# An Experimental Evaluation of Computer Graphics Imagery

GARY W. MEYER, HOLLY E. RUSHMEIER, MICHAEL F. COHEN, DONALD P. GREENBERG, and KENNETH E. TORRANCE
Cornell University

Accurate simulation of light propagation within an environment and perceptually based imaging techniques are necessary for the creation of realistic images. A physical experiment that verifies the simulation of reflected light intensities for diffuse environments was conducted. Measurements of radiant energy flux densities are compared with predictions using the radiosity method for those physical environments. By using color science procedures the results of the light model simulation are then transformed to produce a color television image. The final image compares favorably with the original physical model. The experiment indicates that, when the physical model and the simulation were viewed through a view camera, subjects could not distinguish between them. The results and comparison of both test procedures are presented within this paper.

Categories and Subject Descriptors: I.2.10 [**Artificial Intelligence**]: Vision and Scene Understanding—*intensity, color, photometry, and thresholding*; I.3.3 [**Computer Graphics**]: Picture/Image Generation—*display algorithms; viewing algorithms*; I.3.6 [**Computer Graphics**]: Methodology and Techniques—*ergonomics*; I.3.7 [**Computer Graphics**]: Three-Dimensional Graphics and Realism—*color, shading, shadowing, and texture*; I.4.8 [**Image Processing**]: Scene Analysis—*photometry*

General Terms: Experimentation, Human Factors, Measurement, Verification

Additional Key Words and Phrases: Color science, image science, light reflection models, radiosity

## 1. INTRODUCTION

The creation of realistic images requires an accurate simulation of light propagation within an environment, as well as a perceptually accurate method for displaying the results of the simulation. The need for physically based illumination models and perceptually based imaging techniques means that the lighting calculations and the production of the final simulation are separate tasks, each having different objectives to be met. If a scientific basis for the generation of images is to be established, it is necessary to conduct experimental verification on both the component steps and the final simulation.

This research was funded in part by National Science Foundation grant DCR 8203979, "Interactive Computer Graphics Input and Display Techniques."
Authors' present addresses: G. W. Meyer, Department of Computer and Information Science, University of Oregon, Eugene, OR 97403; H. E. Rushmeier, M. F. Cohen, D. P. Greenberg, and K. E. Torrance, Program of Computer Graphics, 120 Rand Hall, Cornell University, Ithaca, NY 14853-5501.

Early realistic image synthesis techniques and the lighting models that they employed were severely limited by processing and storage constraints, as well as by the display hardware characteristics. The need for computational simplicity substantially influenced the illumination algorithms that were originally developed. The results were light models that did not make direct use of established physical behavior; reflection models arbitrarily assigned ambient, diffuse, and specular portions to the reflected light. The perceptual significance of the monitor's primaries was not recognized as colors were directly computed in terms of the RGB (red, green, blue) primaries. Given specific viewing parameters, the intensity of each picture element was determined only on the basis of the single surface "seen" through that pixel and its direct relationship to light sources. These approaches, which do not simulate the global illumination effects and the interreflections among surfaces in an environment, result in pictures that are obviously computer generated.

Recently, ray-tracing techniques, which attempt to model the global illumination effects of specular surfaces, have been introduced. Ray tracing is still a view- and resolution-dependent approach, but employs a more comprehensive lighting model. Each picture element can receive light directly from the surface immediately behind it and indirectly by ray reflection (and/or refraction) from other objects. However, each participating surface still receives its illumination only in a direct path from light sources or from an arbitrary constant ambient term. In most cases the light model is expressed in terms of the RGB primaries and is not based on sound physical principles. Although the technique is quite expensive computationally, the pictures produced can be impressive and are a substantial improvement over those generated by previous techniques.

The introduction of the radiosity method has led to a complete decoupling of the light reflection simulation from the final imaging technique. An illumination model based on energy conservation principles is used to account for all interreflection of light in an environment. The illumination calculations are independent of viewing parameters and can be performed on a wavelength basis rather than the particular red, green, and blue channels provided by the phosphors of a specific raster display device. The results of the global illumination calculations are used in conjunction with the principles of color science to convert the resulting spectral energy distributions to the RGB primaries of the display device.

What has emerged from this sequence of events is the need for a clear distinction between the physical and perceptual portions of the image synthesis process and the need for experimental verification of each of these steps. The first step in the image synthesis process should be to model correctly the transport of light in the environment. This is inherently a physically based step where the flow of energy is modeled as accurately as possible. To verify the light model, physical measurements should be made on a real scene and should be compared with the simulated values. The second step of the image synthesis process should be to use the results from the physical modeling of the propagation of light to produce the final simulation to be observed. This is inherently a perceptually based step, where the objective is to satisfy the final observer. To verify the final simulation and thereby the overall objective of realistic image synthesis, the simulation should be visually compared with the real scene.

In this paper a simple environment is used to demonstrate an approach to image synthesis that has distinct physical and perceptual portions and that employs experiments to verify both parts of the process. In Section 2, the radiosity method is used to do the light modeling, and the results are compared against physical measurements made on an actual model. In Section 3, the principles of color science are used to produce an image of the same model on a color television monitor, and this picture is visually compared with the real scene by a group of experimental subjects. The observations and conclusions of the paper, summarized in Section 4, indicate that by using a rigorous scientific methodology a good match can be obtained for both the physical and perceptual comparisons.

## 2. RADIOMETRIC COMPARISON

In this section an example is presented of the use of a physical experiment to verify the first part of the image synthesis process—the simulation of reflected light intensities. The distribution of radiation in simple scenes is considered, and the particular theoretical procedure for calculating the radiant transfer to be verified is outlined. This technique, known as the radiosity method, is used to generate all of the synthetic computer images in this paper. An experimental apparatus is also described. This apparatus allows simple, real-world scenes to be tested and is used for all of the scenes presented in this study. Measurements of radiant energy flux densities on a wall of the physical model are compared with the predictions of the radiosity method; a method for measuring the radiant flux densities, which are directly related to the light intensities, is detailed, and measurements on three environments of varying complexity are presented.

### 2.1 Overview of Experimental Design

In an ideal experiment for verifying the accuracy of light intensity calculations on an image plane, an instrument would be used that could be positioned at the "eye" position with respect to the real environment. This instrument would have an angular resolution that would allow it to measure the light energy reaching the "eye" through the solid angle subtended by each pixel in the image plane. This instrument would also have the ability to measure each wavelength band of light reaching the eye.

The instrument defined above would need precise angular and spectral resolution. The associated measurements would be geometrically difficult and time consuming, and would require high photon sensitivity under very carefully controlled lighting conditions. Since the present study is an initial effort to compare a real environment with a synthetic computer image, such a refined experimental study was not carried out. Indeed, a relatively inexpensive and simple radiation measuring instrument (a radiometer) was employed. The instrument gave a single reading corresponding to the hemispherically incident radiant flux over the range of visible wavelengths.

Measuring an entire environment or scene from the "eye" position with this instrument would yield a single reading on a meter. Since this single reading represents a spatial and spectral average of the flux incident on the radiometer, it would not be sufficiently discriminating to allow an evaluation of simulation methods. It represents a point measurement, which is indicative only of the

magnitude of the radiant field. A more discriminating approach requires measurements at several locations to assess the spatial distribution of light energy. Thus measurements would be needed for many different viewing locations.

Furthermore, the light received by a radiometer varies continuously with position and depends on the geometric and optical properties of the entire radiant environment. Although such measurements do not allow direct verification of the detailed predictions of a lighting model, they do allow verification of the integral predictions (i.e., integrated over wavelength and the incident hemisphere) of a lighting model. Such integral measurements at several locations are employed in this study to assess a particular lighting model. In general, a lighting model must be capable of predicting the relative values of these integral quantities if it is to be relied upon to simulate accurately the more detailed light intensities required for image synthesis.

## 2.2 Radiosity and Irradiation

The above radiometric method is used in the present article to evaluate the standard radiosity method and one variation of the radiosity method. The radiosity method is a theoretical procedure for predicting light intensities in a totally diffuse environment. The method was developed in the field of heat transfer to calculate the heat exchange by means of electromagnetic radiation in enclosures. It can also be applied to visible light. The method was first applied to synthetic image generation by Goral et al. [5], and extended by Cohen and Greenberg [2]. In this paper, the radiosity method is used to predict the light energy impinging on, and measured by, the radiometer. A brief summary of the radiosity method is included as background material for the experiments.

In the radiosity method, all emission and reflection processes are assumed to be perfectly diffuse (Lambertian). The scene or enclosure is divided up into discrete surfaces, each of which is assumed to be of uniform radiant intensity. With these assumptions, the intensity of radiation leaving a particular surface is directly proportional to the radiant flux density (energy per unit area per unit time) or radiosity $B$ leaving the surface. The radiosity of a surface $i$ in an enclosure is related to the radiosities of all the surfaces in an enclosure by

$$B_{i\lambda} = E_{i\lambda} + \rho_{i\lambda} \sum_j F_{ij} B_{j\lambda}, \tag{1}$$

where $\lambda$ denotes wavelength, $E_{i\lambda}$ denotes the energy emitted from the surface per unit time and area, $\rho_{i\lambda}$ is the diffuse reflectance of the surface, and the summation $j$ is over all the surfaces in the enclosure. There is one such equation for each surface. The form factor $F_{ij}$ depends only on geometry and represents the fraction of energy leaving surface $i$ that arrives at surface $j$. The energy source term $E_{i\lambda}$ is zero for surfaces that are not light sources. Equation (1) holds for a particular wavelength. However, it also applies for discrete wavelength bands in which $B_{i\lambda}$, $E_{i\lambda}$, and $\rho_{i\lambda}$ are constant, as long as energy is not exchanged between the bands.

The basic radiosity method can be extended to account for directional variations in the light source. The extension is achieved by computing the amount of emitted light directly reaching the surfaces illuminated by the light source. If reflections off the light source are neglected, the equation corresponding to the

light source can be set aside. The radiosity equation for the other surfaces becomes

$$B_{i\lambda} = \rho_{i\lambda}\{\mathrm{df}\ F_{i,\mathrm{light}}\ \max[E_{\mathrm{light},\lambda}] + \sum_j F_{ij}B_{j\lambda}\}, \tag{2}$$

where the index $i$ and the summation $j$ do not include the light source, df is a light source directional factor, and the maximum directional radiosity of the light source is denoted by $\max[E_{\mathrm{light},\lambda}]$ and is found from measurements. The directional factor (df) is zero for surfaces not directly illuminated by the light source.

As discussed above, it is difficult to make spatially and spectrally detailed measurements of the radiant energy in an environment. In the heat transfer literature, direct measurements of radiosity (the energy leaving a surface) are rarely found [10]. The approach of measuring irradiation is more commonly used [9]. The irradiation $H_{i\lambda}$ incident on a surface $i$ is given in terms of the radiosities of all the other surfaces by

$$H_{i\lambda} = \sum_j F_{ij}B_{j\lambda}. \tag{3}$$

For surfaces that are not light sources, comparison with eq. (1) shows that the radiosity of surface $i$ is directly proportional to the irradiation onto the surface. The constant of proportionality is the reflectivity $\rho_{i\lambda}$ of the surface. Since intensity is proportional to radiosity, the intensity of the surface is directly proportional to the irradiation. Thus the relative spatial distributions of the incident irradiation and reflected intensity are the same.

Irradiation can be measured by a radiometric probe. If the sensitivity of the probe varies with direction, however, the probe response cannot be compared directly with eq. (3). Instead, the radiosity $B$ at a particular angle of incidence must be multiplied by an angle-dependent correction factor cf. The predicted response of the probe is then given by

$$H_{i\lambda} = \sum_j \mathrm{cf}(\theta)F_{ij}B_{j\lambda}, \tag{4}$$

where $\theta$ denotes the angle of incidence on surface $i$ of irradiation coming from surface $j$. The predictions of this equation are compared later with radiometric measurements.

### 2.3 Experimental Apparatus

The test environment was the five-sided cube shown in Figure 1. The dimensions of the cube are shown in Figure 2, as are the dimensions of two small boxes that were placed within the cube for some observations. All five sides of the cube could be removed independently so that the color of each side could be changed. All of the surfaces of the cube were painted with flat latex house paints, which are close to being ideal diffuse reflectors. The spectral reflectances of the paints were measured using a Varian Cary 219 spectrophotometer with an in-cell space diffuse reflectance accessory. These reflectances are shown in Figure 3a.

The light source consisted of a 150-watt incandescent flood light mounted at the top of a 15-inch-high metal cone. The interior of the cone was covered with a flat white paint. The light shone through a piece of 4.5 by 3.5-inch flashed opal

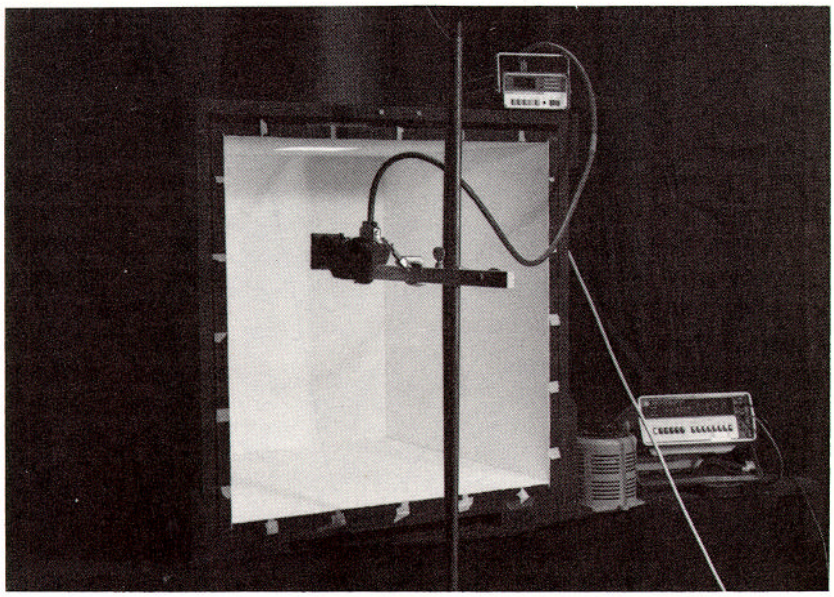

Fig. 1. Experimental setup used to make radiometric measurements.

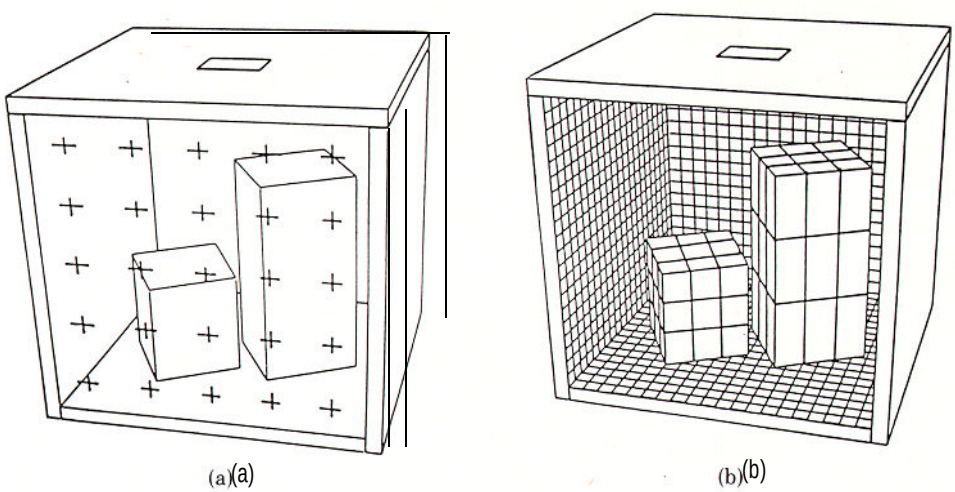


Fig. 2. Schematic of test environment. (a) The crosses indicate positions for radiometric measurements. (b) The surfaces are discretized for radiosity calculations.

| Dimensions: | Width (inches) | Height (inches) | Depth (inches) | Color |
|---|---|---|---|---|
| Enclosure | 21.6 | 21.5 | 22.1 | White; extra red/blue walls |
| Large box | 6.5 | 13.0 | 6.5 | White |
| Small box | 6.5 | 6.5 | 6.5 | Yellow |
| Light | 4.5 | — | 3.5 | See Figure 3b. |

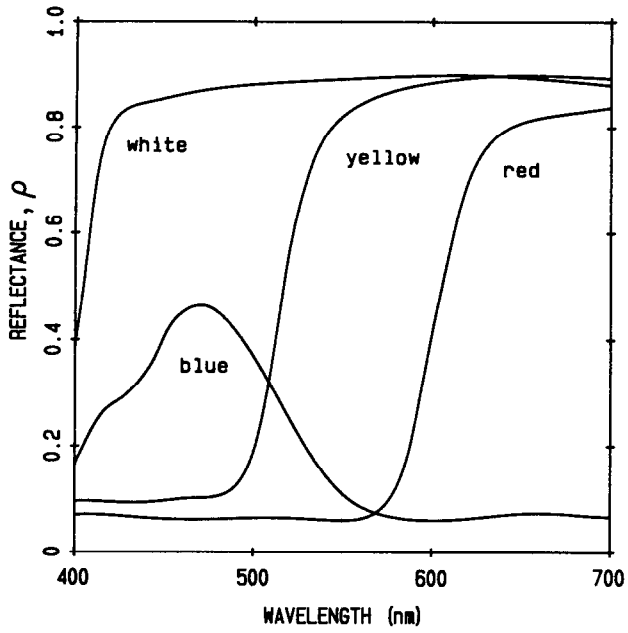


(a)

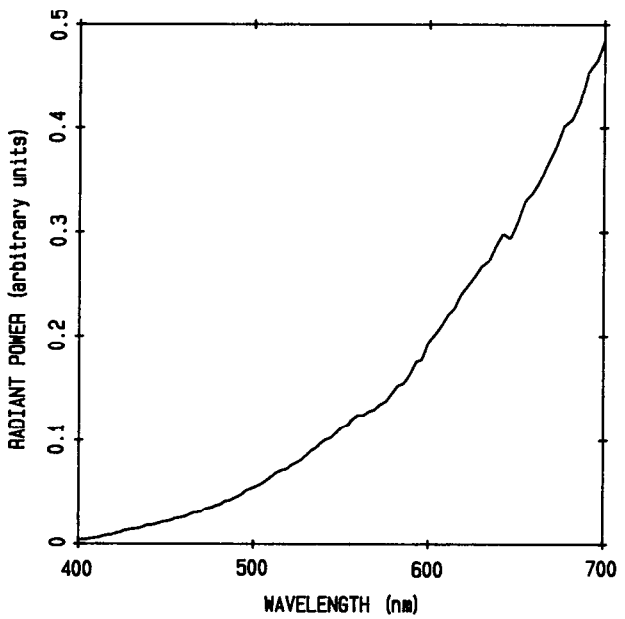


(b)

Fig. 3. (a) Reflectances of paints used to paint cube and small boxes. (b) Spectral energy distribution of light after passing through opal glass.

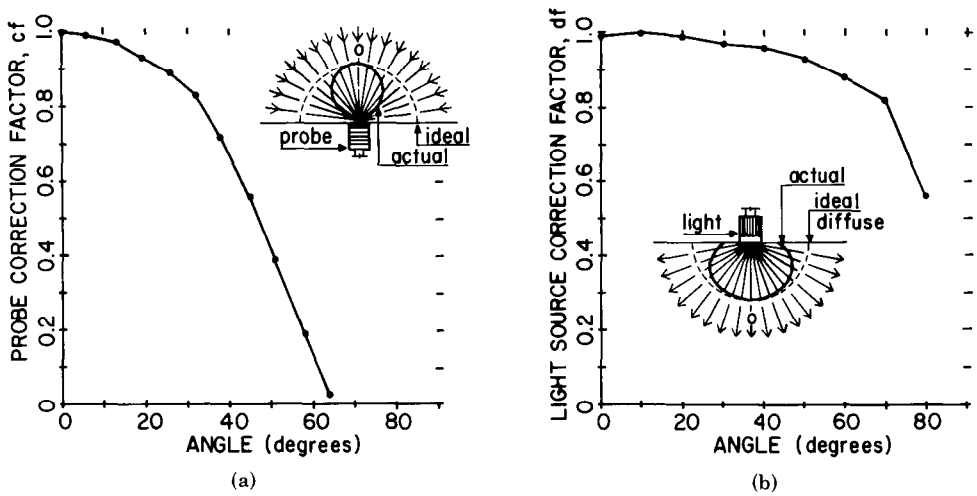


Fig. 4. (a) Probe directional sensitivity. (b) Normalized light source intensity versus angle.

glass, which was mounted in the ceiling of the cube. An autotransformer and digital voltmeter were used with the light source to maintain a constant 115 volts. The spectral energy distribution of the entire light source assembly was measured using equipment described by Imhoff [6] and is shown in Figure 3b.

The enclosure was placed on a flat black table in a small room. The walls of the room were covered with black fabric so that essentially no visible radiation entered the cube through the open side. From the inside of the enclosure, the open side appeared as a black wall.

Irradiation was measured using a Tektronix J16 photometer with a J6502 irradiance probe. This probe has a flat spectral sensitivity in the visible and near infrared ranges. A Corning Glass 1-56 filter was placed on the front of the probe to filter out the infrared energy emitted by the light source. The directional sensitivity of the probe (and filter) was determined by rotating the probe while illuminating it with collimated light. The correction factor cf as a function of incident angle is shown in Figure 4a.

To complete the apparatus specification, an additional measurement of the light source strength is required. This was measured for visible light by holding the J6502 probe flush against the opal glass. For a perfectly diffuse light source, the measured irradiation $H$ can be related to the total light source emission $E_{\text{light}}$ by

$$E_{\text{light}} = \frac{H}{F_{\text{sensor,light}}}. \tag{5}$$

The form factor between the probe sensor and the light, $F_{\text{sensor,light}}$, was estimated to be 0.52 [11, p. 826].

To use eq. (2), the directionality of the light source df is also needed. The light intensity at various angles from the normal was measured by using the photometer with the irradiance probe. The probe was fitted with a long tube to restrict

the acceptance angle of the probe. The results of these measurements are shown in Figure 4b and indicate that the light source is not perfectly diffuse. A perfectly diffuse light source would have a df of unity.

### 2.4 Procedure

Measurements of irradiation were made at 25 locations in the plane of the open face of the cube (shown in Figure 2a) and compared with the simulations. Measurements were made for three scenes: the empty white cube, the empty white cube with the left panel replaced by a blue panel, and the all-white cube with the large white box inside it.

The measurement locations were chosen for two reasons: (1) to maximize the light energy incident at any point, and thus to minimize the uncertainty in each reading, and (2) to minimize the effect of the probe on the environment. The probe should cast no shadows and should reflect little light back into the environment. The foregoing criteria are satisfied by placing the probe at the open side of the cube.

Tests were made to examine the potential sources of error in the measurements. The light source voltage could be controlled so that variations in light source emission changed by less than 1 percent. Movement of objects within the room surrounding the cube and changing the position of the cube within the room had no measurable effect on the irradiation at the open face of the cube. Doubling or tripling the size of cracks between the panels had no measurable effect.

Another potential source of error in the measurement was the position of the probe. The three-dimensional position and angular orientation of the probe were carefully controlled. Very small variations in these parameters could result in large variations in the measured irradiation. This error was estimated by repositioning the probe at each measurement location several times and recording the measured irradiation. The maximum deviation at each point was approximately ±5 percent of the mean value of the readings.

The error in the photometer itself is given by the manufacturer as less than ±5 percent. Combining the estimated error due to all of the foregoing factors leads to a total root mean square estimated error of ±7 percent.

### 2.5 Results

For each of the three scenes, measurements were made at the 25 test locations and compared with theoretical predictions based on the radiosity method.

For the radiosity calculations, the form factors were determined as described by Cohen and Greenberg [2]. In every case, each of the five walls of the cube was divided into 225 elements of equal area (see Figure 2b), the light source was divided into nine elements, and the sides of the large rectangular box in the third scene were each divided into nine elements. The reflectances used for the walls were averages of the measured spectral reflectance curves in Figure 3a. The open side of the cube was modeled as a black wall with zero reflectance. This side was divided into 25 surfaces, the center of each surface corresponding to a measurement position. The average measured irradiation when the probe was held flush against the light source was 240 microwatts per square centimeter. Using eq. (5), the total emission of the light source was estimated to be 460 microwatts/per square centimeter. This value, $\max[E_{\text{light}}]$, was used for all calculations. The

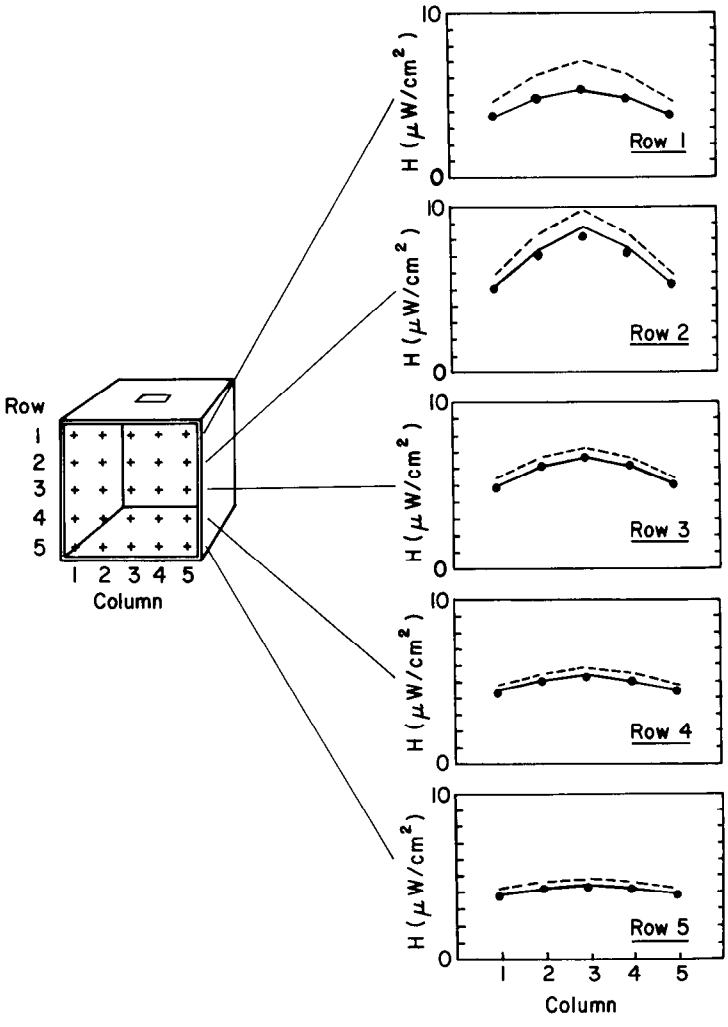


Fig. 5. Comparison of the measured and calculated irradiation at the open side of the empty white cube: ----, radiosity calculation with diffuse light source; ———, radiosity calculation with directional light source; •, radiometer measurement.

spectral distribution, $\max[E_{\text{light},\lambda}]$, was obtained from Figure 3b. The results predicted by the radiosity calculations were converted to an equivalent probe response by using eq. (4).

The measurements made with the empty all-white cube are shown in Figure 5 (filled circles). Each of the plots is for one horizontal row of locations. In general, the results show that the irradiation is highest near the center of the open side of the cube. This area has the best view of the light source and the other walls.

Figure 5 also shows the results of calculations using a completely diffuse light source (dashed lines). These calculations were made using eqs. (1) and (4). The calculated values are much higher than the measurements.

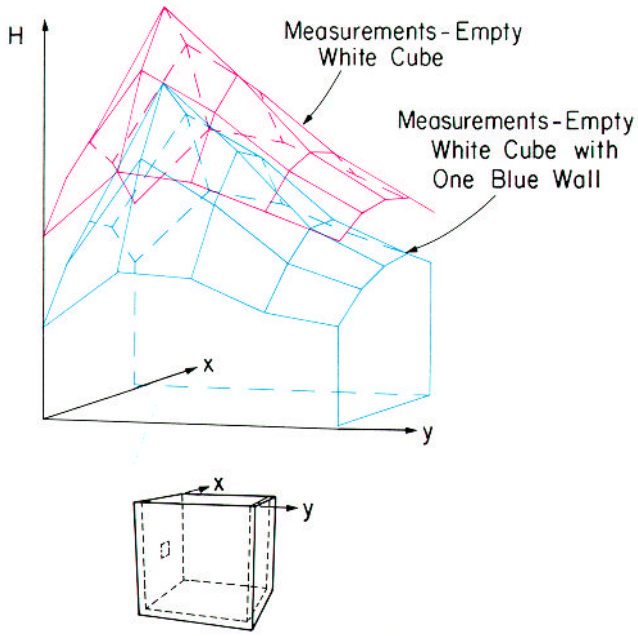


Fig. 6. Irradiation $H$ emerging from the open side of the cube. The sketch shows the cube with the open side facing upward and the top side to the left; the blue wall is on the side closest to the reader.

The major source of discrepancy between the measured results and those calculated for the purely diffuse environment is the directionality of the light source. This causes a difference in both curve shape and overall illumination level. Calculations using the diffuse radiosity method extended to include the directionality of the light source, as described by eq. (2), are also included in Figure 5 (solid lines). The calculated values are obviously lower than those for a purely diffuse light source, since less energy reaches each surface directly, and less energy is interreflected within the cube. The calculated results for the upper row show the largest changes since this row has the largest angle of incidence, and thus the greatest deviation, with respect to the light source emission. When light source directionality is accounted for, the root mean square difference between the calculated and measured results at the open face of the cube is less than 4 percent and the root mean square difference between normalized results is less than 3 percent. These values compare with 18 and 7 percent, respectively when the light source directionality is not accounted for, and are both less than the estimated measurement error of 7 percent. Thus the calculations made by assuming a directional light source are significantly more accurate than those made by assuming a perfectly diffuse light source. The rest of the calculated results presented in this section assumed a directional light source.

Figures 6-9 provide a comparison of the three scenes that were considered. In each figure the irradiation $H$ is shown as a function of measurement position. Figure 6 provides a comparison of the measurements on the empty white cube with measurements on an empty cube with one blue wall, the other walls being

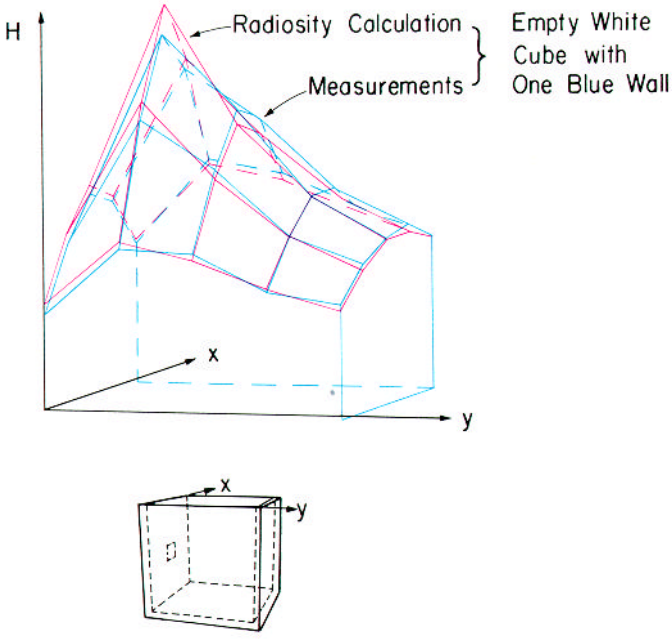


Fig. 7. Irradiation $H$ emerging from the open side of the cube. The sketch shows the cube with the open side facing upward and the top side to the left; the blue wall is on the side closest to the reader.

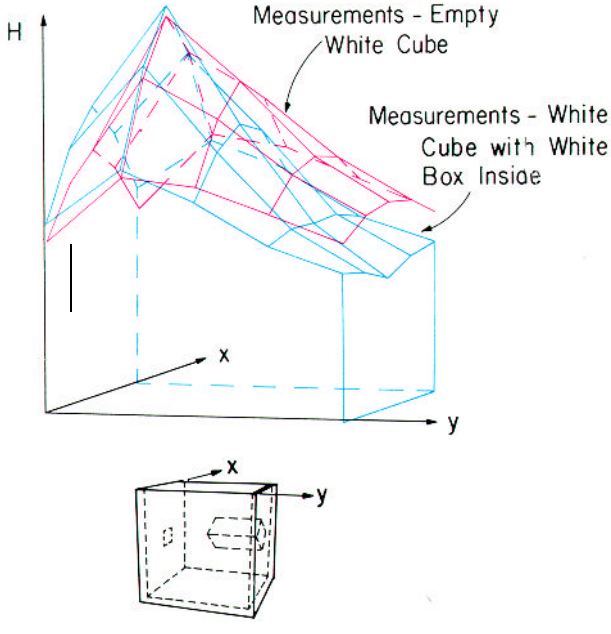


Fig. 8. Irradiation $H$ emerging from the open side of the cube. The sketch shows the cube with the open side facing upward and the top side to the left; the large white box is on the floor opposite the light source.

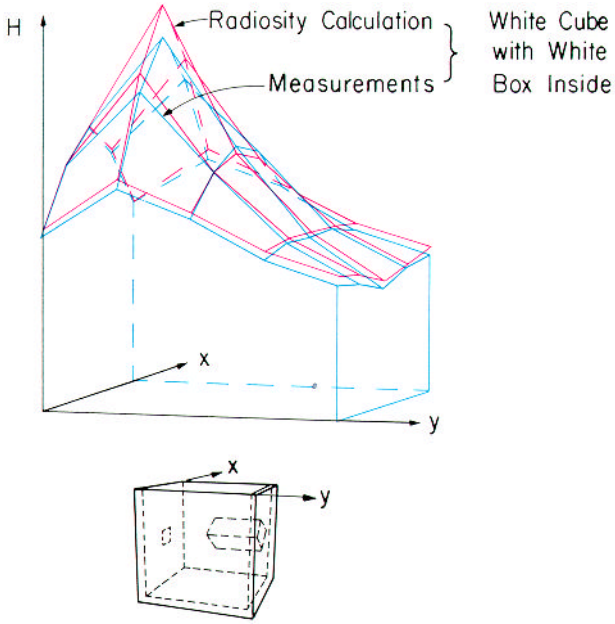


Fig. 9. Irradiation $H$ emerging from the open side of the cube. The sketch shows the cube with the open side facing upward and the top side to the left; the large white box is on the floor opposite the light source.

white. The low reflectance of the blue wall (Figure 3a) reduces the overall illumination in the cube. Clearly, this reveals that a large proportion of the light incident on any surface in the cube is due to reflection from other surfaces rather than direct illumination from the light source. Furthermore, measurements along the leftmost column of measuring positions (see Figure 2a) are only about 75 percent of the corresponding measurements on the right side of the cube. Thus small surfaces located on the left side of the open side of the cube would appear darker than surfaces located on the right side. This influence of the neighboring surfaces on the intensity of a point is related to the color bleeding effect discussed in earlier work [2, 5].

Figure 7 shows a comparison of measured and calculated results for an empty cube with one blue wall and four white walls. The calculated results differ from the measured results by a root mean square difference of less than about 7 percent. On comparing relative values, the results differ by 4 percent. The calculated and measured results are lower for the left side than for the right side of the cube. A method for calculating light intensities that only takes into account the location of the light source, and not surface interreflections, would have given equal reflected intensities on the left and right sides.

A comparison of measurements in an all-white cube with and without an internal all-white box is shown in Figure 8. The large white box was placed in the center of the floor below the light source and was turned at a 45° angle to the walls of the cube. The intensity of the top two horizontal rows is higher for

the case with the internal box because light is reflected from the top of the white box to the upper edges of the cube. The bottom two rows of the open side are less intense than those for an empty cube. These rows see the dark sides of the internal box, which receive relatively little energy since they, in turn, face a black wall and have no direct view of the light source.

Figure 9 shows a comparison of measured and calculated results for an all-white cube with an internal white box. The calculated and measured results have a root mean square difference of about 4 percent and a root mean square difference in relative values of 3 percent. The presence of hidden surfaces did not appear to reduce the accuracy of the calculations. The calculated results follow the same trend as the measured results in having values for the top row that are higher than those for the open white cube, and values for the bottom two rows that are lower. Calculation methods that do not account for diffuse interreflections would not predict the increase in intensity near the top caused by reflection off the box.

In summary, there is good agreement between the radiometric measurements and the predictions of the lighting model. A full summary of results is deferred until Section 4. The perceptual experiments are described in the next section.

## 3. PERCEPTUAL COMPARISONS

Given the experimentally verified output of a light model, the next step in the image synthesis process is to use this information to produce the final simulation. In this section color science methods are used to create a color television image of the simple cubical environment from the output of the radiosity method. This picture is then compared by a group of experimental subjects against a real model as seen through the back of a view camera. This step is taken to evaluate the simulation and thereby to determine whether the overall objective of realistic image synthesis has been achieved.

There is some precedent for performing comparisons between pictures and reality. O. W. Smith constructed an experiment to study depth perception in which a subject viewed a picture and a real scene through a peephole [12]. In computer graphics, comparisons have been made between photographs of reality and photographs of computer-generated images [5, 8], the value of synthetic images for interior illumination design has been studied by an indirect comparison against a real scene [4], and two computer-generated pictures have been compared in order to determine how many polygons are necessary to represent a surface [1].

This section begins with a discussion of the rationale for viewing the model through a view camera while making the comparisons. Next, the experimental apparatus is described, and the procedures that were used to compute the color television picture and compare it against a view of the real model are discussed. Finally, the results of having a group of human observers make the comparison are presented.

### 3.1 Selecting the View of the Real Model

Although the field of view is restricted and some perceptual cues are eliminated, the view camera has been selected for several reasons: (1) it allows simultaneous side-by-side comparisons to be made without introducing the effect of the

observer's memory (a factor that is unavoidable in an alternative viewing scheme such as a pinhole), (2) it corresponds closely to the "synthetic camera" approach employed in computer graphics, (3) it is an experimental setup that can easily be controlled, (4) it is a starting point that must be mastered before other standards can be evaluated, and (5) the degree of restriction is relatively unimportant once the unavoidable step of limiting the view has been taken. In order to present the real and synthetic scences to the observer in the same way, the color television picture was also observed through a view camera.

### 3.2 Apparatus

The simple cubical enclosure that was described in the previous section was used for the perceptual experiments (see Figure 2). In this test case, the model was set up to have a blue wall on the right and a red wall on the left, with the rest of the walls white. The small yellow block was placed on the left and the large white block on the right. The blocks were turned at a slight angle with respect to one another.

The imaging media consisted of a frame buffer and a color television monitor. The frame buffer (Grinnel Systems GMR-27) had a resolution of 480 vertical by 512 horizontal pixels, had eight bits of intensity information in each of its three channels, and produced an interlaced video signal with a frame rate of 30 hertz and a field rate of 60 hertz. The monitor (Barco CTVM 3/51) had a 20-inch display tube with phosphor chromaticity coordinates:

$$x_R = 0.64, \quad x_G = 0.29, \quad x_B = 0.15,$$
$$y_R = 0.33, \quad y_G = 0.60, \quad y_B = 0.06.$$

The individual brightness and contrast controls for each of the monitor guns were adjusted to yield a D6500 white point, and the individual gamma correction functions were measured for each of the guns. The luminance ratios necessary to set the white point were found to be

$$Y_R : Y_G : Y_B = 0.3142 : 1.0 : 0.1009.$$

By determining the proportional relationship between luminance and radiance for each of the guns, these luminance ratios were converted to radiance ratios and were used to balance the guns over their entire dynamic range. The luminance of the white point was set to 24 foot lamberts.

Two Calumet 4 × 5 view cameras were used to view the model and the monitor. The two lenses used were a Schneider–Kreuznach Symmar f1 : 5.6/150mm and a Schneider–Kreuznach Symmar-S f5.6/150mm. Fresnel lenses ruled with 110 lines to the inch and with 10-inch focal length were placed in front of the ground glass of each camera to act as image intensifiers. The combination of the Fresnel lenses and the ground glass introduced some image degradation that made construction and imaging artifacts in both images less obvious.

The positions of the view cameras, the model, and the monitor are shown in Figure 10. The cameras were positioned so that the images were identical in size ($3\frac{5}{8}$ inches by $3\frac{5}{8}$ inches), and the f-stop settings of each camera were adjusted so that the intensities were the same. The combination of f-stop setting and camera-to-model distance were such that the entire depth of the model was in focus,

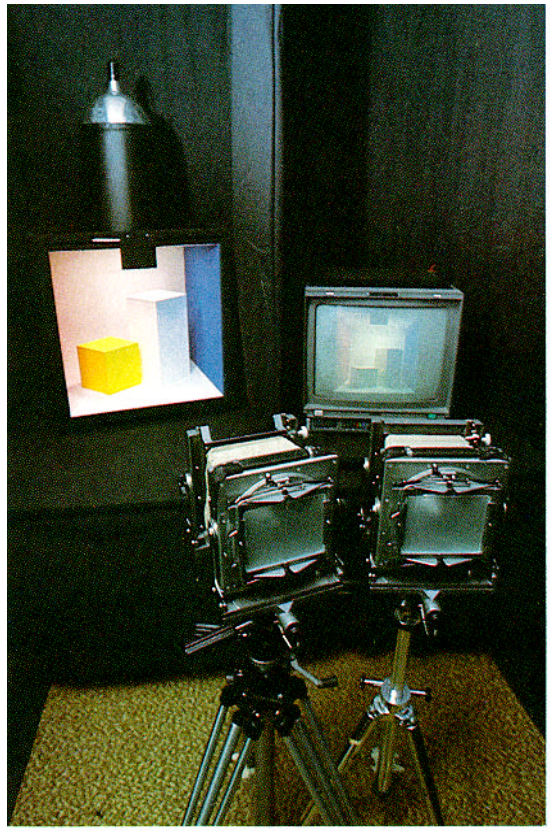

Fig. 10. Experimental setup with the partitioning curtains removed.

thereby minimizing depth of field problems. To minimize reflections, all of the walls were draped in black, and a black curtain (only part of which can be seen in Figure 10) split the room lenghtwise and separated the two view cameras. Another black curtain was hung across the width of the room to separate the subject from the model and the monitor, and the view cameras protruded through holes cut in this curtain.

Figure 11 shows the experimental setup with the widthwise curtain in place and an experimental subject evaluating the view-camera images. The centers of the view-camera backs were $8\frac{1}{2}$ inches apart and were 44 inches off the ground. The subjects were positioned so that their eyes were 25 inches from the view cameras and 48 inches off the ground. The scene, as viewed by the subjects, was inverted, and observations were made under dark ambient conditions.

### 3.3 Procedure

The image was computed using the radiosity software described above [2]. The frustum angle and eye-point position were selected to properly simulate the 150-millimeter lens on the 4 X 5 view camera. Radiosity computations were performed in 15 evenly spaced wavelength bands between 400 and 700 nanometers, and the resulting spectral energy distributions were converted to CIE *XYZ*

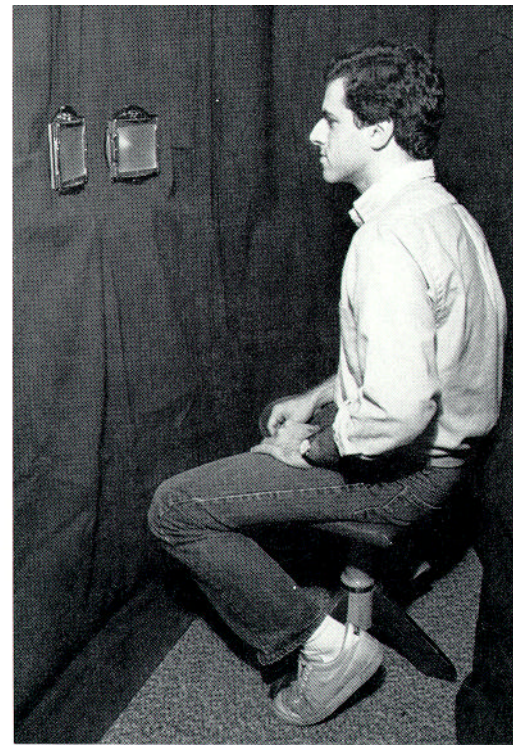

Fig. 11. A subject comparing the real and the simulated images.

tristimulus values. The RGB triplets were found by applying a matrix based on the chromaticity coordinates of the monitor phosphors and the monitor white point [3, 7]. These RGB triplets were subsequently gamma corrected and loaded into the frame store.

Preliminary observations indicated that the limited dynamic range of the monitor would not allow the light source in the ceiling to be rendered convincingly. To avoid this problem, it was decided to alter the experimental design by adding a 4.5 x 2.75-inch opaque flap at the top of the open side in order to obscure the light source when viewing the cube. This minimized the range of light intensities in the scene.

Figure 12 is a black and white picture taken from the position of the observer. It gives an approximate idea of what was seen. Further documentation was obtained by exposing color negative film in the view cameras and producing the color prints shown in Figure 13. No attempt was made to compensate for distortions caused by the photographic process or for the fact that reflection prints seen under bright ambient conditions present an entirely different mode of viewing than self-luminous images seen in a dark ambient environment. Thus these photographs should not be used to evaluate the responses given by the subjects during the comparison experiment.

The subjects for this test consisted of 10 members of the Cornell University Program of Computer Graphics Laboratory who had extensive experience evaluating computer graphics images, and 10 people with little or no experience with

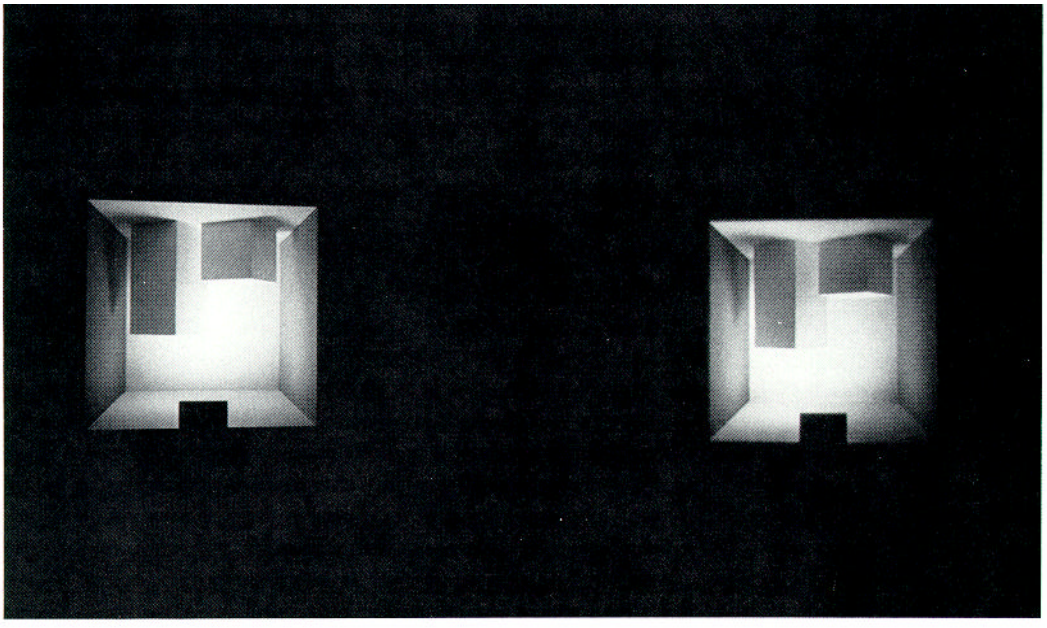

Fig. 12. Photograph taken from the position of the observer that gives an approximate idea of what was seen. The real scene is on the left and the simulation is on the right.

computer-generated pictures. In order to factor out the possible effect of differences between the two lenses, five of each group did the experiment with the lenses on particular sides and five of each group did the experiment with the lenses switched. Because no color vision test was available, the subjects were all taken at their word regarding the normalcy of their color vision.

### 3.4 Results

In trying to decide which was the picture of the model and which was the computer-generated picture, 9 out of 20 people, or 45 percent, selected the wrong answer. The subjects did no better than they would have by guessing.

In all cases, the subjects considered the match between the model and the simulation to be quite good. Specifically, the overall match was rated as being between good and excellent, the color match was rated as being slightly better than good, and the shadow correspondence was rated as being slightly less than good.

Two differences between the pictures were pointed out quite frequently in the written comments. The shadows were described as being fuzzy in the computer-generated image but distinct in the image of the model. This may be due to not discretizing finely enough the surfaces in the environment. It was also noted that the ceiling corners were brighter in the computer-generated image than in the image of the model. Given the results of the radiometric study, where it was discovered that the actual light emits more radiation downward than it does to the sides, this comment is not surprising, since the image was computed with the assumption that the light source emits evenly in all directions.

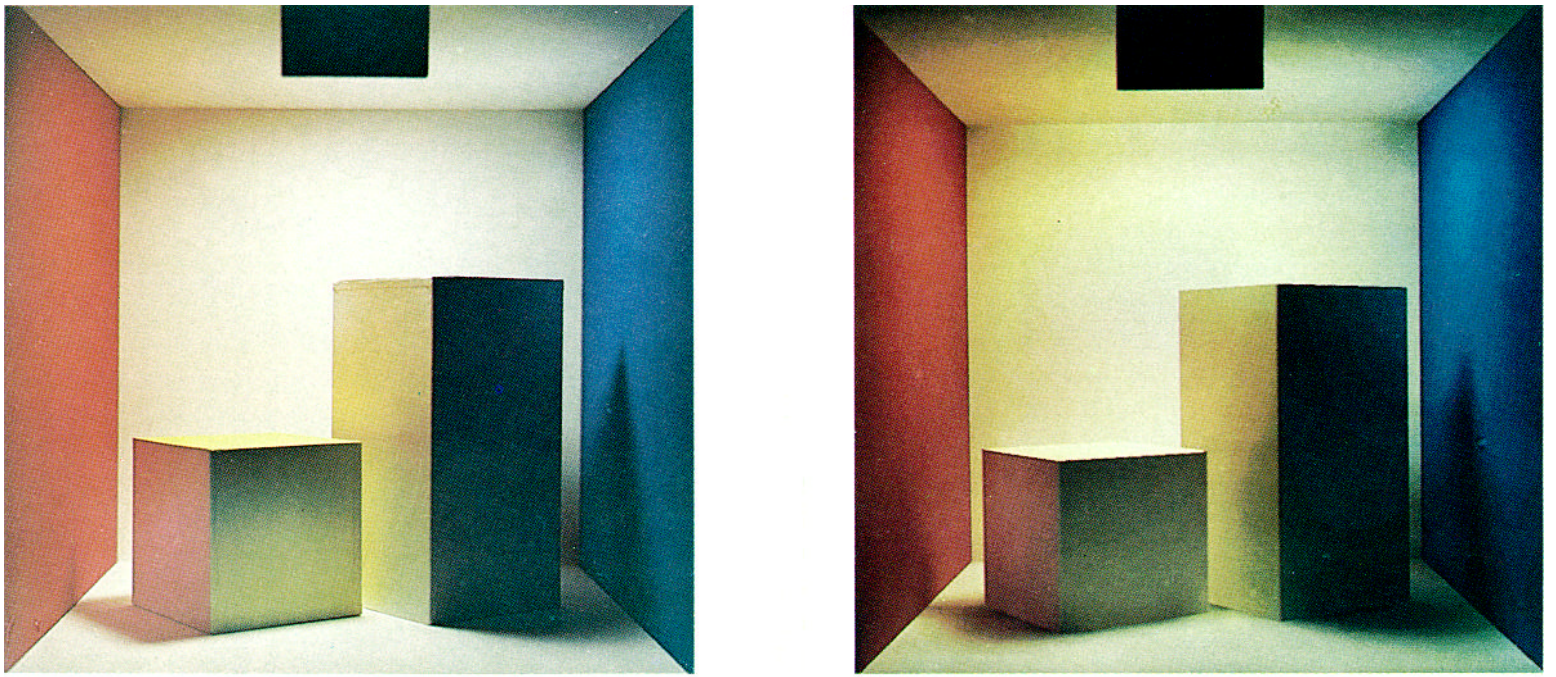

Fig. 13. Photographs taken by exposing color negative film in the view cameras. The real scene is on the left and the simulation is on the right.

## 4. SUMMARY AND CONCLUSIONS

Two experimental studies were carried out to assess the physical and perceptual aspects of the image synthesis process. A physical model was built using diffusely reflecting materials. The physical model was compared in two different ways to the predictions of a diffuse lighting model (the radiosity method). The first study allowed a test of the physical aspects of the lighting model (energy transfer rates were compared). The second study allowed a perceptual test of a rendered image against the physical model.

In the first study, radiometric measurements were made on the physical model and were compared with the predictions of the radiosity method. A simple, hemispherically and spectrally integrating radiometer was used. Three different environments were considered. Some general guidelines emerged for creating a lighting model that accurately describes light (and energy) transport processes in the physical scene. First, the spectral reflectance of materials in the scene must be measured and used as input to the lighting model. Similarly, the spectral and directional characteristics of the light source must be measured and used as input. It is especially important that any directionality in the light source be accounted for. If the light source is not ideal diffuse, an extension to the radiosity method as described in the paper can be used. In the experimental measurements it is necessary to account for the spectral and directional characteristics of the radiometer. With the foregoing factors accounted for, and with care in conducting the experiments, the radiometric measurements and lighting model predictions were found to be in good agreement (see Figures 5, 7, and 9). This agreement lends strong support for the radiosity method as an accurate simulation of the light transfer processes that occur in diffuse environments.

In the second study, the physical model was compared with an image on a color television monitor. The image was synthesized by applying the radiosity method on a spectral basis, selecting the viewing direction, and then converting the predicted spectral energy distributions to $XYZ$ tristimulus values and rendering the image. A single scene was considered (see Figure 10). A perceptual experiment was carried out by asking a group of experimental subjects to compare the simulated image against the physical model. A restricted mode of viewing was employed by asking the subjects to observe the scenes through two view cameras (see Figures 11–13). In comparing the physical scene against the monitor, the subjects did no better than they would have by simple guessing. Although they considered the overall match and the color match to be good, some weaknesses were cited in the sharpness of the shadows (a consequence of the discretization in the simulation) and in the brightness of the ceiling panel (a consequence of the directional characteristics of the light source). The overall agreement lends strong support to the perceptual validity of the simulation and display process.

The present experiments provide a first step in assessing the physical and perceptual aspects of the image synthesis process. Future work should be directed to refining these comparisons. Possible steps include radiometric measurements with high directional and spectral resolution, perceptual tests with alternative modes of viewing, and extensions to more complex environments and to other lighting models.

ACKNOWLEDGMENTS

We express our appreciation to the people from both inside and outside the Computer Graphics Laboratory who donated their time to be experimental subjects. We also acknowledge the assistance of several departments and individuals at Cornell: Physics (David Weidman), Design and Environmental Analysis (Yarrow Namaste), and Mechanical and Aerospace Engineering. Photographic assistance was provided by Emil Ghinger and Rebecca Slivka. The anonymous referees contributed many constructive and clarifying suggestions for which we are grateful. All calculations were performed on VAX computers provided by a generous grant from the Digital Equipment Corporation.

REFERENCES

1. ATHERTON, P., AND CAPOREAL, L. A subjective judgement study of polygon based curved surface imagery. *CHI'85 Conference on Human Factors in Computing Systems* (San Francisco, Calif., Apr. 14–18). ACM/SIGCHI, New York, 1985.
2. COHEN, M., AND GREENBERG, D. P. The hemi-cube: A radiosity solution for complex environments. *Comput. Graph. 19*, 3 (July 1985), 31–40.
3. COWAN, W. B. An inexpensive scheme for calibration of a colour monitor in terms of CIE Standard coordinates. *Comput. Graph. 17*, 3 (July 1983), 315–321.
4. DAVIS, R. G., AND BERNECKER, C. A. An evaluation of computer graphic images of the lighted environment. *J. Illum. Eng. Soc. 14*, 1 (Oct. 1984), 493–514.
5. GORAL, C., TORRANCE, K. E., GREENBERG, D. P., AND BATTAILE, B. Modeling the interaction of light between diffuse surfaces. *Comput. Graph. 18*, 3 (July 1984), 213–222.
6. IMHOFF, E. A. Raman scattering and luminescence in polyacetylene during the *cis–trans* isomerization. Ph.D. dissertation, Physics Dept., Cornell Univ., Ithaca, N.Y., May 1983, pp. 92–95 and Appendix A.
7. MEYER, G. W. Colorimetry and computer graphics. Rep. 83-1, Program of Computer Graphics, Cornell Univ., Ithaca, N.Y., Apr. 1983.
8. MILLER, N. J., NGAI, P. Y., AND MILLER, D. D. The application of computer graphics in lighting design. *J. Illum. Eng. Soc. 14*, 1 (Oct. 1984), 6–26.
9. SCHORNHORST, J. R., AND VISKANTA, R. An experimental examination of the validity of the commonly used methods of radiant heat transfer analysis. *J. Heat Transfer 90* (Nov. 1968), 429–436.
10. SHIH, S. H., LOVE, T. J., AND FRANCIS, J. E. Direct measurement of the radiosity of a nonisothermal hemispherical cavity. In *AIAA Progress in Astronautics and Aeronautics: Heat Transfer, Thermal Control, and Heat Pipes*, vol. 70, W. B. Olstad, Ed. American Institute of Aeronautics and Astronautics, New York, 1980.
11. SIEGEL, R., AND HOWELL, J. R. *Thermal Radiation Heat Transfer*. Hemisphere Publishing, Washington D.C., 1981.
12. SMITH, O. W., AND GRUBER, H. Perception of depth in photographs. *Percept. Motor Skills 8*, (1958), 307–313.

Received May 1985; accepted April 1986