Tools for Lighting Design and Analysis

Gregory J. Ward
Lawrence Berkeley National Laboratory
Berkeley, CA

ABSTRACT

Herein we describe some of the more useful tools and methods for applying computer technology to lighting design problems. We discuss input tools for 3D geometry, materials and photometry, simulation and rendering tools and methods, and output devices and formats. Numerous examples are given.

DISCLAIMER

A number of specific products and manufacturers will be mentioned in the text of these course notes. They are merely meant to serve as examples, and our references to them should not be construed as endorsements of their suitability or merit for any particular application. The only exceptions to this rule are products associated with the organizers of this course, which we will plug shamelessly.

Introduction

In this course, we have glimpsed some of the possibilities and applications of global illumination in architecture and entertainment. However, it is clear that there are still many details to be worked out. For instance, how do we generate the input necessary for a good global illumination calculation? What is the most appropriate calculation tool for our application? How can we best represent and present our results? Though these issues may seem peripheral, they will eventually determine the success or failure of our efforts.

In this section of the course, we will discuss some of the hardware and software tools available for generating 3D models with physical properties, performing appropriate lighting and rendering computations, and presenting the results. We will start with input tools and methods, then look at simulation and rendering, followed by output and communication. As part of the discussion of output, we will give some examples from lighting design work conducted at LBNL and elsewhere.

Input

There are three interrelated classes of data that are important to global illumination: geometry, materials and photometry. Most of us are familiar with geometry and its input, but few computer graphics practitioners have experience with physical materials, and even fewer understand photometry. This is because, even today, most rendering programs ignore physical lighting properties in favor of simplistic or artistic views of the world. Not so in the case of global illumination -- here we need to know exactly how light interacts with surfaces in order to predict their true appearance. The material models describe how light reflects off and transmits through surfaces, and the photometry describes how light emanates from light sources.

Tools for Geometry

Many of the tools and methods available for describing geometry will be familiar to those with experience in generating 3D computer graphics imagery. Nevertheless, we will go over some of the more useful tools here for the sake of completeness.

CAD Systems

Computer-Aided Design systems are the most common means of generating 3D input for global illumination. Sometimes, a CAD system is included in the rendering system itself, which ensures that the data is compatible and complete. However, writing a good geometric editor is very difficult, so the more common approach is to import 3D geometry in some standard format, such as IGES, DXF or VRML, and rely on a commercial CAD program to create the geometric description.

Among the commercial CAD programs, AutoDesk's *AutoCAD* is probably the most popular worldwide, because it has been around a long time and many companies have successfully marketed value-added software functionality and libraries for it. Unfortunately, because it was originally designed strictly as a 2D drafting tool, its 3D editing abilities are somewhat awkward and difficult to use. Many other CAD systems are available, some of them as shareware and freeware and some are much better than *AutoCAD*, but the choice of CAD system is usually made in advance by personal preferences or company requirements, so we will not discuss this further here.

The main issue when considering which CAD system to use it is whether or not it can produce a usable model for global illumination. Unfortunately, the only way to know for sure is to try it. In some cases, compatibility depends on *how* the system is used -- i.e., sticking to exportable primitives and labeling surfaces in such a way that the appropriate materials may later be linked to them. Other issues related to exportability include surface normal orientation, planarity, meshing and vertex accuracy.

3D Scanners and Digitizers

As fun as CAD programs are, we often wish for an easier method to input a model of an object we have on hand. After all, if the 3D geometry of a telephone or a stapler is right in front of us, why should we have to recreate it all over again in a CAD program?

Two types of devices can help us out, 3D scanners and 3D digitizers.

A 3D scanner is a device that automatically scans in an object to create a 3D model of that object. This is very difficult to do in general, because objects may have concavities and reflectance properties that cannot be captured or adversely affect the capturing process. There are only a few devices currently on the market for capturing full 3D geometry. The best known products come from Cyberware (http://cyberware.com), which use an off-axis laser scanning method. A London company manufacturing scanners is 3D Scanners Ltd. (http://www.3dscanners.com). In general, these devices are bulky and expensive.

3D digitizers are much cheaper and simpler in concept, but generating detailed models with them is time-consuming. They typically have an arm with a stylus attached, which is used to locate vertex points in 3-space. Some devices are wireless and cover a larger volume, but accuracy can be an issue. 3D digitizers are frequently advertised in graphics and CAD magazines.

Graph Paper

Sometimes the simplest solution is the best. A method that has worked for me for many years is one of projecting an object onto a piece of graph paper and tracing its outline. Coordinates can then be located and, provided the object has some symmetry, a nice 3D model can be constructed. An overhead projector may be used to get a nice silhouette, or if the object is reasonably flat, it may even be traced directly. The data points may then be entered into a CAD program or your favorite text editor.

FTP and Web Sites

Of all the methods for creating 3D models, the quickest and easiest will always be downloading them off the net. In some cases, models are offered free of charge by magnanimous individuals or organizations. Reputably the largest free repository was assembled by the Navy's Vislab at China Lake, and dubbed Avalon. This site also contains useful tips and information on converting models from various formats. Avalon is currently maintained by Viewpoint Datalabs, who has their own proprietary data models. Another point of interest is the MGF web site, who also offers complete models with all the necessary materials and photometry for global illumination in a well documented format with an ANSI C parser. Here is an abbreviated list of 3D model sites (* means all data is free of charge):

http://www.viewpoint.com/avalon	Avalon, maintained by Viewpoint*
ftp://wuarchive.wustl.edu/graphics/graphics/mirrors/avalon	US mirror*
ftp://ftp.flashnet.it/avalon	Italian mirror*
ftp://sgi.felk.cvut.cz/pub/avalon	Czech mirror*
http://www.viewpoint.com	Viewpoint main site
http://radsite.lbl.gov/mgf	MGF web site*
http://www.acuris.com	Acuris web site
http://204.191.254.145	QuickDraw 3D Clipart

Tools for Materials

The materials in a scene or object description describe how light interacts with each surface. The important thing for global illumination is that the material models must be physically plausible, that is, they must adhere with some consistency to physical laws governing light transport. Unfortunately, the most commonly used material models in computer graphics have a very weak physical basis, and in most cases they do not yield plausible results if applied in their original form. (See [He91], [Ward92], [Schlick93] and [Lewis93] for good counterexamples.) Therefore, most materials that accompany computer graphics models must be reinterpreted in a physical context, or recreated from scratch.

In the case where the actual object or scene is available, one may wish to measure the materials directly. The methods and tools used for such measurement depend on the type of material being measured and the desired application and accuracy. Here are some examples.

Cameras and Scanners

When a material has a prominent texture or pattern and we want to capture this for a more realistic rendering, we can scan this pattern into our computer. If we have a sample of the material handy, we may scan it in directly by placing it on a flatbed scanner. This usually yields the highest quality results. If the object is in the field and unmovable, we might rather take a picture of it and scan in the picture using a 35mm scanner, or go directly to digital with a digital camera. Alternatively, we might capture frames from a video, though this tends to compromise resolution and quality quite a bit.

However we get the data to our computer, we will have difficulty getting accurate color and reflectance information unless we are very careful. The simplest method for maintaining accurate color and contrast in this kind of process is to transfer a reference chart in parallel, such as the Macbeth ColorChecker Chart[1]. When lighting is uniform and controlled, as in the case of a flatbed scanner, absolute color and reflectance accuracy can be quite good. Figure 1 shows the result of a calibration process using the Macbeth chart.

[1] Available in most photo supply stores, or contact Munsell Color, 405 Little Britain Rd., New Windsor, NY 12553, 1-800-MACBETH (USA) or 1-800-COLOUR (Canada).

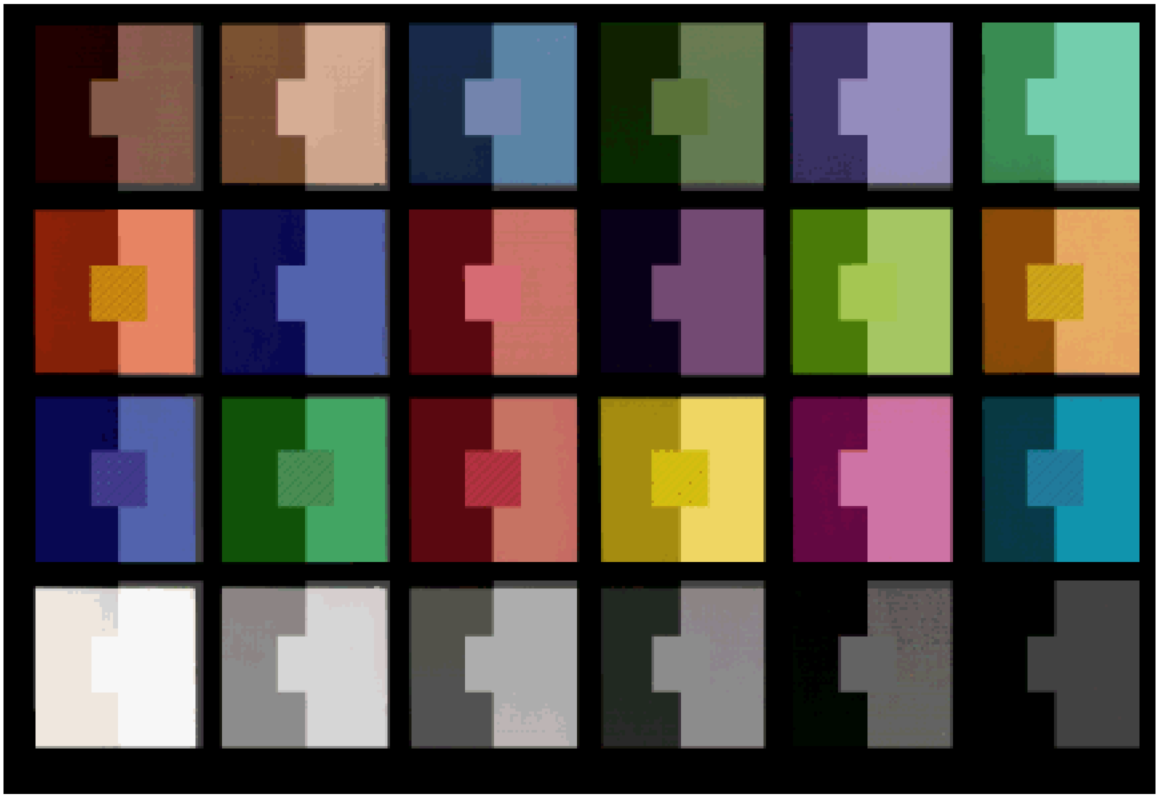

Figure 1. Automatic contrast and color calibration using a Macbeth ColorChecker. The left side of each square shows original scanner results. The midsection of each square shows target Macbeth color, according to published CIE chromaticity and reflectance values. The right side of each square shows results achieved through color/contrast correction step. Notice that grays are matched very accurately, as are most colors, but some highly saturated colors are off, which may indicate that these colors are outside of the device's gamut.

Spectrophotometers

If a surface has fairly uniform color, more accurate measurements of its reflectance are possible. A *spectrophotometer* (or *spectrometer* for short) is a device that measures the reflectance or transmittance (or emittance) of a surface at multiple wavelengths. Some devices, such as the Colortron by Lightsource (http://www.lightsource.com) use a diffraction grating to separate visible wavelengths. Figure 2 shows an example of a reddish surface measured in this way. Other devices use a set of highly selective color filters, such as those in the Minolta CM-2002. The CM-2002 has the added advantage of being able to measure reflectance with and without the specular (mirror) component. Figure 3 shows a diagram of the internal workings of this device.

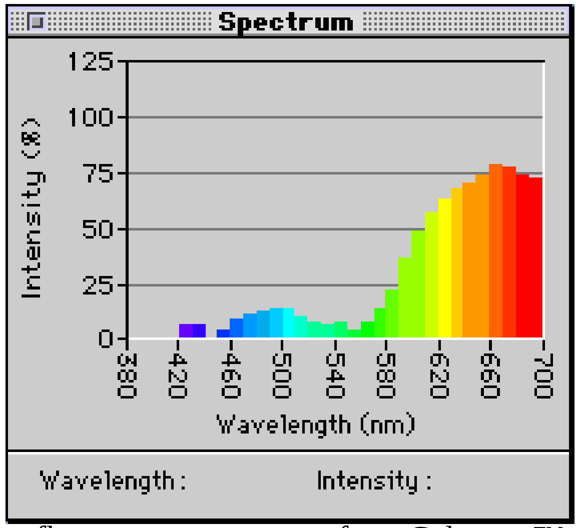

Figure 2. Spectral reflectance measurement from Colortron™ spectrometer. Each waveband is about 10 nm wide.

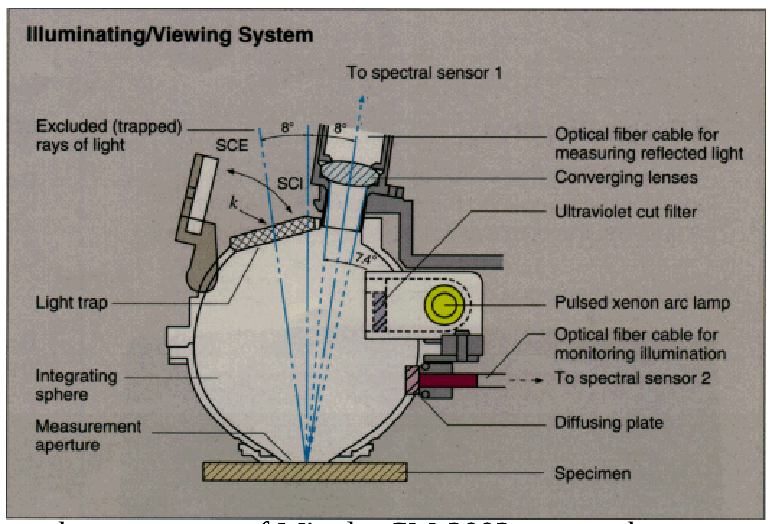

Figure 3. Internal arrangement of Minolta CM-2002 spectrophotometer. Note the trap door for excluding the specular component.

Gloss Meters

Cameras and spectrometers work fine for diffuse materials (or in the case of the Minolta CM-2002 described above, diffuse materials with some ideal specular component), but what about surfaces with some rough specular reflection? How do we measure and characterize material highlights? For many years, auto manufacturers and other industries concerned with the appearance of shiny paints have employed devices called *gloss meters* to do just this. A gloss meter takes a relative measurement of the intensity of a highlight at one or a few different angles. It is a purely relative measurement, and mixes highlight spread with intensity in a way that is difficult to translate to useful material properties. Therefore, we will not say anything more about these measurements here, as we do not recommend them for computer graphics or global illumination.

Gonioreflectometers

A *gonioreflectometer* is a device that measures reflection as a function of incident and reflected angle in order to arrive at the bidirectional reflectance (or transmittance) distribution function for a surface (BRDF or BTDF). Devices that measure this function as it depends on wavelength also are called *goniospectroreflectometers*. (Sometimes people run out of breath and just call them all *goniometers* .) These devices tend to be large, expensive and non-portable. Figure 4 shows an imaging gonioreflectometer developed by the author, which is further described in [Ward92].

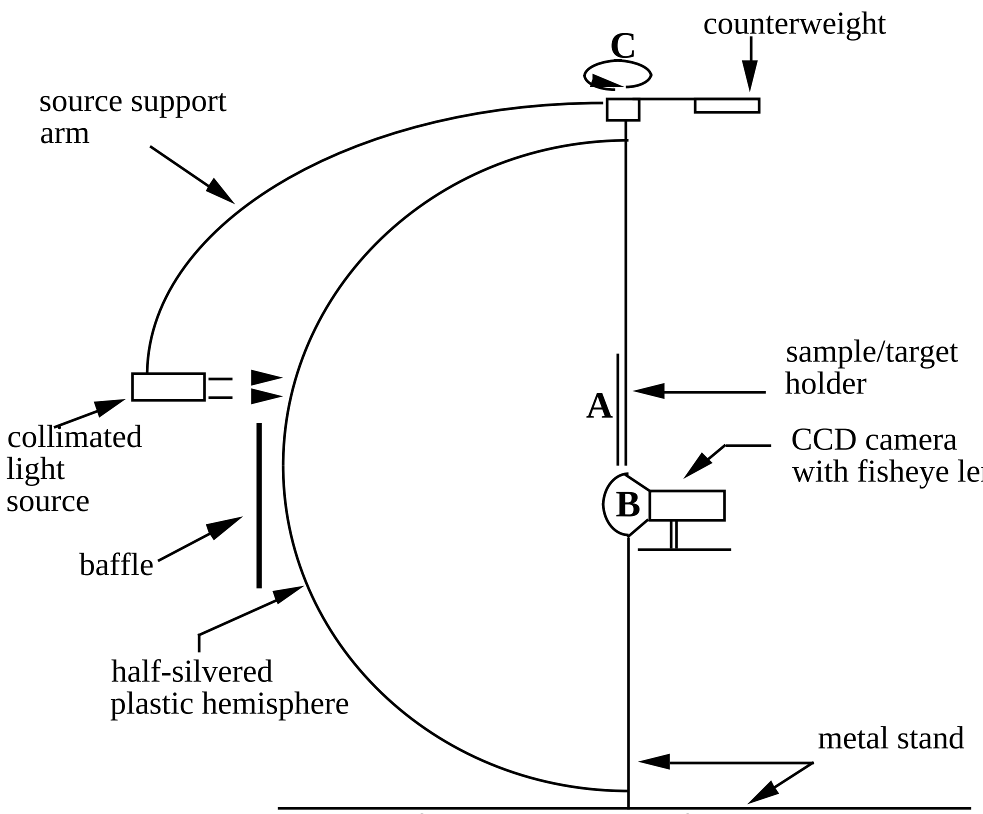

Figure 4. LBNL imaging gonioreflectometer, with half-silvered hemisphere reflecting light into fisheye lens of digital camera [Ward92].

We hope to see simpler and more practical methods developed in the near future. We especially need methods for characterizing BRDFs that are cheap and portable, like the method recently proposed by Karner et al, which uses a digital camera and a fixed source and reference material [Karner96]. Figure 5 shows this sample arrangement.

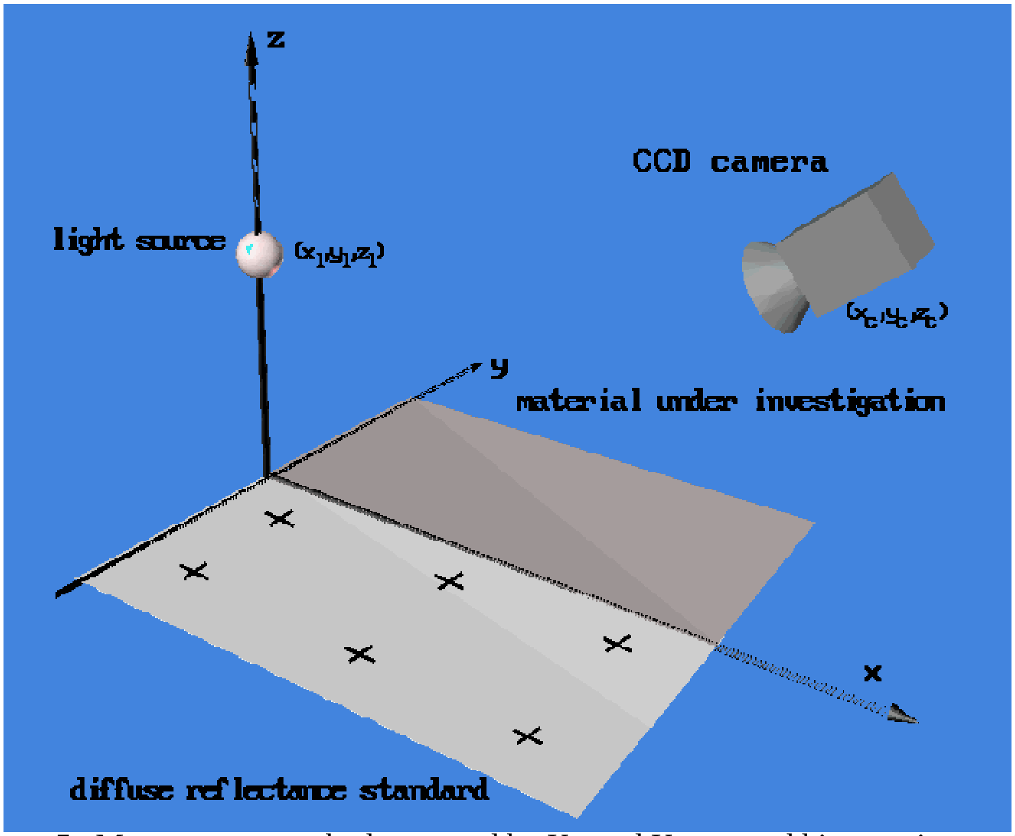

Figure 5. Measurement method proposed by Konrad Karner and his associates at TU Graz in Austria for obtaining BRDF using simple camera, source and reflectance standard arrangement [Karner96].

Tools for Photometry

Photometry is the measurement of light output from sources in an environment. In most scenes, this might include daylight or electric light. Daylight is generally described by the sun and a model sky distribution. The sun varies only in overall intensity, as determined by its altitude and the atmospheric conditions. Under cloud cover, the sun may not be visible at all. The sky model can be very simple (such as a uniform sky) or very sophisticated (such as a measured sky). Electric lighting usually means luminaires in the environment, each having a photometric output distribution. The photometry of a luminaire is generally available from its manufacturer, which relieves us of the nasty business of goniophotometric measurements.

Daylight Models

The International Commission on Illumination (CIE) has adopted three reference models for sky luminance distribution: the CIE overcast sky, the CIE clear sky and the CIE intermediate sky (http://chelsea.ios.com/~kdtjek/cieusa.html, pub #110). These models offer reasonable approximations to a completely overcast day, an ideal clear sky, and a sky that has some uniform cloud cover but still allows direct sun to seep through. Some global illumination programs may also provide a choice for a uniform sky distribution, which is merely a constant value over the entire hemisphere. This distribution is not very realistic, but may serve as a point of reference in some calculations. The CIE overcast sky model is most often used in *daylight factor* calculations (the ratio of interior to exterior horizontal illuminance), since the distribution does not vary with sun angle.

It is very important to understand the limitations of reference sky models. Since they serve only as ideals, they are not very good at predicting actual interior (or even exterior) illumination values. For this, we really need measured sky data. At the very least, we need measurements of the global and diffuse sky components, which allow us to set approximately correct values for the sky and solar intensities. However, to get truly accurate results for a specific time of day and year, we need measurements of the sky distribution and solar component for that moment. Devices exist for taking such measurements, but like the gonioreflectometers discussed in the previous section, then tend to be expensive, bulky, and difficult to use.

Luminaire Photometry

If sky measurement is difficult, measuring the output of luminaires is even worse. Fortunately, most luminaire manufacturers do this for us, often sending the work out to independent laboratories to retain some semblance of objectivity. The data is then made available to us as photometric reports or standard luminaire data files. The latter is obviously preferable if our intent is to apply some calculation tool such as a global illumination or rendering program. All we need to check is that our program takes the provided format as input.

In the U.S., there is really only one generally agreed upon standard for luminaire data, put forth by the Illuminating Engineering Society of North America (IESNA or IES for short). It is available as IES publication LM-63-91, *Standard File Format for Electronic Transfer of Photometric Data*. (Replace the last two digits of the code with the year a new one becomes available.) Connect to the IESNA web site (http://www.iesna.org) for more information. A free parser for this format, written by Ian Ashdown, is available from the Ledalite website (http://www.ledalite.com).

In Europe, the CIE has put forth their own standard for luminaire data, available as pub #102. (See the CIE web site at http://chelsea.ios.com/~kdtjek/cieusa.html. Also, Andrew Glassner's book, *Principles of Digital Image Synthesis*, offers excellent summaries of this and the IESNA standard [Glassner95].) This is a relatively new standard, and there are still many other country-specific standards in use in Europe, such as CIBSE and LTLI. Hopefully, one day the world will agree on a single standard.

Combining Input

It would be nice if there were a common format for storing all the essential input information for global illumination. Currently, we must import everything into the simulation or rendering program we wish to use, with little hope of sharing this information between different calculations once it is combined.

Recently, several new formats have emerged for representing 3D geometry and materials, including VRML (a network extended version of SGI's Inventor format), QuickDraw 3D (from Apple) and 3D Studio. Unfortunately, none of these formats employ physically plausible material models, and none of them allow for light source photometry, so they do not really provide complete descriptions for global illumination.

There is currently one finished standard that makes a reasonable attempt to provide complete descriptions for global illumination in a neutral format, called the Materials and Geometry Format (MGF). (See http://radsite.lbl.gov/mgf.) MGF shall soon be integrated into the IESNA standard for luminaire data, described in the previous section. Two other efforts are underway to provide more ambitious and complete descriptions of objects for all types of simulation, the Industry Alliance for Interoperability (IAI) in the U.S. and STEP in Europe. These last two efforts will likely merge to one day provide a single, global standard for describing everything computers want to know about manufactured objects. (See http://elib.cme.nist.gov/pub/nipde regarding STEP or http://www.interoperability.com regarding IAI.)

Lighting Simulation and Rendering

The type of global illumination calculation we perform depends on the kind of input we have, the degree of accuracy we require, and the kind of output we want. Usually, it works best to think this through backwards, that is, decide *first* what kind of output we need for our application, then decide what simulation tool works best for this, then go about gathering the necessary input.

Lighting simulation and rendering can be thought of as a sort of continuum, starting with the most basic calculations and moving on towards the ultimate in predictive reality. Figure 6 shows this continuum with three example software packages. The horizontal axis shows the relative effort required on the input side, and the vertical axis shows the quality and detail produced on the output side.

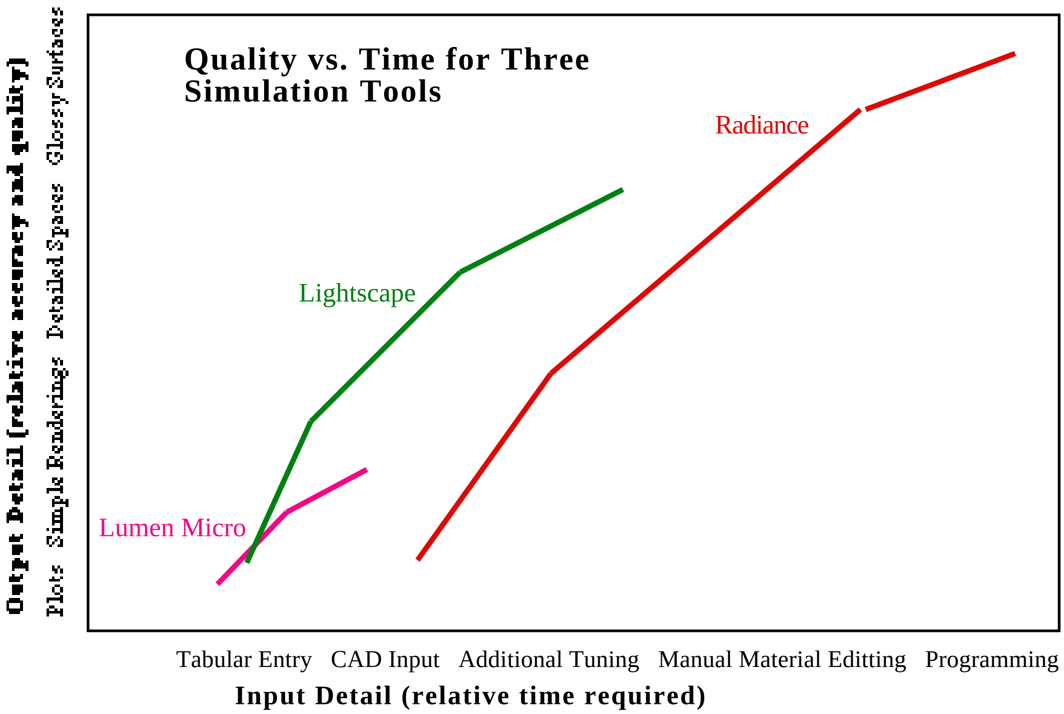

Figure 6. Three curves showing output quality versus input time for three lighting simulation tools.

Figure 6 shows that the output desired affects the level of effort required on the input side as well as the choice of simulation software. *Lumen Micro* is good for simple calculations and spaces, *Lightscape* works well for intermediate level of detail, and *Radiance* for high detail spaces [Ward94b]. *Lightscape* offers a nice user interface, which is why its curve falls slightly above the other two -- it takes less time to get the same quality of output. The fact that these are curves instead of straight lines reflects the law of diminishing returns as the user expends additional effort on the input side.

For more information on *Lumen Micro*, contact Lighting Technologies in Boulder, Colorado (303) 449-5791. *Lightscape* has a web page at http://www.lightscape.com. The official web site for *Radiance* is http://radsite.lbl.gov/radiance.

Output and Communication

From what we just noted in the previous section, we are treating last the topic that should be considered first, namely what sort of output is needed for a particular application. For some applications, such as routine lighting design, simple line graphs or contour plots may be sufficient. In some cases, the designer or client may wish to see some crude renderings of the space without furniture or other details. More critical applications may require daylighting analysis or the inclusion of partitions and some furniture, though a diffuse approximation to interreflection may still be adequate. For critical visual applications such as interior design or control operations environments, full modeling of detailed furniture and reflection is often necessary.

In addition to considering the types of calculations required, it is important to think about what output media best convey the information to the intended audience. If the audience is an experienced lighting designer, very crude renderings or even point values may be sufficient. A typical design client (building owner or buyer) usually prefers color renderings. If the design is part of a competition or other demonstration, it may be worthwhile to produce an animation.

The time required to produce a particular result is generally a product two axes: input detail and output resolution. By "output resolution, " we mean the number of individual calculations needed. Point calculations tend to be sparse, so this is the cheapest output to produce. High resolution video is the most expensive output, since every pixel of every frame must be computed somehow. For simple environments, even high-resolution video can be produced in a reasonable time frame, but few people are interested in walk-throughs of empty spaces. Likewise, point calculations can be done fairly quickly even in a complicated environment, but this does not make good use of the detail in the model.

Let us summarize and demonstrate some of these output possibilities.

Numerical Lighting Values

Numerical lighting values are often used to decide if light levels are adequate in a space and to check that there are no problematic sources of visual discomfort or glare. Figure 7 shows a daylight factor contour plot generated from a small set of workplane illuminance values computed by the *ADELINE* software package (http://radsite.lbl.gov/adeline). Similar plots may be obtained for illuminance. Figure 8 shows a glare calculation, which locates extremely bright regions in the scene and determines visual discomfort probabilities for horizontal view directions. (Accurate glare calculations may require non-diffuse reflection models, and are among the more difficult numerical analyses used in lighting.)

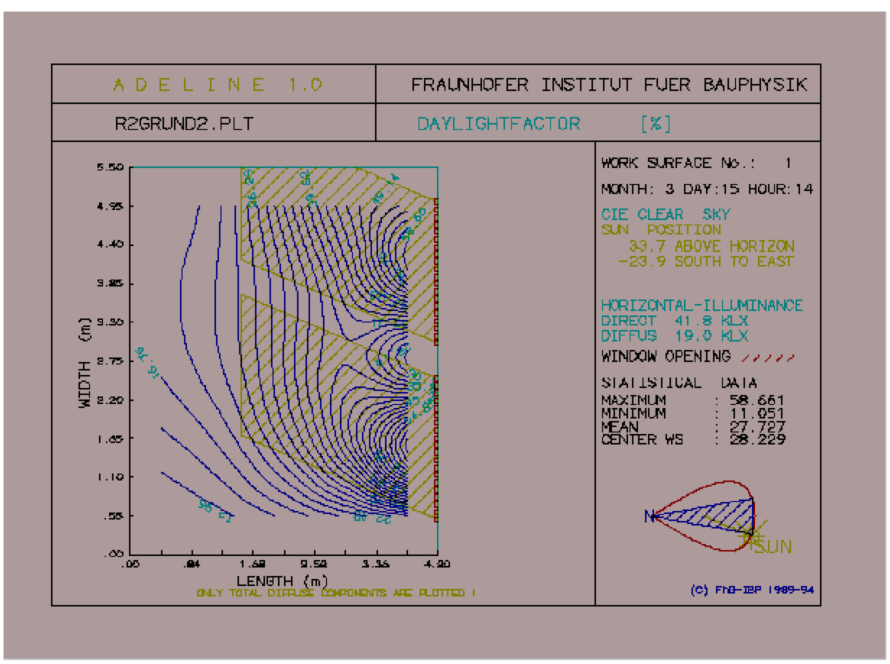

Figure 7. Daylight factor contours for workplane of two-windowed office space (empty), computed by *ADELINE*. Sun patches are shown as yellow, hatched regions.

Figure 8a. Visual comfort probability (VCP, glare) calculation for cubicle office space, showing position and intensity of glare sources, computed with *Radiance*.

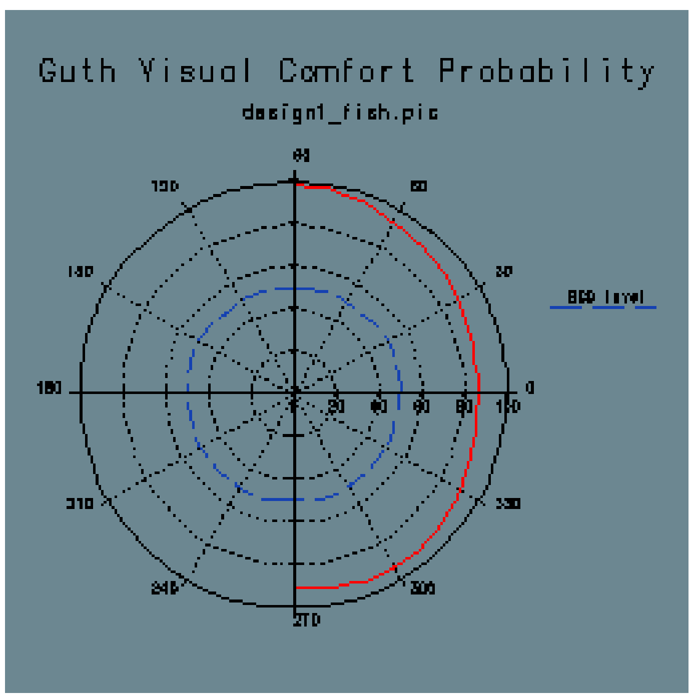

Figure 8b. Visual comfort probability (VCP, glare) calculation for cubicle office space, showing a numerical plot of VCP. The zero degree view is looking straight ahead in the image shown, and positive angles look to the right. A VCP of 75% means 75% of occupants are comfortable.

Renderings

The most natural medium for presenting lighting simulation results is of course the rendered image. Renderings are easy for anyone to interpret, as they attempt to show what a design space will look like. How well they achieve this aim depends on the modeled environment, the software, and the chosen display method.

The chief limitation of most display methods is dynamic range. A typical scene has luminance ratios on the order of 1000:1, with many daylight scenes exceeding 10000:1. Unfortunately, most display monitors have a maximum dynamic range of around 100:1, with film recorders doing slightly better and other hardcopy devices doing slightly worse. There simply are no generally available display methods that approach the dynamic range of a real scene. What this means is, what you see in an image is not always what you get in real life.

A few different researchers have addressed the "tone mapping" problem as it is called. The pioneering work in computer graphics was conducted by Tumblin and Rushmeier [Tumblin93], who developed an exponential function to produce roughly the same apparent brightness as real environments. This author took a slightly different approach [Ward94a], attempting to reproduce visible contrast differences with a linear operator. Ferwerda et al followed similar logic, and went on to include scotopic/photopic response (i.e., color sensitivity) and time adaptation [Ferwerda96].

Unfortunately, none of these approaches really solves the problem entirely, since the experience of a display with limited dynamic range will never duplicate the experience of a real environment. For example, an extremely bright source like an oncoming car on the highway will just appear as two white dots, rather than causing discomfort in the viewer [Spencer95]. We therefore need some other indication of luminance, such as the glare display shown in Figure 8 of the previous section, or a false color representation as shown in Figure 9.

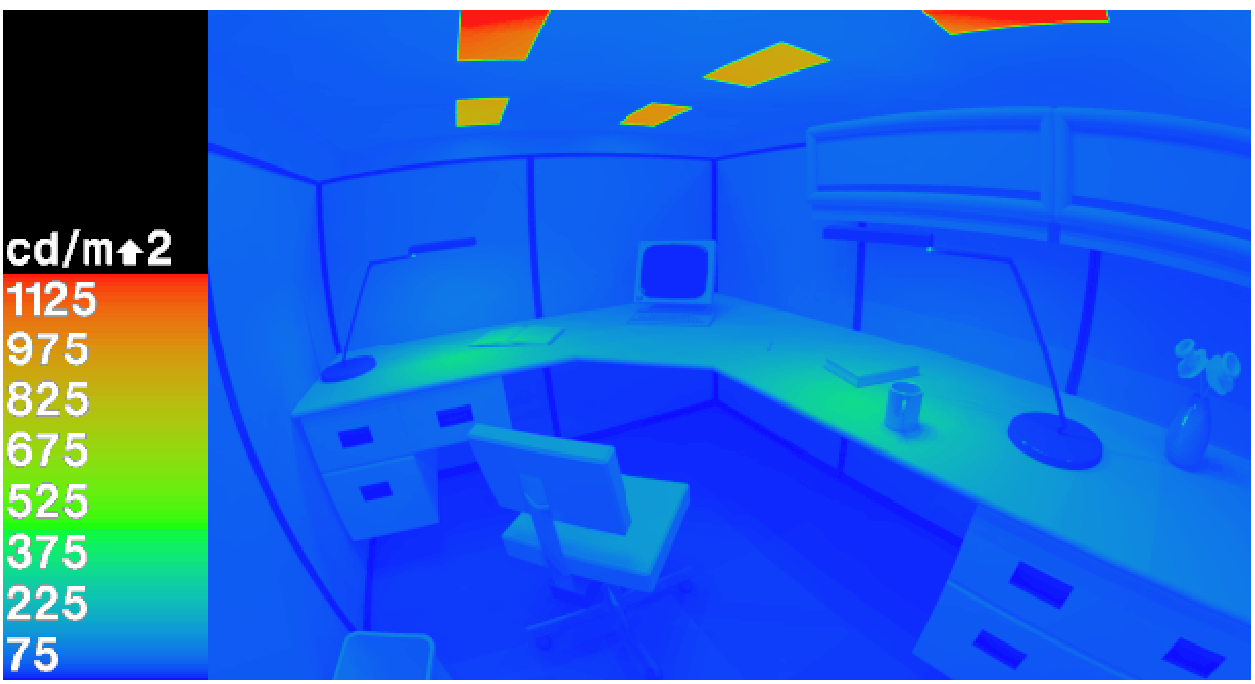

Figure 9. A "false color" image representation may be used to overcome limitations in the dynamic range of displays, replacing tone mapping with a numerical legend.

Another alternative on the horizon is high intensity VR display hardware, tied into the viewer's direction and adjusting the exposure dynamically and continuously. We will discuss this further in the final section on real-time interaction. Next, let's take a look at animation.

Animation

While plots and still images work well enough for most lighting design problems, there are occasions when the client wants to see more than can be shown in a few pictures. At its simplest, animation is merely stringing together many still images to get the illusion of motion. Sometimes, we may wish to move objects in the scene, though for lighting design, it is more likely we will want to simulate different lighting conditions (e.g., different times of day and year or different light settings). Let us discuss some of these possibilities in turn.

Camera Animation

The simplest kind of animation, and the cheapest to compute, is camera animation (also called "walk-through" animation). In this case, the scene and lighting remain static, while the viewpoint and direction move through space. In diffuse environments, such an animation can be computed in real-time for moderately complex spaces (10,000 polygons) on modern hardware. (We will touch on this again in the final section.) For very complex scenes and scenes with non-diffuse surfaces, however, other methods are necessary. Movie 1 shows a lower deck of a U.S. Navy cruiser, modeled with a few hundred thousand surfaces (which would expand to over a million polygons). A walk-through animation of this environment was computed using *Radiance* with a Z-buffer interpolation scheme similar to the one described by Chen and Williams [Chen93].

Movie 1. Walk-through animation of lower deck of U.S. Navy cruiser, computed with *Radiance* using Z-buffer frame interpolation.

A compromise approach that does not require computing multiple frames is Apple's QuickTime VR [Chen95], where a single panoramic view provides the user with a pivot point in the environment. By computing panoramas at several positions in the design space, a certain degree of freedom is provided. We will mention this again in the final section.

Scene Animation

A favorite pastime in computer graphics, scene animation moves objects in the environment around. This is not used as much in lighting design and analysis, since it is expensive and not terribly informative.

Lighting Animation

Lighting animation shows what happens to a view as the lighting is changed over time. This is especially useful in the case of daylighting, which is highly variable throughout the day and year. Figure 10 shows four times of day for a small office with venetian blinds on the window, again computed with *Radiance*. In this study, different blind settings were examined as were algorithms for automatically controlling the blind angle throughout the day, and the results were produced on a QuickTime video. Some researchers have even optimized this type of solution using steerable filters [Nimeroff94].

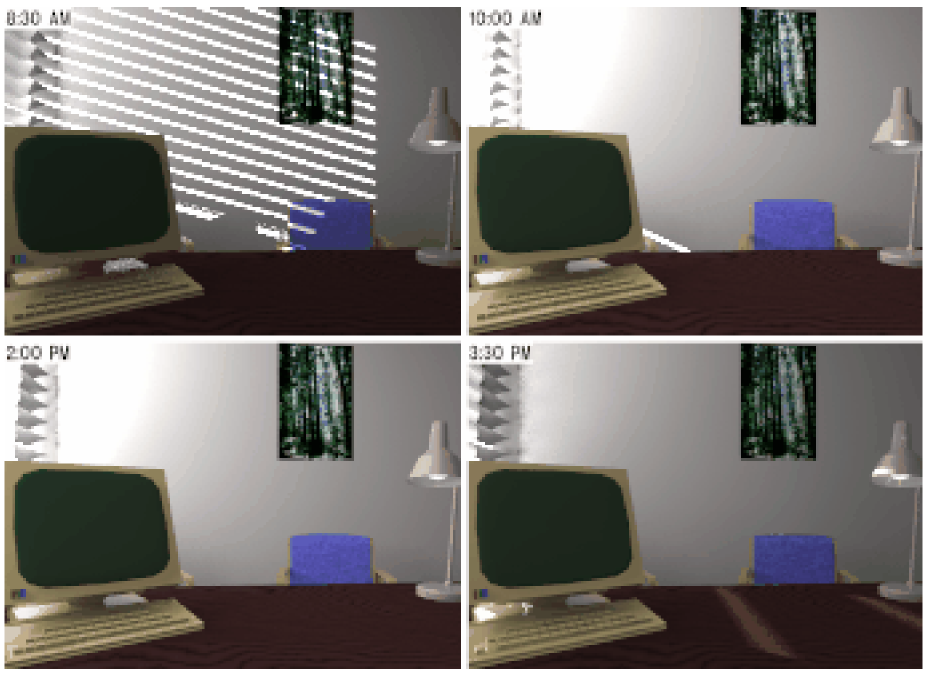

Figure 10. Four frames from a time-lapse animation of a small office with venetian blinds, computed by *Radiance*.

Real-time Interaction

Although animations are nice, they require a lot of preparation and production time, not to mention expensive video conversion equipment or, in the case of compressed video, compromises in quality. Modern Z-buffer graphics hardware is capable of displaying simple 3D spaces interactively in real-time, which is a much nicer way to view a scene. If our global illumination solution assumes diffuse surface reflection, we can use Gouraud shading to get a fairly accurate screen representation, provided the dynamic range is not too great. Figure 11 shows a radiosity rendering of an operating room, computed by *Lightscape*. The model contains a few thousand polygons, and a user may move around in this space freely on a modern graphics workstation. Displaying textures is more difficult, requiring a much more expensive class of machines for real-time interaction. Also, very complex geometry is unmanageable even for the most expensive graphics systems, leading a number of researchers to work on automatic model culling for real-time display [Funkhouser93] [Teller93].

Figure 11. Radiosity solution of operating room, computed by *Lightscape*. Real-time movement through this scene is possible on a moderately-priced graphics workstation, without textures.

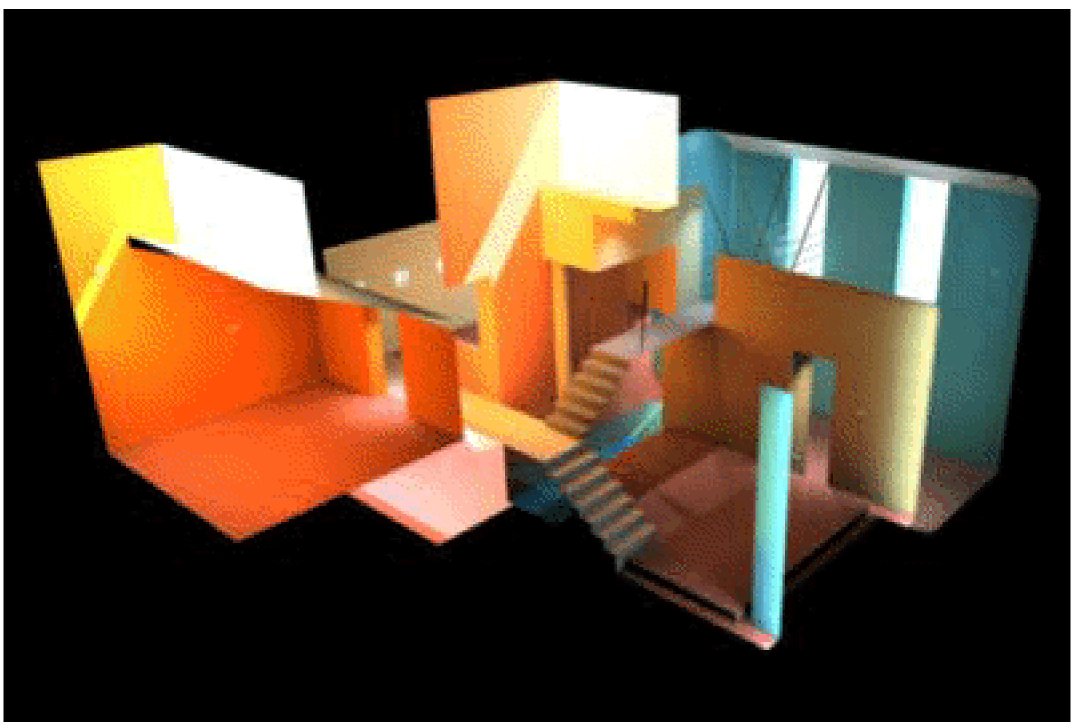

Figure 12. A cutaway view of a small, daylighted house, computed by *Lightscape*. The limited dynamic range of graphics hardware means that bright areas get washed out.

Figure 12 shows a cutaway view (back faces removed) of a small house, illuminated by daylight. This is a more interesting model to walk through, and is still simple enough to interact with in real-time on a graphics workstation. However, the dynamic range of the space cannot be adequately represented, so some discretion must be applied in interpreting what one sees in this space. (This limitation is due primarily to the integer color computations of graphics hardware, not the radiosity solution itself, computed here by *Lightscape*.)

An even cheaper alternative to graphics hardware interaction, and one not limited to diffuse environments, is Apple's QuickTime VR [Chen95]. In this system, 360˚ cylindrical views (panoramas) are generated at specific view points throughout the environment. The user may move from one view point to the next, and at each point may look at the scene in any direction (except straight up or down). Because the view point does not move, specular interactions can be displayed correctly, and there is no limit placed on geometric complexity, though these two factors will of course influence the initial computation time. Also, the same limit in dynamic range applies to this system as well as the other hardware methods.

To bring global illumination into the world of truly visceral virtual reality, we need a wide-field, high-resolution, high dynamic range display system hooked to a real-time calculation of light transport in complex, textured, non-diffuse environments. As you might suspect, we are still some distance from attaining this goal. Even with massively parallel computer architectures, the time required to calculate specular and diffuse interactions in complex environments is too long for real-time update of high-resolution, wide-field displays. However, in the case of a head-mounted display, we know that the foveal region is the only thing that needs high resolution, and we can skimp on the rest. It may be possible, therefore, to optimize our calculation based on the viewer's gaze direction, and combine this with quick redisplay methods such as Z-buffer interpolation to reach real-time interactivity in complex environments.

This is the challenge that lies before us.

Acknowledgments

Saba Rofchaei produced most of the cruiser model shown in Movie 1 and the office model shown in Figure 11. Figure 12 was produced by Dewoolf Partnership, Architects of Rochester, NY. Figure 13 was produced by David Hileman of Toronto, Canada. The model in Figure 14 (color plate) was created by Charles Ehrlich for the Federal Aviation Administration.

References

[Chen93] Chen, Shenchang Eric, Lance Williams, "View Interpolation for Image Synthesis," *Computer Graphics Proceedings, Annual Conference Series*, 1993.

[Chen95] Chen, Shenchang Eric, "QuickTime VR -- An Image-Based Approach to Virtual Environment Navigation," *Computer Graphics Proceedings, Annual Conference Series*, 1995.

[Ferwerda96] Ferwerda, James A., Sumant Pattanaik, Peter Shirley, Donald P. Greenberg, "A Model of Visual Adaptation for Realistic Image Synthesis," *Computer Graphics Proceedings, Annual Conference Series*, 1996.

[Funkhouser93] Funkhouser, Thomas A., Carlo H. Sequin, "Adaptive Display Algorithm for Interactive Frame Rates During Visualization of Complex Virtual Environments," *Computer Graphics Proceedings, Annual Conference Series*, 1993.

[Glassner95] Glassner, Andrew S., *Principles of Digital Image Synthesis*, Appendix F., Morgan Kaufmann Publishing, 1995.

[He91] He, Xiao Dong, Kenneth E. Torrance, Francois X. Sillion, Donald P. Greenberg, "A Comprehensive Physical Model for Light Reflection," *Computer Graphics Proceedings,* Volume 25, Number 4, 1991.

[Karner96] Karner, K.F., H. Mayer, M. Gervautz, "An Image based Measurement System for Anisotropic Reflection," *EUROGRAPHICS Annual Conference Proceedings*, 1996.

[Lewis93] Lewis, Robert R., "Making Shaders More Physically Plausible," *Fourth EUROGRAPHICS Workshop on Rendering,* June 1993.

[Nimeroff94] Nimeroff, Jeffry S., Eero Simoncelli, Julie Dorsey, "Efficient Re-rendering of Naturally Illuminated Environments," *Fifth EUROGRAPHICS Workshop on Rendering*, June 1994.

[Schlick93] Schlick, Christophe, "A Customizable Reflectance Model for Everyday Rendering," *Fourth EUROGRAPHICS Workshop on Rendering*, June 1993.

[Spencer95] Spencer, Greg, Peter Shirley, Kurt Zimmerman, Donald P. Greenberg, "Physically-Based Glare Effects for Digital Images," *Computer Graphics Proceedings, Annual Conference Series*, 1995.

[Teller93] Teller, Seth, Pat Hanrahan, "Global Visibility Algorithms for Illumination Computations," *Computer Graphics Proceedings, Annual Conference Series*, 1993.

[Tumblin93] Tumblin, Jack, Holly E. Rushmeier, "Tone Reproduction for Realistic Images," *IEEE Computer Graphics and Applications*, November 1993.

[Ward92] Ward, Gregory J., "Measuring and Modeling Anisotropic Reflection," *Computer Graphics*, Volume 26, Number 2, July 1992.

[Ward94a] Ward, Greg, "A Contrast-Based Scalefactor for Luminance Display," *Graphics Gems IV*, editted by Paul Heckbert, Academic Press, 1994.

[Ward94b] Ward, Gregory J., "The RADIANCE Lighting Simulation and Rendering System," *Computer Graphics Proceedings, Annual Conference Series*, 1994.