A RADIOSITY METHOD FOR NON-DIFFUSE ENVIRONMENTS

David S. Immel, Michael F. Cohen,
Donald P. Greenberg
Cornell University
Ithaca, N.Y. 14853

ABSTRACT

A general radiosity method accounting for all interreflections of light between diffuse and non-diffuse surfaces in complex environments is introduced. As contrasted with previous radiosity methods, surfaces are no longer required to be perfectly diffuse reflectors and emitters. A complete, viewer independent description of the light leaving each surface in each direction is computed, allowing dynamic sequences of images to be rendered with little additional computation per image. Phenomena such as "reflection tracking", reflections following a moving observer across a specular surface are produced. Secondary light sources, such as the light from a spotlight reflecting off a mirror onto a wall are also accounted for.

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

General Terms: Algorithms

Additional Keywords and Phrases: Radiosity, intensity, bi-directional reflectance, hidden-surface, depth buffer, non-diffuse reflection.

INTRODUCTION

The production of realistic images requires the ability to simulate the propagation of light in an environment. This requires two steps: modeling the interaction of light with an individual surface and combining the effects of emission, transmission and reflection between all the surfaces in the environment. In the past, computer generated pictures were obtained by simulating how light is reflected from an individual surface, ignoring the interreflections from surface to surface. A number of reflection models were developed, ranging from a simple Lambertian diffuse model in which all light is reflected with equal intensity in all directions, to more complex models including non-uniform specular distributions based on geometric and electromagnetic properties of the surface. For any given reflection model, images were created by determining the visible surface at each screen pixel location, and then computing the intensity of light leaving that surface in the direction of the eye. This intensity was typically found from the geometric relationships between the visible surface, the light sources, and the eye.

Ray-tracing was the first method introduced which attempted to include the effects of shadowing and reflections from neighboring surfaces [1]. In ray-tracing, an additional intensity contribution from secondary reflected and transmitted rays is added from the mirror and refracted directions. The fact that a limited number of rays are traced from the eye into the environment precludes the ability to completely account for the global illumination arising from complex interreflections within the environment.

Although ray-tracing techniques have been able to produce strikingly realistic images, only radiosity methods can account for the complex diffuse interreflections within an environment. The radiosity method, borrowed from thermal engineering techniques [2] [3], was introduced in the context of computer graphics by Goral in 1984 [4] and was extended for use in complex environments by Cohen [5] [6] and Nishita [7]. This method is based on the fundamental Law of Conservation of Energy within closed systems. It entails performing a simultaneous global solution for the intensity of light leaving each surface by constructing and solving a set of linear equations describing the transfer of diffuse light energy between all surfaces. Calculations are performed independently for any number of discrete wavelength bands. The advantages of this method are the increased accuracy in simulating the global illumination effects and the fact that all of the intensity calculations are independent of viewing parameters, thus allowing efficient rendering of dynamic sequences.

An underlying assumption in all previous radiosity methods has been that all surfaces exhibit perfect Lambertian diffuse reflections. Environments could not contain surfaces with specular properties. This paper demonstrates a method which removes this restriction by generalizing the radiosity formulation to include

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.

© 1986 ACM 0-89791-196-2/86/008/0133 $00.75

directional reflection properties, while still maintaining the advantages of accurate global illumination and view independence.

To correctly determine the illumination of an environment, one must simulate the spatial and spectral intensity distribution of light leaving each differential surface area (or point) within the environment. There are two requirements to predict this intensity distribution. First, the spatial and spectral distribution of the light arriving at the point must be specified. This "incident" light is a function of the point's geometric relationship to all other surfaces in the environment and the light leaving those surfaces in the direction of the point. Second, a reflection model must be specified which describes how light of any spectral content, which arrives from any arbitrary direction, is scattered from the surface. This may include transmission as well as reflection and is usually referred to as the bidirectional reflectance of a surface. Since the light arriving at any surface is a function of the light leaving all other surfaces, it is apparent that a simultaneous solution for the global illumination must be performed.

The solution would contain all of the information necessary to render an image of any part of the environment from any viewpoint. As with standard diffuse radiosity solutions, the complex interreflections within an environment are correctly modeled, yielding "color bleeding" effects, which occur when one surface acts to illuminate another surface, as well as soft shadows and penumbrae, which are associated with area light sources. The new procedure for non-diffuse environments inherently contains additional directional phenomena such as specular reflections, which follow the viewer as the eye position changes. This "reflection tracking" can be seen by moving your head from side to side while looking at a shiny surface. Secondary light sources, such as the light from a spot light bouncing off a mirror onto the floor, are also accounted for. This source of secondary illumination has been ignored by all previous image synthesis methods.

THEORETICAL CONCEPTS

The intensity of a point in a certain direction, or "directional intensity", is the sum of the total reflected intensity from the point in that direction and, if the surface is an emitter, the light intensity emitted from the point in that direction:

$$I_t(\Psi_o) = I_e(\Psi_o) + I_r(\Psi_o) \qquad (1)$$

Figure 1 shows the hemisphere of directions surrounding the normal to a point, or differential area, dA_1, through which and from which it receives and reflects light. The intensity emitted directly from dA_1 in outgoing direction o is simply:

$$I_e(\Psi_{o,1}) = \varepsilon_1(\Psi_{o,1}) \qquad (2)$$

The subscript 1 has been added to indicate that the emittance and reflected intensity terms emanate from surface 1.

Table 1
Definition of Terms

Ψ represents a vector direction to or from a surface.

$\Psi_{i,1}$ = incoming direction to surface 1
$\Psi_{i,2}$ = incoming direction to surface 2
$\Psi_{o,1}$ = incoming direction from surface 1
$\Psi_{o,2}$ = incoming direction from surface 2
θ_1 = angle away from the normal to surface 1
ε_1 = directional emission from surface 1
ρ''_1 = bidirectional refletance of surface 1
$d\omega_1$ = differential solid angle around surface 1
$E_{i,1}$ = energy incident to surface 1
$I_{r,1}$ = intensity reflected from surface 1
$I_{e,1}$ = intensity emitted from surface 1
$I_{t,1}$ = total directional intensity leaving surface 1
$I_{t,2}$ = total directional intensity leaving surface 2
$\Omega_{1,2}$ = total solid angle encompassing surface 2 when viewed from surface 1

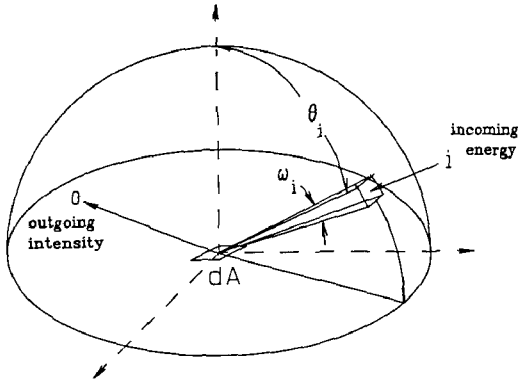

Hemisphere of Directions to a Differential Area

FIGURE 1

The spectral and spatial distribution of the reflected light is a function of the material properties of the surface from which it reflects. The light intensity reflected from a point can be found given a full description of the light arriving at that point and the surface's bidirectional reflectance. The bidirectional reflectance of a surface defines the relationship between the light energy arriving at some point and the intensity of light leaving that point:

$$\rho'' = \frac{I_r}{E_i} \qquad (3)$$

The total incoming energy is a function of the incoming intensity I_i, the solid angle $d\omega_i$ through which it arrives, and the projected differential area. The incoming energy per unit area from direction Ψ_i is thus:

$$E_i(\Psi_i) = I_i(\Psi_i)\cos\theta_i d\omega_i \qquad (4)$$

The reflected intensity in direction Ψ_o from energy incident from direction Ψ_i is therefore:

$$I_r(\Psi_o, \Psi_i) = \rho''(\Psi_o, \Psi_i) I_i \cos\theta_i d\omega_i \qquad (5)$$

Consider an environment consisting of only two surfaces, A_1 and A_2 (Figure 2). Intensity radiates from differential area dA_1 in direction o due to both emission and reflection. The reflected intensity in the outgoing direction results from energy arriving at dA_1 from A_2. Let $I_{i,1}$ be the incident intensity through solid angle $d\omega_1$. By integrating Equation 5 over the solid angle encompassing area A_2, the total intensity reflected in direction o due to the light energy arriving from A_2 is:

$$I_r(\Psi_{o,1}) = \int_{\Omega_{1,2}} \rho_1''(\Psi_{o,1}, \Psi_{i,1}) I_{i,1}(\Psi_{i,1}) cos\theta_{i,1} d\omega_{i,1} \quad (6)$$

In Equation 6, the incoming intensities, $I_{i,1}$, result from and are equal to the outgoing intensities, $I_{t,2}$, from A_2:

$$I_{i,1}(\Psi_{i,1}) = I_{t,2}(\Psi_{o,2}) \quad (7)$$

Substituting the outgoing intensities from A_2 for the incoming intensities to dA_1, Equation 6 becomes:

$$I_r(\Psi_{o,1}) = \int_{\Omega_{1,2}} \rho_1''(\Psi_{o,1}, \Psi_{i,1}) I_{t,2}(\Psi_{o,2}) cos\theta_{i,1} d\omega_{i,1} \quad (8)$$

Therefore, the directional intensity from surface dA_1 in direction o is the sum of the emitted intensity (Equation 2) and the reflected intensity (Equation 8):

$$I_{t,1}(\Psi_{o,1}) = \varepsilon_i(\Psi_{o,1}) + \int_{\Omega_{1,2}} \rho_1''(\Psi_{o,1}, \Psi_{i,1}) I_{i,1}(\Psi_{i,1}) cos\theta_{i,1} d\omega_{i,1} \quad (9)$$

Similar equations define intensities leaving dA_1 for all other directions within the hemisphere in the simple two-surface environment considered above. Each outgoing directional intensity from dA_1 is a function of the directional intensities leaving area A_2. In an equivalent manner, the directional intensities leaving each differential area of A_2 are a function of the directional intensities leaving area A_1. Thus, all of the outgoing directional intensities from both surfaces must be solved simultaneously.

In addition to the directional variation, intensities may vary continuously according to spatial position across a surface, or according to wavelength across the visible spectrum, since the spectral distributions and bidirectional reflectances are wavelength dependent. Thus, there are theoretically an infinite number of unknown intensity values to be simultaneously solved.

The problem is made tractable by discretizing the surfaces, the visible spectrum, and the hemisphere of incoming and outgoing directions. As with previous radiosity approaches, surfaces are divided into small discrete areas, or "patches", with the assumption that for each direction, the intensities arriving and leaving the patch are constant over the whole patch. The visible spectrum is likewise divided into a small number of wavelength bands. Each wavelength band is assumed to be independent from each other, allowing a separate solution for the light interation for each band.

The hemisphere above the patch is also discretized into a finite number of directions, or solid angles. Thus, for a particular patch, the

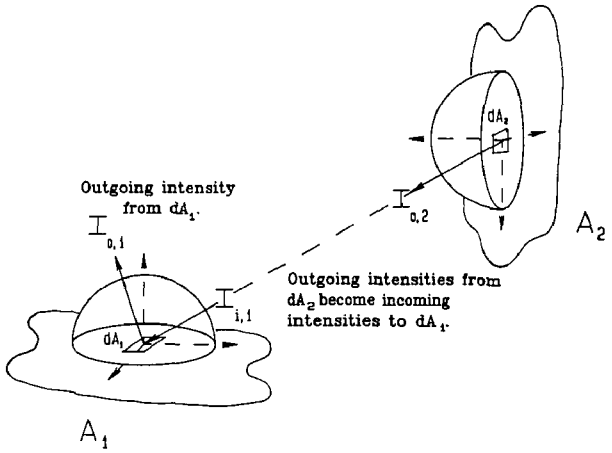

A simple environment of two surfaces, each with its own set of directions.

FIGURE 2

directional intensity in direction f then becomes its emission in direction f plus the sum of the reflected intensities from discrete directions d, which encompasses the hemisphere above the patch (Equation 10). d^{-1} represents the outgoing directions from patch 2 which point towards patch 1, and are thus in the opposite direction from d.

$$I_{t,1}(f) = \varepsilon_1(f) + \sum_{d=1}^{D} \rho_1''(f, d) I_{t,2}(d^{-1}) cos\theta_d \omega_d \quad (10)$$

Figure 3 shows the results of two incoming directions with several outgoing directions from a patch.

A set of simultaneous linear equations can now be formed, defining the unknown directional intensities for all directions, given the bidirectional reflectivities and emissions of each patch.

So far, the method has been illustrated by an environment consisting of only two surfaces. The form of the relationships remains the same for more complex environments except that light

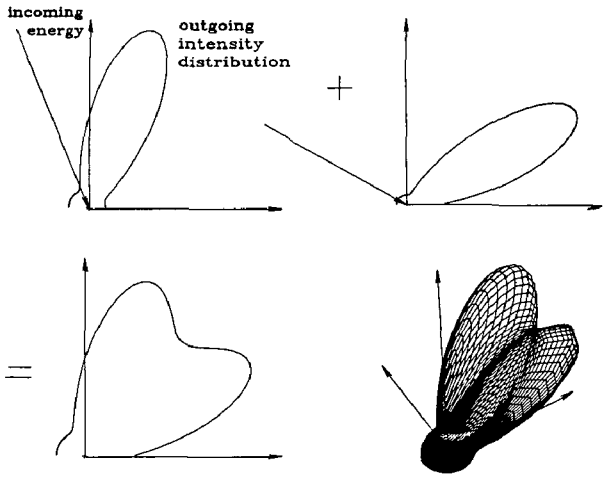

The outgoing distributions from all incoming energies are summed, resulting in an aggregate distribution.

FIGURE 3

arrives at a patch from many surfaces. An additional problem arises from the possible occlusion of one patch from another due to intervening surfaces. Thus, the directional intensities of each patch depend on the directional intensities of all other patches and the visibility between them. Let N be the number of patches in an environment and D be the number of directions on the hemisphere. The outgoing intensity of a patch i is given by:

$$I_{i,1}(f) = \varepsilon_i(f) + \sum_{j=1}^{N}\sum_{d=1}^{D}\rho_i''(f,d)I_{i,j}(d^{-1})\cos\theta_d\omega_d HID(i,j,d) \quad (11)$$

where

$$HID(i,j,d) = \begin{cases} 1 & \text{Patch } i \text{ sees patch } j \text{ in direction } d \\ 0 & \text{Otherwise} \end{cases} \quad (12)$$

for environments with hidden surfaces.

At first glance, Equation 11 appears immensely more complicated than previous radiosity solutions, since intensities must be solved for each patch for each direction. However, certain properties of the matrix and the environment can be effectively exploited to obtain tractable solutions. For example, if in an environment, there are 1000 patches, and the hemisphere of directions is divided into 1000 small solid angles, then there are 1,000,000 unknowns and 1,000,000 squared or 10^{12} matrix coefficients. Fortunately, this matrix is VERY sparse. This can easily be seen since each patch receives energy from only 1000 directions. Thus, each row of the coefficient matrix has only 1000 non-zero HID terms, and the matrix is 99.9% sparse. In addition, if some surfaces are diffuse, each of their patches will exhibit a single constant directional intensity over all directions, and the 1000 directional equations can be reduced to a single linear equation. These facts are taken advantage of in the implementation, which is described in the following sections.

PROGRAM OVERVIEW

The directional radiosity analysis is comprised of a series of separate steps analagous to previous diffuse radiosity methods (Figure 4). First, the environment geometry and material properties such as the bidirectional reflectance must be specified. An additional initial requirement for radiosity analyses is that each surface must be subdivided into discrete surface patches.

The second step is to determine the geometric and visibility relationship between every pair of patches in the environment.

Given the set of equations, a simultaneous solution for the directional intensities of each patch is then performed. This global patch solution is then used along with the vertex to patch visibility relationships to determine the directional intensity distribution at each vertex grid point on every surface. The resulting full description of the intensity leaving each surface in all directions provides the means to render an image of the environment from any given viewpoint.

Given an eye position and viewing direction, the intensities pointing back to the eye from each vertex are used to render an image. Succesive images of the environment can be rendered from different viewpoints.

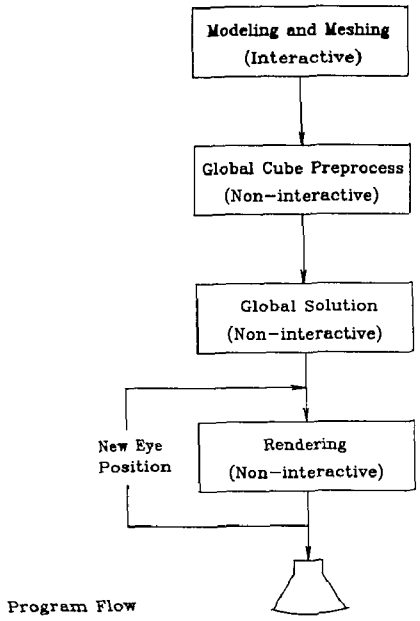

FIGURE 4

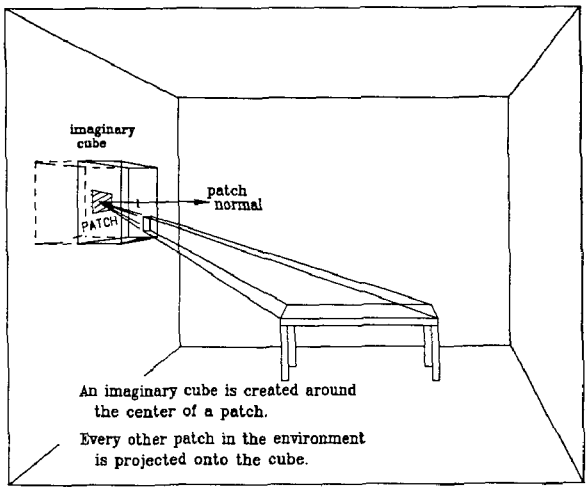

FIGURE 5

THE GLOBAL CUBE

The light leaving a patch depends on the light arriving at it and the direction(s) from which it arrives. This impinging light is the result of light leaving the other patches in the environment, and thus a determination must be made of which patches are visible from each patch. Each patch has a view of its environment across the hemisphere of directions surrounding its normal. This continuum of directions is approximated by finding the visible patches in a finite number of discrete solid angles which cumulatively are equivalent to the hemisphere.

The hemisphere can be replaced by the upper half of a cube, or "hemi-cube" (Figure 5), creating five separate 90 degree viewing frustums through which to project the environment [5].

Each surface of the cube becomes an imaginary screen which can be divided into small square cells. Each of these cells defines a small solid angle of view from the patch at the center of the cube. Depth buffer algorithms are performed for each frustum to determine the patch "seen" through each cell. To obtain an accurate view of the environment from a patch, a series of sample points within the patch from which to view the environment are selected. A grid is established on the patch defining a series of sub-patches or "elements". A view is taken from each grid vertex and may vary across the grid. The final view of the environment from a patch is taken as a weighted average of the views from the individual grid vertices. This is analogous to the "sub-structuring" described in [6].

The hemi-cube can be oriented so that the patch normal always coincides with the center of the hemi-cube's top face. This allows for the same set of small solid angles to be located in the same directions relative to each patch. However, since the orientations are always in the patch's local coordinate system, there is no single directional relationship between hemi-cubes (Figure 6). In previous diffuse radiosity methods, since each patch was defined to have a single intensity equal in all directions, the directional relationship was unimportant.

When the intensities are allowed to vary directionally, it is important to equate the incoming directions of the receiving patch with the reciprocal outgoing directions from the sending patches. Rather than reorienting a hemi-cube for each patch, a full cube surrounding each vertex is maintained, oriented along the global axes (Figure 7). The advantage lies in the fact that if patch i sees patch j through the cell oriented in the X,Y,Z direction, then the reciprocal cell from patch j is simply in the -X,-Y,-Z direction. The cube is stored as a vector of visible patch identification numbers. If the faces of the cube are oriented as in Figure 8a and the vector is ordered as in Figure 8b, then the reciprocal global cube locations may be trivially calculated. For a global cube with D discrete directions, the cell with the reciprocal direction from a cell with index c (0 <= C <= D-1) will simply have index D - c - 1.

The projection of the environment on a "global cube" surrounding a vertex solves the hidden surface problem providing the HID terms in equation 11. The θ terms are found by the dot product of the patch normal and the global orientation of the cell. A small solid angle, ω, is associated with each cell, providing the necessary information to simultaneously solve for the directional radiosities of each patch.

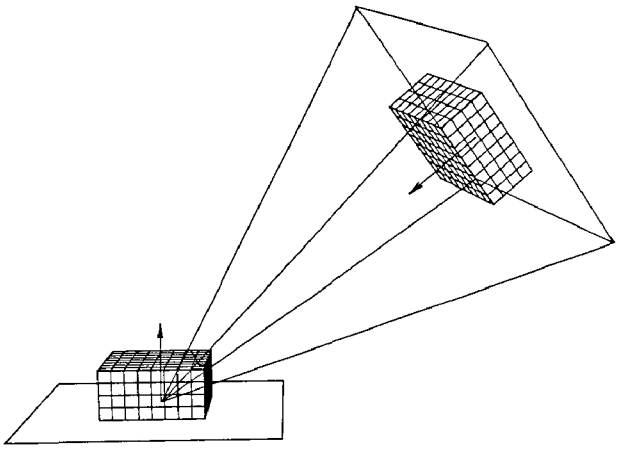

Each hemi-cube has its own local orientation.

FIGURE 6

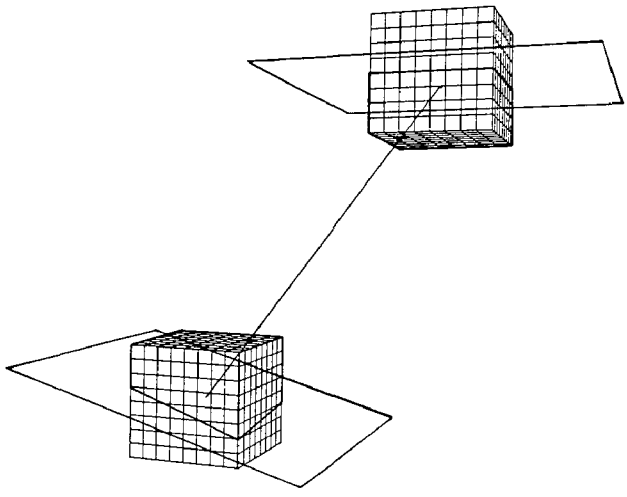

Each global cube is placed in a global orientation.

FIGURE 7

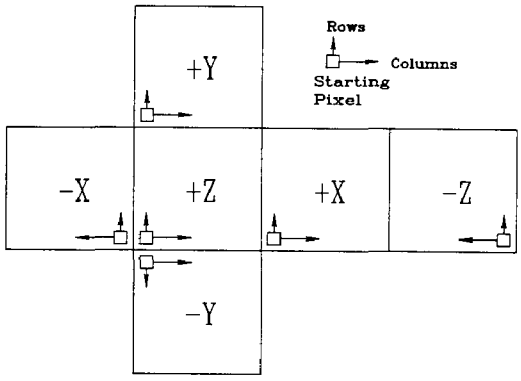

The cells within the faces of the global cube are ordered so that the reciprocal global cube locations may be trivially calculated.

(a)

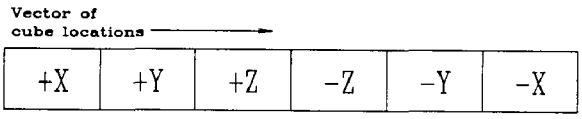

The ordering within each face is in row major order. The rows and columns are determined by the diagram above.

(b)

FIGURE 8

BIDIRECTIONAL REFLECTANCE

Until now, the actual bidirectional reflectance, which is derived from a model describing how light scatters from a surface, has not been specified. The choice of a specific reflection model is independent of the methods introduced in this paper. However, in this context there are two requirements a reflection model must satisfy. The first requirement is that the model must exhibit reciprocity. That is, the fraction of outgoing intensity to incoming energy for any pair of directions must be the same if the incoming and outgoing directions are reversed.

In other words:

$$\rho''(\Psi_r, \Psi_i) = \rho''(\Psi_i, \Psi_r) \qquad (13)$$

The second requirement is that energy must be conserved. The total amount of incident energy to a patch must be greater than or equal to the total amount of reflected energy from the patch. A patch cannot reflect more energy than it receives. Thus,

$$E_i \geq \int_{\Omega_h} E_r = \int_{\Omega_h} I_r \cos\theta d\omega \qquad \Omega_h = \text{hemisphere} \quad (14)$$

Two different bidirectional reflectance relationships were tested. One relationship is a mirror-like reflection, in which the incoming intensity leaves in the mirror direction tempered only by a reflectance between 0 and 1. The other relationship is based on work done by Phong [8]. The "Phong-like" bidirectional reflectance used is:

$$\rho'' = \frac{k_d}{S_d} + \left[\frac{k_s}{S_s}\right] \cos^n\phi \qquad (15)$$

where k_d = the fraction of energy diffusely reflected, and k_s = the fraction of energy specularly reflected: $k_d + k_s \leq 1$. ϕ is the angle between the mirror direction and the reflection direction, r, and S_d and S_s are constant factors which must be included to conserve energy:

$$S_d = \int_{\Omega_h} d\omega = \pi \quad \text{and} \quad S_s = \int_{\Omega_h} \cos^{(n+1)}\phi \, d\omega \quad (16)$$

GLOBAL SOLUTION

A simultaneous solution to the matrix of Equations 11 is required to find all of the directional radiosities for all of the surfaces of an environment. The matrix to be solved is too large to be solved using direct solution methods. However, the Gauss-Siedel iterative technique, which was found effective for previous radiosity methods is well suited for this problem [9]. The iteration method allows the matrix coefficients to be computed as they are needed. The solution converges in a small number of iterations since the matrix is diagonally dominant.

An initial estimate for the radiosity solution is set equal to the emissions of the light sources with all other patch radiosities initially zero. The global cube analysis provides the patches and directions from those patches which illuminate a particular patch. Based on this information and the current solution estimate, a global cube of incoming intensities to a patch through all directions is assembled. From Equation 11, each outgoing intensity can be computed from these incoming intensities and the given bidirectional reflectance of the patch. These new directional patch radiosities are used in the current solution estimate for subsequent patches. Every patch is solved in turn, and the whole process iterates until the solution converges.

After every iteration step, each patch receives a more accurate set of incoming intensities, and therefore, reflects a more accurate set of outgoing intensities (Figure 9). Every directional intensity value is compared with the same directional intensity value from the previous iteration. If no pair of values differ by more than a prespecified margin, the solution process is said to have converged, and no further iterations are required. In pseudo-code, the procedure is as follows:

```
/* Initial estimate for directional intensities */
  FOR each direction d and patch p DO
    Int[p][d] = Emiss[p][d]
  END
  FOR each iteration DO
    FOR each patch p DO
    /* Set each new intensity to be its emission */
      FOR each direction d DO
        Newint[p][d] = Emiss[p][d]
      END
    /* Take energy from incoming directions in and */
    /* distribute it to all outgoing directions out */
      FOR each direction in DO
      /* Calculate energy from in */
        Energy[in] = Int[Cube[p][in]][D-in-1] cosθω[in]
        For each direction out DO
          Newint[p][out] += ρ''(p, in, out)Energy[in]
        END
      END
    /* Replace old intensities with new and test */
    /* for convergence */
      FOR each direction d DO
        Compare Int[p][d] with Newint[p][d]
        Int[p][d] = Newint[p][d]
      END
    END
  END
```

The number of iterations required for convergence of the patch solution depends on the order in which patch equations are solved since the incoming intensities to a patch result from the outgoing intensities of previously solved patches. It thus makes sense to solve for the patches in roughly the same order that light is propagated through the environment (Figure 10). Thus, the emitters are solved first, because they are the originators of energy for the solution; next, the patches that see the emitters are solved, and thus become secondary emitters; next, those that see the secondary emitters; etc. The results of the global cube provide the visible surfaces seen from each surface. A list is compiled of the secondary light sources by examining the global cubes of the primary emitters. Analogously, a list of tertiary light sources is compiled from the global cubes of

the secondary sources, etc.

The patch solution provides the intensity of light leaving each patch in each direction. At this point, an analogous solution must be found for each grid vertex, as it is the vertex directional intensities that will be used during the rendering process. Each grid vertex has an associated global cube of visible patches, from the original global cube analysis. A cube of directional intensities for each grid vertex is filled by back substitution of the previously solved patch intensities. The outgoing intensities are calculated from the incoming intensities and the bidirectional reflectance of the vertex's surface.

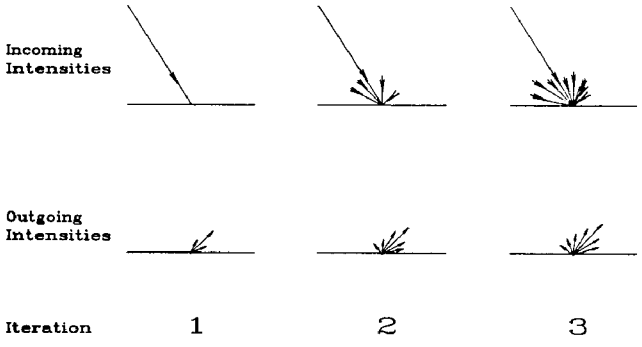

After each iteration, each patch receives a more accurate set of incoming intensities, and therefore, reflects a more accurate set of outgoing intensities.

FIGURE 9

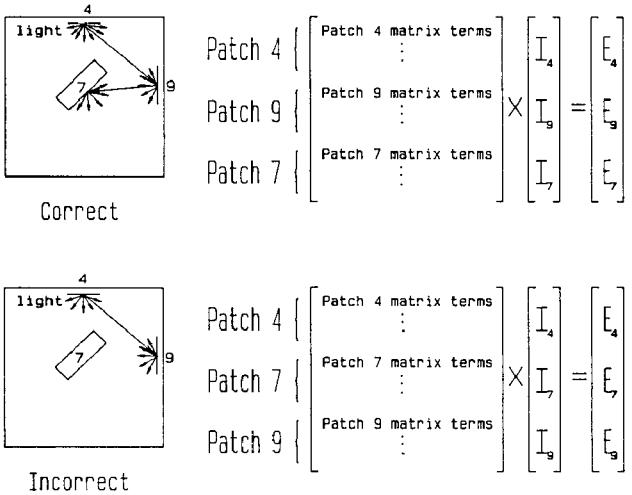

To decrease the number of iterations, solve for the patches in roughly the same order that light is propogated through the environment.

FIGURE 10

RENDERING

At this point, a view independent radiosity solution has been calculated. The final rendering stage is simply a post process to display intensities pointing in the direction of the eye. To render an image, an eye position and viewing direction are specified which establishes the directions from each vertex back to the eye. The particular intensities pointing to the eye are extracted from the global cubes of directional intensities (Figure 11). These are obtained by bilinearly interpolating between the nearest vertex cube directional intensities. This results in a single intensity value for each grid vertex in the environment from that viewpoint. At this point the rendering process becomes identical to the process described in [6], where the pixel intensities are bilinearly interpolated from the vertex intensities. Additional views of the environment are rendered from the same global radiosity solution simply by specifying a new view point.

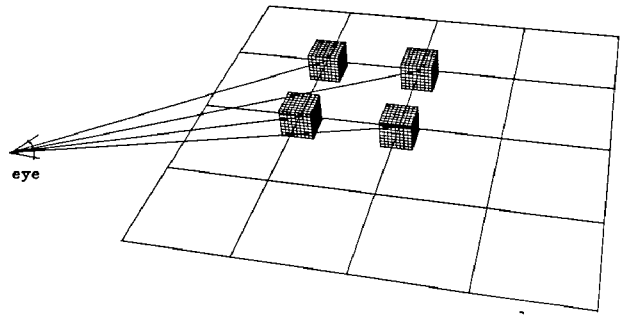

The particular intensities pointing to the eye are extracted from the global cubes of directional intensities.

FIGURE 11

RESULTS

The following images were rendered from two global solutions of the same environment. Each set of figures contains a series of views (a) illustrating the view independence of the solution. Notice the reflections in the floor tracking the eye as the viewer position is changed. The two larger images (b and c) in each set are enlargements from the series above them. The primary difference between the two sets of images is the bidirectional reflectance models. The algorithms outlined in this paper are independent of any particular reflectance model. The solution for the images in Figure 12 used the Phong-like reflectance model. For Figure 13, the solution used the simpler mirror-like bidirectional reflectance.

The environmental modeling and meshing for the images was performed using the testbed image synthesis system at Cornell's Program of Computer Graphics. The system was written in C under a VMS operating system and ran on a VAX 11/780. An Evans and Sutherland picture system was used for the modeling and meshing, a Floating Point Systems FPS264 attached to a VAX 11/750 was used for the solution process, and the pictures were displayed using a Rastertek 1280x1024x24 bits frame buffer.

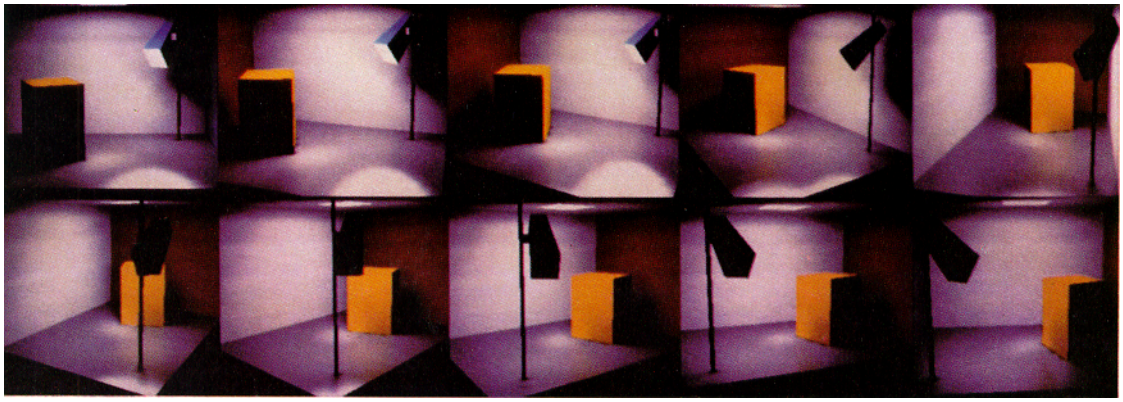

(a)

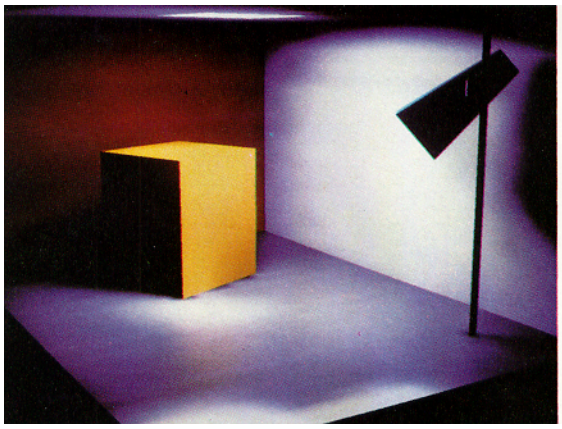

(b)

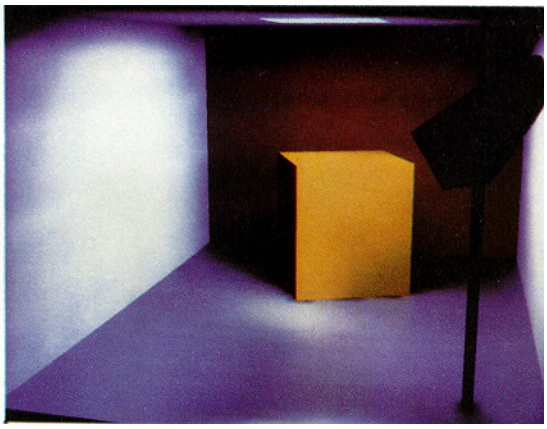

(c)

PHONG-LIKE REFLECTION

Number of Specular Patches	: 64
Number of Diffuse Patches	: 93
Number of Specular Vertices	: 1089
Number of Diffuse Vertices	: 1400
Number of Cube Directions	: 24576
Number of Wavelengths	: 4

PROCESS	CPUTIME (Hrs)
Visible Surfaces	8
Solution	192
Render	1

FIGURE 12

As can be seen from the statistics listed below each set of images, the computation time, particularly for the solution process, is non-trivial. The time required for the solution process is a function of the number of cube directions and the number of patches and vertices, more specifically, the number of specular patches and vertices. The amount of time is significantly reduced by using the simpler mirror-like reflection model in Figure 13 since incoming energy is not scattered in all directions.

An artifact due to the nature of the algorithm becomes visible when the reflection becomes more mirror-like. Recall that the rendering is based on the vertex solution which is in turn based on the patch solution. Thus, the discrete patches become visible in the reflection. This resulted in the increased number of diffuse patches in Figure 13.

CONCLUSION

A general radiosity method which accounts for all interreflections from both diffuse and non-diffuse surfaces has been introduced. The procedure computes an intensity distribution for every surface in an environment, consisting of directional intensities for a number of discrete directions. The rendering process then becomes one of only looking up the intensities that point back to the eye and displaying them.

This method makes it possible to precalculate the energy leaving specular surfaces regardless of the viewing direction. Since this global solution is view independent, successive views of the same environment can be calculated using the same directional intensity information, and the phenomenon of reflection tracking can be observed.

The addition of specular surfaces to the radiosity method is significant, since radiosity methods need no longer be limited to environments

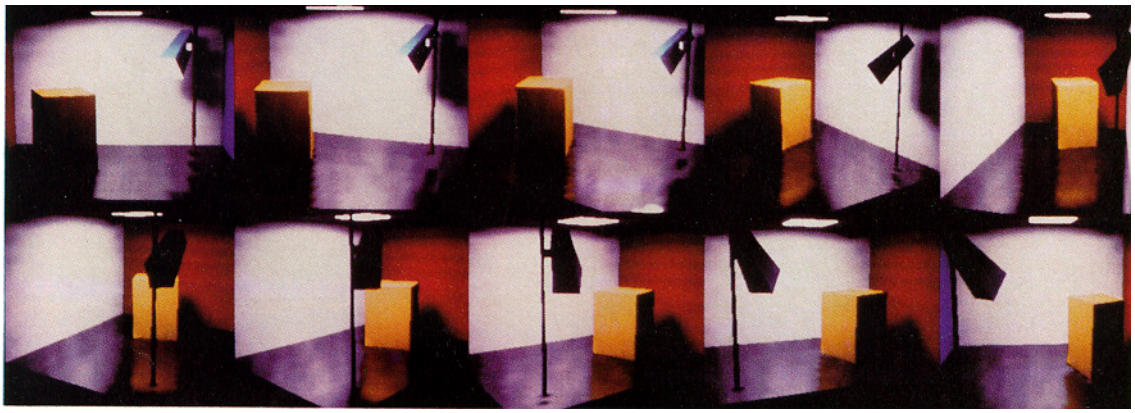

(a)

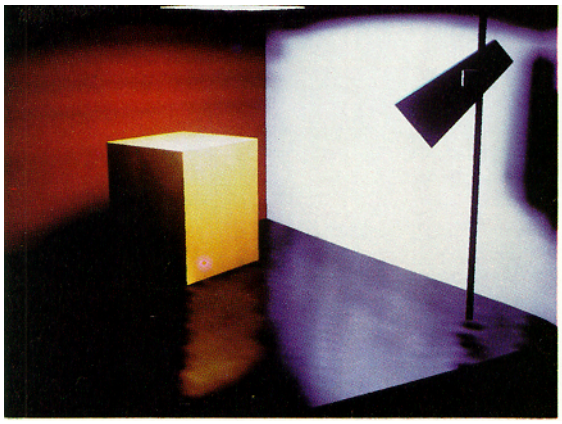

(b)

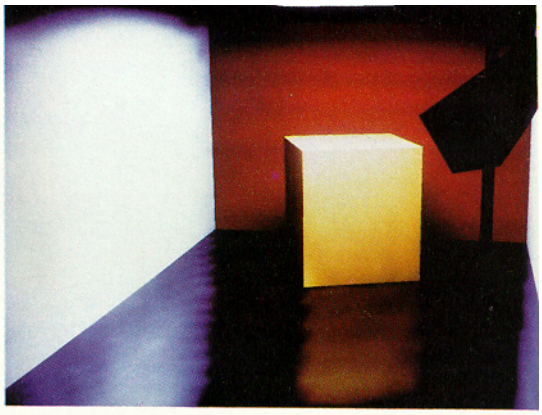

(c)

MIRROR-LIKE REFLECTION

		PROCESS	*CPU TIME (Hrs)*
Number of Specular Patches	**: 64**		
Number of Diffuse Patches	**: 237**		
Number of Specular Vertices	**: 1089**		
Number of Diffuse Vertices	**: 1400**	**Visible Surfaces**	**7**
Number of Cube Directions	**: 15000**	**Solution**	**25**
Number of Wavelengths	**: 4**	**Render**	**1**

FIGURE 13

containing only purely diffuse surfaces. The methodology is independent of a particular reflection model, thus any reflection model which relates incoming energy from one direction to outgoing intensity in another direction can be used. The only restriction placed on the model is that it must conform to the physical laws of reciprocity and the conservation of energy as discussed in the paper.

A few problems exist that must be dealt with in the future. Errors in approximation occur because environments have been discretized. Discretizing the number of directions on the global cube results in misrepresenting the incoming energy because of aliasing on the faces of the global cube. Reflections from a highly reflective surface may appear blurred due to discretizing the surface, as well as interpolating between outgoing intensities in rendering.

Future work should include the addition of non-planar and textured surfaces, as well as transparent and translucent surfaces. The framework of the method is general enough to accomodate these additions by simply using different reflection functions. Although this method currently requires very large computational resources for complex environments due to the volume of information that must be processed, this research was undertaken to provide a scientific basis for modeling complex interreflections within an environment. In the future, as storage and computational power increases, global solutions of more complex environments will be within reach.

ACKNOWLEDGMENTS

The authors would like to thank all those who helped in the preparation of the article. Thanks go to Professor Kenneth Torrance, Holly Rushmeier, and Kevin Koestner for their help in laying the theoretical groundwork. Thanks also go to Dan Baum, Janet Brown-Aist, and Emil Ghinger for their help in coding and preparation of the article. This research was conducted at the Cornell University Program of Computer Graphics, under a grant from the National Science Foundation #DCR-8203979. The VAX computers used were made possible by a large grant from the Digital Equipment Corporation and use of the Floating Point Systems Advanced Processor through Cornell University's Center for Theory and Simulation in Science and Engineering.

REFERENCES

[1] Whitted, Turner, "An Improved Illumination Model for Shaded Display," Communications of the ACM, Vol. 23, No. 6, June 1980, pp. 343-349.

[2] Siegel, Robert and John R. Howell, Thermal Radiation Heat Transfer, Hemisphere Publishing Corp., Washington DC., 1981.

[3] Sparrow, E. M. and R. D. Cess, Radiation Heat Transfer, Hemisphere Publishing Corp., Washington DC., 1978.

[4] Goral, Cindy M., Kenneth E. Torrance, Donald P. Greenberg, Bennet Battaile, "Modeling the Interaction of Light Between Diffuse Surfaces," ACM Computer Graphics (Proceedings 1984), pp. 213-222.

[5] Cohen, Michael F. and Donald P. Greenberg, "A Radiosity Solution for Complex Environments," ACM Computer Graphics (Proceedings 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, March 1986.

[7] Nishita, Tomoyuki, and Eihachiro Nakamae, "Continuous Tone Representation of Three-Dimensional Objects Taking Account of Shadows and Interreflection," ACM Computer Graphics (Proceedings 1985), pp. 22-30.

[8] Phong, Bui Tuong, Illumination for Computer Generated Images, Ph.D. Dissertation, University of Utah, 1973.

[9] Hornbeck, Robert W., Numerical Methods, Quantum Publishers, New York, NY, 1974, pp. 101-106.