Computer Graphics, Vol. 19, No. 3, July 1985, pp. 69-78.
11. Dunkle, R. V., "Radiant interchange in an enclosure with specular surfaces and enclosures with window or diathermanous walls," in Heat Transfer, Thermodynamics and Education, edited by H. A. Johnson, Boelter Anniversary Volume, New York: McGraw-Hill, 1964.
12. Eckert, E. R. G., and Sparrow, E. M., "Radiative Heat Exchange Between Surfaces with Specular Reflection," International Journal of Heat and Mass Transfer, Vol. 3, pp. 42-54, 1961.
13. Goral, Cindy M., Kenneth E. Torrance, Donald P. Greenburg, Bennet Battaile, "Modeling the Interaction of Light Between Diffuse Surfaces," Proceedings of SIGGRAPH'84, In Computer Graphics, Vol. 18, No. 3, July 1984, pp. 213-222.
14. Gouraud, H., "Continuous Shading of Curved Surfaces," IEEE Transactions on Computers, Vol. 20, No. 6, June 1971, pp. 623-628.
15. Hall, Roy A. and Donald P. Greenberg, "A Testbed for Realistic Image Synthesis," IEEE Computer Graphics and Applications, Vol. 3, No. 10, Nov. 1983, pp. 10-20.
16. Heckbert, Paul S. and Pat Hanrahan, "Beam Tracing Polygonal Objects," Proceedings of SIGGRAPH'84, In Computer Graphics, Vol. 18, No. 3, July 1984, pp. 119-128.
17. Immel, David S., Michael F. Cohen, Donald P. Greenberg, "A Radiosity Method for Non-Diffuse Environments," Proceedings of SIGGRAPH'86, In Computer Graphics, Vol. 20, No. 4, Aug. 1986, pp. 133-142.
18. Kajiya, James T., "The Rendering Equation," Proceedings of SIGGRAPH'86, In Computer Graphics, Vol. 20, No. 4, Aug. 1986, pp. 143-150.
19. Nishita, Tomoyuki and Eihachiro Nakamae, "Continuous Tone Representation of Three-Dimensional Objects Taking Account of Shadows and Interreflection," Proceedings of SIGGRAPH'85, In Computer Graphics, Vol. 19, No. 3, July 1985, pp. 22-30.
20. Phong, Bui Tuong, "Illumination for Computer Generated Pictures," Communications of the ACM, Vol. 18, No. 6, June 1975, pp. 311-317.
21. Rushmeier, Holly E., "Extending the Radiosity Method to Transmitting and Specularly Reflecting Surfaces," Master's thesis, Cornell Univ., Ithaca, 1986.
22. Siegel, Robert and John R. Howell, Thermal Radiation Heat Transfer, Hemisphere Publishing Corp., Washington DC., 1981.
23. Swanson, Roger W. and Larry J. Thayer, "A Fast Shaded-Polygon Renderer," Proceedings of SIGGRAPH'86, In Computer Graphics, Vol. 20, No. 4, Aug. 1986, pp. 95-102.
24. Whitted, Turner, "An Improved Illumination Model for Shaded Display," Communications of the ACM, Vol. 23, No. 6, June 1980, pp. 343-349.