A Ray Tracing Framework for Global Illumination Systems

Peter Shirley *
Department of Computer Science
Indiana University

Kelvin Sung †
Department of Computer Science
University of Illinois at Urbana-Champaign

William Brown ‡
Department of Computer Science
University of Illinois at Urbana-Champaign

Abstract

The fundamental software components useful for a zonal ray tracing system are described. The interface protocols and some implementational observations are outlined for each of the key components. An example global illumination system assembled from the components is presented.

CR Categories and Subject Descriptors: I.3.0 [Computer Graphics]: General; I.3.6 [Computer Graphics]: Methodology and Techniques.

Additional Key Words and Phrases: Ray tracing, radiosity, object-oriented design, zonal method, visual realism.

1 Introduction

In recent years many researchers have investigated ray tracing solutions to the global illumination problem [47, 13, 12, 26, 46, 45, 1, 40, 23]. Other authors have applied object-oriented design to ray tracing programs [28, 22]. In this paper, we discuss the major components[1] of an object oriented ray tracing system for global illumination.

There are several existing systems that enable programmers to construct user interfaces from basic components (e.g menus, scrollbars, windows). In a similar fashion, graphics programmers should be able to compose a global illumination renderer from basic components. Such components are specified only by their interface protocols, so different implementations of the same component can be invisibly exchanged in a particular system, which allows isolated testing of particular component implementations. This, of course, is simply an instance of abstract data typing.

The design strategy outlined above is not novel. We are simply applying this strategy to the global illumination problem, and thus identifying the key components of such a system. In this paper, we focus on the functionality and interface protocols of these components. We also describe some particular component implementations, and discuss interesting observations about these implementations. Finally, the paper summaries by providing an example of a full global illumination system assembled from the basic components.

The fundamental components of a distributed ray tracing system are described in Section 2. This extends the basic framework presented by Kirk and Arvo [28]. In Section 3 we describe the new components that are needed for zonal calculations, and discuss how the ray tracing components can be reused. An example global illumination system assembled from these components is described in Section 4.

2 Ray Tracing Components

Heckbert observed that geometrical objects should be viewed as basic components with common interface protocols, so that a ray tracing system can be designed independently from the geometrical primitive types (e.g. spheres, polygons) [22]. Kirk and Arvo have generalized this idea, by observing that ray-object intersection search structures, such as octrees, have the same interface protocols as geometrical primitives. This generalization allows the ray tracing system to be designed independently of both geometrical primative types, and ray-object intersection optimization strategy [28]. Kirk and Arvo also described a shader component, which encapsulates the local illumination model as a basic compo-

*Department of Computer Science, Lindley Hall, Indiana University, Bloomington, IN 47405. Email: shirley@cs.indiana.edu.

†Dept. of Computer Science, University of Illinois, Urbana, Illinois 61801. Email: ksung@cs.uiuc.edu

‡Dept. of Computer Science, University of Illinois, Urbana, Illinois 61801. Email: brown@cs.uiuc.edu

[1] Components are defined by their behavior and their interface protocols. This is the same as the *black box framework* described by Johnson and Foote [25].

nent. A similar shader concept has been advanced by various researchers [44, 21].

In this section, we describe the basic components of our ray tracing framework. In Section 2.1 we present several simple classes that are useful in graphics programs. The components which manage sample distributions on the pixel are discussed in Section 2.2. The ray-object intersection component is described in Section 2.3. The ray-material interaction component, our variation of shader, is outlined in Section 2.4. In Section 4 we show how these components can be assembled in a specific ray tracing system.

2.1 Graphics Utility Classes

Some utility classes were obviously needed: points, vectors, 4 by 4 transformation matrices, and colors. Operations involving these classes are defined using operator overloading mechanisms. Only operations that have mathematical meaning are allowed as suggested by Ronald Goldman. The allowed expressions are:

1. vector + vector = vector
2. vector - vector = vector
3. point + vector = point
4. point - vector = point
5. rotation-matrix * rotation-matrix = rotation-matrix
6. point * rotation-matrix = point
7. vector * rotation-matrix = vector

Denying operations such as the addition of two points makes mistakes in expressions easier to catch. Unfortunately, it also makes it difficult to take the centroid of a set of points. In our experience, the benefits of these operational restrictions have greatly outweighed the deficits.

It is often useful to restrict a vector to be of unit length, such as when dealing with direction vectors in a ray. The unit-vector class has been defined as a subclass of vector. The unit-vector class can be confusing because some operations such as a matrix multiplication can change its length. A partial resolution to the problems of rigor versus efficiency is to let the unit-vector inherit all of the operations of the vector type. For example, we can have 'vector = vector + unit-vector'. A unit vector can be initialized by an assignment operator 'unit-vector = vector', where the assignment occurs after automatic normalization of the vector value. Unit vectors themselves have a subclass, surface-normal-vector, which responds to transformations differently than unit-vectors.

Colors have the operators of addition, multiplication, division, and subtraction with themselves, and multiplication and division with scalars. Initially, a RGB color model was used. This was later replaced by a very general spectral model, where each color was approximated by a piecewise linear approximation with arbitrary node locations. When two colors were combined by an operation, a new color with possibly more nodes was generated. This meant that unbounded lists had to be used to store the node locations and amplitudes. Though this general color representation was well suited for complicated spectral distributions, such as the emission spectrum of a fluorescent light, and was spatially efficient for simple spectra, the time needed to manage the color variables was too large to justify the switch from RGB (typically, the run time more than doubled for a 'typical' image). Next, we switched to twenty evenly spaced nodes were used to represent color spectra. Because this had a high storage cost (especially for images using zonal calculations), the four unevenly spaced node locations suggested by Meyer was used [30]. There was no qualitative loss in image quality observed going from the twenty to four samples, but this may say more about the arbitrary nature of most of the input spectral curves than about the quality of the spectral approximation. We suspect that filtering the input spectra before point sampling them will avoid most problems associated with only using a few sample locations.

Another basic utility class is the ray class, represented by a point of origin, and a unit-vector representing direction. The ray class has no allowed operations, but does have the basic 'method' (member function) that finds the point a certain distance along the ray. We have found that it is helpful to associate two other characteristics with a ray. The first is a material id which stores the material the ray is in (e.g. 'this ray is now traveling through glass'). The second characteristic is the attenuation of the ray. This makes the bookkeeping associated with adaptive ray tree pruning [20] straightforward.

A utility class that has been surprisingly useful is an orthonormal basis of three vectors. Though any two of these vectors uniquely defines the basis, all three are explicitly stored, trading space for execution time. These bases are used in viewing calculations, and to assign a local coordinate system to a surface. The basis is defined by three orthonormal vectors, $\vec{\mathbf{u}}$, $\vec{\mathbf{v}}$, and $\vec{\mathbf{w}}$.

Another utility class is the (u, v) pair. This class is useful for pixel sampling and texture mapping. The texture class itself adds another useful utility. In this implementation, the texture abstract class takes both a point in 3D, and a (u, v) pair. Surface textures and solid textures [34] are both subclasses of the texture class. A solid texture will use the point for texture generation, and the surface texture will use the (u, v) coordinate for color lookup. This 'send all information, needed or not' strategy can be inefficient, but it is a simple way to guarantee that the needed information is passed to a texture module.

A very useful utility class is a solid noise generator. We have implemented the generator given in Perlin's 1989 paper [35], and found it to be quite mechanical to implement. While it has been demonstrated that solid noise is useful in generating uniform textures such as marble, we have found that these techniques are too simple for semistructured patterns

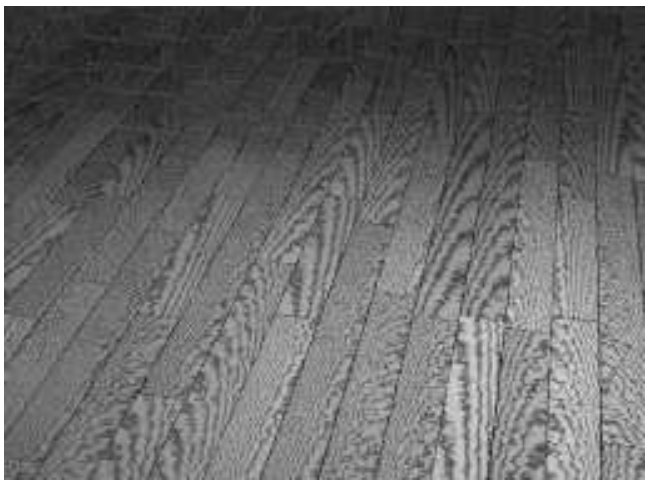

Figure 1: A board floor generated from a procedural solid texture.

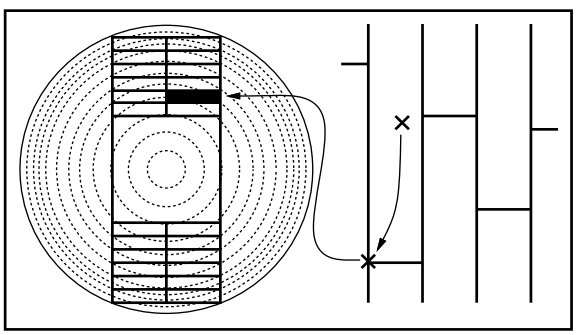

Figure 2: Use one corner of o board to find a solid noise seed for an index into the tree.

such as board floors, brick walls, and carpeted floors. Each element of these patterns (e.g. a specific brick) can be modeled using various transformations on solid noise. This needs to be added one or more randomly or semi-randomly varying parameters to make the elements vary among themselves.

As an example, a single board has a random grain pattern associated with a planer surface passing through a tree. The woodgrain of a tree is modeled in a manner similar to Perlin's marble: the basic structure of the tree is light and dark concentric rings, where the area (not distance) between rings is roughly constant. The radius of a given ring will decrease slightly as it moves up the tree (it is visually important that the rings are not perfect cylinders). The noise is used to add some irregularity to the geometry of the rings. The particular location of a board in a tree is what gives its unique character. Since the tree is of finite volume there are a finite set of specific boards that come form a tree, and each can be given a specific integer id.

We create our board texture by putting down an algorithmic (predefined) board pattern, and associating a specific board id to each board. This is easily accomplished by using the solid noise at a particular corner of the board as a seed to generate the board location in the tree (Figure 2). This guarantees that all points on a board are mapped to the same id in the tree. This procedure gives some irregularity to what types of boards (fine-grained versus loose-grained) are adjacent in the pattern. To ensure that there is no visible correlation between the ids of adjacent boards, we scale the corner points by some large factor before calling the solid noise function. An example floor generated by this method is shown in Figure 1.

This same basic idea can be used to generate bricks. Solid noise evaluated at one corner of a brick will give a seed to control some global parameter for that brick. If more than one parameter is needed, the other corners (or midpoints, etc.) can be also used for solid noise seed points. To generate the basic colors of wood or brick, we attempt to use physical spectral curves. Unfortunately this data can be hard to come by, and directly specifying spectral curves is not intuitive. In these cases we generate curves by 'mixing' standard artist's pigments, the curves for which can be found in [9]. When we have existing rgb data, we use Glassner's conversion method [17]. It has been our experience that using smooth curves for this conversion is highly preferable to using impulses.

2.2 Sample Point Generation

A basic module in any ray tracing code selects sample points on the pixel (in some systems the samples are chosen more glabally on the screen). In a stochastic ray tracing system, sample points must be chosen from a larger multidimensional space (e.g. screen-lens-reflection-shadow-time space). Cook recommends choosing sample point module one or two dimensions at a time and then combining these choices for full multidimensional sample distributions. For example, we

might choose screen locations uniformly on the pixel area, and lens points uniformly on the lens disk. These points could then be paired randomly or according to some fashion to form sample points in four-dimensional screen-lens space. Cook calls his specific version of this method *uncorrelated jittering*. There has been much discussion on how to get a 'good' set of sample points on a square two-dimensional region [13, 31].

Many spaces we need to sample (e.g lens area, reflection ray direction) are not square. One way to generate sample points on a non-square region is a special purpose algorithm, such as the one used by Cook for reflected rays [12]. The other is to generate points on the square and then apply a warping function so that their distribution is changed. An example of this method is the transformation used by Ward et al. to generate a cosine distribution of sample points on a hemisphere [46]. The first abstraction is to make the sampling distribution come from an abstract class. This lets the user flexibly choose and add sampling methods and filter functions in a natural way.

In this section we discuss two basic modules useful for sample point generation. The first generates sample points on the unit square, and the second module acts as a filter to create specific sample point distributions by 'warping' uniform distributions.

2.2.1 Uniform Sample Point Generation

The sample-generator module is passed a non-negative integer, n, and generates n (u, v) sample points that are in some sense uniformly distributed on the unit square $(0 \leq u \leq 1, 0 \leq v \leq 1)$. The basic implementations of this module are commonly regular sampling, jittering, and poisson disk sampling. All of these strategies have problems. Regular sampling can cause extreme aliasing. Jittering requires n to be a perfect square for optimum performance, and requires some factoring system otherwise ($n = 17$ would imply a 1 by 17 sampling partition). Poisson disk has non-deterministic running time, and choosing the disk radius is not straightforward. Mitchell addressed the first problem with an approximation to poisson disk sampling based on error diffusion [31].

We have implemented these strategies and one new strategy, *n-rooks sampling*. N-rooks is a variation of Cook's uncorrelated jittering. We generate the u coordinates of the n points by jittering in one dimension, and then generate n v coordinates in the same way. We then randomly link the u and v coordinates to form n (u, v) pairs. This forms n sample points in a way that is unbiased in relation to the probability of selecting any particular points. Another way to look at the sample generation is to draw a n by n chessboard, and randomly place n rooks on the board subject to the constraint that the rooks cannot capture each other in one move. This results in a pattern where every row and column hass exactly one rook. We now select one (u, v) pair from each square where a rook sits.

The n-rooks strategy is really a simplification of uncorrelated jittering (rather than a generalization), because none of the dimensions are really linked. N-rooks has the advantages that it can generate a set of n samples for any n, that it takes a deterministic amount of time, and that there are no parameters that need to be set other than n. It is not surprising that n-rooks performs well for pixels containing horizontal or vertical edges, because it fully jitters each dimension. To our surprise, it has so far almost always been more accurate than two-dimensional jittering or poisson disk sampling on the images we have tested. In our experience, two-dimensional jittering usually is more accurate than poisson disk sampling, regardless of the disk radius. These test results, and some analysis on why this might be true can be found in [38, 41].

2.2.2 Sample Distribution Transformations

Suppose we want to apply weighted area averaging to a pixel rather than simple area averaging [14]. We could simply apply weights to each sample point based on its (u, v) coordinate. If we sample each pixel independently it is better to apply importance sampling by placing the sample points in a way that is distributed according to the weighting function [12]. In this section, we describe a module that takes a set of n uniformly distributed points, and somehow warps them in a way that they are distributed according to the desired weighting function.

As a simple example, suppose we want the trivial 'box' weighting function, which simply samples the pixel area uniformly ($w(x, y) = 1$ inside the pixel). The warping module would take n sample points generated by a module described in the last section, and apply the mapping $x = u - 0.5$, $y = v - 0.5$. This assumes we are using a coordinate system where the pixel center is the origin and the pixel width is 1.

In practice we will want non-uniform weighting functions. For example, suppose we have the width 2 weighting function:

$$w(x, y) = (1 - |x|)(1 - |y|) \tag{1}$$

If we just wanted to generate independent random points with density w, we can apply standard techniques to transform canonical[2] random numbers into these points [29]. Since this particular w is separable, we can generate x according to $w(x) = (1 - |x|)$, and y the same way. We define the distribution function, F, associated with w:

$$F(x) = \int_{-1}^{x} (1 - |x'|)dx' = \frac{1}{2} + x - \frac{1}{2}x|x| \tag{2}$$

To get our desired independent (x, y) distributed according to w, we simply take canonical (ξ_1, ξ_2) and apply $x = F^{-1}(\xi_1)$, and $y = F^{-1}(\xi_2)$, where F^{-1} is the inverse function of F.

[2] A canonical random number ξ is uniformly distributed between 0 and 1.

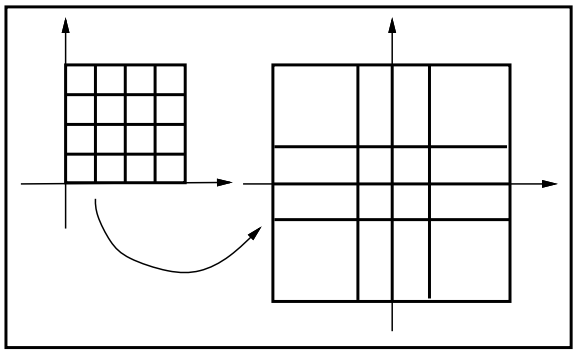

Figure 3: A transformation for nonuniform filter sampling.

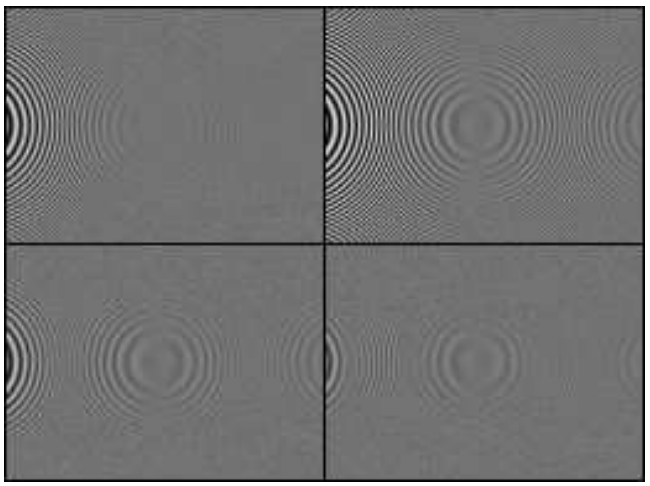

Figure 4: Images with different weighting functions. Left to right, top to bottom: $(1-|x|)(1-|y|)$, width one box, width 2 box, width 3 box.

This idea can be extended by taking our set of sample points and transforming them as if they were canonical. This will preserve some of the good qualities of the original sampling distribution, but will have our desired distribution properties. Applying this idea to the filter in Equation 1, we can take our n uniform sample points and apply the transformation:

$$x = \begin{cases} -1+\sqrt{2u} & \text{if } u < 0.5 \\ 1-\sqrt{2(1-u)} & \text{if } u \geq 0.5 \end{cases}$$

followed by a similar transformation involving y and v. This transformation on 16 jitter cells is shown in Figure 3. Note that this will transform nonuniformly within each cell, so the sample point selection must take place before the transformation.

Figure 4 shows two different weighting functions applied to the image intensity function $I(r) = 0.5 * (1 + \sin(r^2))$, where r is the distance from the middle of the center of the left edge of the image. Note that the non-uniform weighting function reduces aliasing.

The uniform sampling transformation can also apply to non-separable density functions, and to functions defined on non-Cartesian manifolds. This requires dealing with the joint distribution function and non-constant metrics, but otherwise the same techniques apply. Details can be found in [37, 38].

We have found several other transformations to be useful. For example, to choose points uniformly from a disk of radius R, apply the transformation $\theta = 2\pi u$, $r = R\sqrt{v}$. To choose points in a cosine distribution on a hemisphere $(p(\theta,\phi) = (1/\pi)\cos\theta)$, apply the transformation $\theta = \arccos(\sqrt{1-u})$, $\phi = 2\pi v$. A generalization is to choose directions according to a 'phong' distribution $(p(\theta,\phi) = ((n+1)/(2\pi))\cos^n\theta)$, by applying the transformation $\theta = \arccos((1-u)^{1/(n+1)})$, $\phi = 2\pi v$. To choose random points on a triangle defined by vertices p_0, p_1, and p_2, apply the transformation: $a = 1-\sqrt{1-u}$, $b = (1-a)v$, and the random point p will is: $p = p_0 + a(p_1 - p_0) + b(p_2 - p_0)$.

2.3 Ray-Object Intersection

From a functional point of view, the ray-object intersection module can be defined as: given a ray, is there an object hit by this ray; if so, what is the object hit and what is the point.

With this simple definition, this component can be implemented as a black box framework [25]. Our goal here is to present the design of a framework to facilitate the independent development of different algorithms for this component. We use the *Faster Ray Intersection Techniques* as described by [8] as a base to discuss our design and show that our design unifies the approaches. This section concludes with some interesting implementation observations.

2.3.1 Class Hierarchy

Like Kirk and Arvo [28], we implemented geometrical primitives (e.g. sphere, polygon) and collection structures (e.g. octree, bounding volume) as subclasses of the same *geom-object* abstract class. In this way, they can have the same interface protocols and behaviors. This allows a particular collection structure (e.g. octree) to contain, as members, other collection structures (e.g. bounding volume, regular grids). For example, both geometrical primitives and collection structures respond to the *Hit(ray)* message by returning the *geom-object* hit by the input ray. Since both geometrical primitives and collection structures are subclass of geom-object, the result of *Hit* message could be either. In this way, it is possible to have nested octrees or have regular grid nested inside an octree structure. Intelligent use of this generality can reduce running time. Unfortunately, just how to attain intelligent use is not obvious.

This general grouping of spheres, polygons, octrees, etc. emphasizes an important point: classes with common access (member) functions should be subclasses of the same ab-

stract class, even if their underlying representations are very different. Classes with different access rules, even if their underlying representations are identical, should not be grouped together, as seen with points and vectors in the previous section.

2.3.2 Functional Modules

It is well known that linear exhaustive testing of primitives with rays for potential intersection is inadequate. The geometry model should be processed into some internal representation (e.g. collection structures like octree) for efficient candidate primitive look up during intersection testings. Since current ray tracers and zonal renderers only model geometrical optics [26], the *processed internal representation* could be viewed as an efficient *primitive/geometry-characteristic* association list storage and look up mechanism.

From the design level, the ray-object intersection component consists of three functional modules: Builder, Structurer, and Traverser. For each primitive, the builder module inquires certain geometry characteristic from the primitive, and sends requests to the structurer module to associate the characteristic with the primitive. The traverser module extracts the relevant geometry characteristic from an input ray and requests the structurer module to look up the associated primitives stored. It is a primitive's responsibility to calculate for possible intersections. Under this design, the ray-object intersection component becomes the *coordinator* that assists and thus speeds up the process of searching for potential intersecting primitives.

The acceleration techniques for faster intersection calculations as described in [8] are unified under this design:

Bounding Slabs [27]. The *Builder module* inquires the bounding slab distances from the primitive. In this case, the bounding slab is the "geometry characteristic". The *Structurer module* associates the primitive with the corresponding slab and slab distances. The construction of the bounding slab hierarchy is also the responsibility of the structurer module. The *Traverser module* takes a ray geometry and pierce the bounding slab hierarchy to locate potential intersecting primitives. Notice that a bounding slab is also a *primitive* in the sense that it must respond to the *Hit* protocol message by returning the primitives (including children slabs) bounded by the current slab.

Spatial Subdivision [16, 3]. The *Builder module* inquires the extent from a primitive and determines which spatial cell should be associated with the primitive. The actual implementation of spatial cell and the association between the primitive is performed by the *Structurer module*. The Structurer module stores the association of primitives and spatial cell units in such a way that later retrieval of primitives can be done efficiently (when given a ray). The *Traverser module* takes a ray geometry and traverses the spatial structure to locate the potential intersection primitives. Again, a spatial cell unit must be able to respond to the *Hit* protocol message by returning primitives (including children spatial cells in case of adaptive subdivision algorithms) contained in the current spatial cell.

New algorithms could be formulated by identifying the geometry characteristic that the approach is taking advantage of. For example, the *ray coherence theorem* [32] uses a direction and an acute angle as the geometry characteristic. In this case, the builder module finds out the list of candidate primitives that are visible from a direction/angle pair of a primitive. When tracing a secondary ray, the traverser module uses the ray direction and reflecting angle to request the structurer module to extract the candidate primitives associated with the corresponding direction/angle pair for potential intersection calculations.

The reason to separate this component into three modules is such that the implementation of each module can be isolated and replaced without affecting the others. For example, for an octree spatial subdivision implementation, the builder can build the octree structure statically before ray tracing starts, or dynamically during ray tracing. The structurer module is the underlying implementation of an octree, this could be done by a hash table, or a hierarchy of pointers. The traverser module knows how to traverse an octree, but does not need to know the detail implementation of the octree structure. Some of the examples of octree traversers are: linear octree walking[3], Glassner's algorithm [16], and DDA octree traverser [43].

2.3.3 Implementation Notes

Following the design approach described in this section, it is possible to replace a module in this component with the rest of the system remaining unchanged. As a result, we are able to isolate and observe the effect of different algorithms. Some interesting observations we thus get are described here:

Octree Traversal. It is believed that the original octree traversal algorithm [16] can be improved by realizing *tree location coherence* [15]. In his original octree paper [16], Glassner proposed to always begin the search for next octant from the root of an octree. It has been pointed out by various researchers [24, 15] that theoretically, starting the next octant search from the parent of current octant should be faster (this approach has been called *tree walking* [15]).

However, Sung observed that in a zonal renderer, tree walking may be slower than that of the Glassner's original algorithm [43]. We note that the only difference between the two algorithms is in the process of getting the next octant: tree walking reconizes the tree location coherence and start searching from the parent of the current octant, while the Glassner's algorithm always start searching from the root of the octree. A fact that has been generally overlooked is that it takes extra time for the tree walking algorithm to ascend

[3] Arvo suggested and discussed the *linear octree walking* algorithm in Volume 2 of *The Ray Tracing News*, which is currently available under anonymous ftp from drizzle.cs.uoregon.edu

the octree. When traversing between two octants of different first level parents, the two algorithms take the same amount of time to descent the octree, while tree walking needs extra time to ascend the octree. The energy transferring process of a zonal renderer is often implemented by sending *energy rays* between patches. Typically, patches are separated by empty spaces. This implies that sending energy rays involve traversing from deep[4] in an octree across shallow octants (empty space) to deep octants (target patches). Because it takes extra time for tree walking to ascend from the deep octants, Glassner's algorithm may be faster in these cases.

Mail box. The mail box concept was proposed independently in [6, 5]. The idea is to avoid multiple ray-object intersection calculation between the same object and the same ray in different spatial cell units. The underlying assumption is that the primitive objects usually stretch across many spatial cell units, thus the time saved in avoiding multiple intersection calculations offsets the overhead involved in maintaining the extra information. This assumption is generally true for a traditional ray tracer. For example, it is possible to define and ray trace a large cylinder primitive. In a zonal renderer, all primitives must be turned into patches. In order to generate high quality final image, it is often necessary to generate relatively small patches [11]; thus all large primitives need to be break down into small patches. The probability of these small size patches stretching across spatial units are smaller. Also, since there are a larger amount of primitives (e.g. a cylinder primitive is turned into many small primitive patches), the overhead involved in performing the pre-intersection checking and post-intersection information recording for every primitive becomes greater. Sung observed [43] that implementing mailbox in a zonal renderer does not achieve significant speed up (less than 5%). This analysis may not apply if the decomposition into zones is simply done in the objects parametric space with no geometric changes being made. In summary, the benifits of the mailbox technique may be highly dependent on the general geometric character of the geometric primatives used.

2.4 Ray-Material Interaction

Several researchers have noted that reflection behavior should be encapsulated as one unit of a rendering system [28, 44, 21]. We treat light-material interaction as a component, where the reflection behavior is determined strictly from a set of material parameters. Traditionally this might be accomplished with one shading model with parameters including ks, kt, and kd [19]. One problem with such an approach is that physically implausible parameter combinations can be chosen by the user (e.g. kd = ks = 0, kt = 1). Implausible combinations may be useful for many applications, but if realism is desired, it is better to limit the user's choices.

We have used the idea that materials can be classed as families, each grouped by the parameters that effect their behavior. This way the user only needs to choose the relevant parameters for a particular material. Once the material is chosen it is treated as a black box component that responds to a limited protocol (much like geometrical objects' for ray-object intersection). The first way in which a material can be queried is by asking, 'given an incoming ray r, a point p on the surface, and a surface normal $\vec{\mathbf{n}}$, what rays r_i are reflected/transmitted by the material, and what is the attenuation k_i for each ray'. This will allow us to handle building the ray propagation code. For other lighting calculations, such as the direct lighting component, we need to query a material, 'what is your radiance, $L(\vec{\mathbf{v}}_{out})$, that comes from a source of radiance $L(\vec{\mathbf{v}}_{in})$ that subtends a solid angle ω?'. We also need to ask a material if it is a luminaire (source of light), and if so, how much light it emits in a particular direction.

The materials that we have implemented are:

conductor: Parameters n (refractive index), k (extinction coefficient), e (phong-style exponent). Example: aluminum.

dielectric: Parameters n (refractive index), a (filter coefficient), e (phong-style exponent). Example: glass.

lambertian: Parameter k_d (diffuse coefficient). Example: matte paint.

polished: Parameters k_d (diffuse coefficient of substrate), n (refractive index of polish), e (phong-style exponent). Example: gloss paint.

translucent: Parameters k_d^1 (diffuse coefficient of first side), k_d^2 (diffuse coefficient of first side), kt (transmission coefficient). Example: lampshade.

luminaire: Parameters k_d (diffuse coefficient), e (phong-style exponent). Example: light bulb.

These basic materials can be extended, but they have proven to be fairly good approximations to common real world materials. Conductors are sometimes a little difficult because the parameters n and k are' not intuitively controllable. We have found most data for conductors in [33]. The behavior of both conductors and dielectrics is determined using the Fresnel Equations, the full form of which can be found in [42, 38]. The polished surface is an approximation to a diffuse substrate with a thin dielectric covering. This means that for a given direction we first calculate the specular reflectivity $k_s(\theta)$, and then the remaining light is reflected diffusely, giving a diffuse reflectance of $(1 - k_s(\theta))k_d$. This allows glare effects to be approximated accurately. The phong-style exponent e is used to allow some spread in the reflected component of conductors, dielectrics, and polished materials. For smooth surfaces e is set to a large number.

The translucent surface reflects light diffusely from either side, and also allows some light to be diffusely transmitted. The luminaire acts as a diffuse reflector, and also emits

[4] All objects in a zonal scene are constructed by patches. The space containing objects usually need more octree subdivision, and thus are *deeper*. By the same token, empty space octants do not contain objects and do not require much subdivision, and thus are *shallower* in an octree.

power in a phong-style distribution. Large exponents are used if spot lights are desired.

The ray reflection/transmission behavior can be summarized as:

conductor: Generate one reflected ray with attenuation determined by Fresnel Equations. Perturb ray stochasticly if $e \neq \infty$.

dielectric: Generate one reflected ray and one transmitted ray with attenuations determined by Fresnel Equations. Perturb both rays stochasticly if $e \neq \infty$.

lambertian/luminaire: Generate one reflected ray stochastically with a cosine distribution.

polished: Generate one reflected ray from the polish with attenuation determined by Fresnel Equations. Perturb ray stochasticly if $e \neq \infty$. Generate one reflected ray from the diffuse substrate stochastically with a cosine distribution.

translucent: Generate two rays stochastically, one reflected, one transmitted, each with a stochastic cosine distribution.

It is useful to be able to turn off reflections from a particular material. We allow this to be done when the material is initialized. A conventional Whitted-style ray tracer would turn off reflections for the lambertain, translucent, and luminaire surfaces, and would turn off reflections off the substrate (but not the polish) of the polished surfaces. To maintain some form of dependent sampling, such as uncorrelated jittering, the reflection protocol should also accept a canonical (u, v) pair (Section 2.2), to be used as a basis for any probabilistic reflection that might occur.

3 Zonal Calculations

Several recent zonal (radiosity: global lighting information is stored at a finite set of 'zones') systems are based on progressive refinement techniques [10, 2]. The theoretical basis for such systems is straightforward to extract. If the progressive refinement is viewed as power transport simulation, which implies fairly direct non-diffuse zonal solutions [7, 40, 18, 39, 36]. These solutions are easy to construct if we view the zone as a black box which collects power carrying rays, and later emits a group of power carrying rays that represent reflected power accumulated since the previous emission step.

This abstraction underlying zonal calculations can be stated: zones should receive, accumulate, and send power, and the mechanics of how this happens should be hidden. This is accomplished by defining a zonal-data module. In addition, the module should, after the zonal calculations are completed, be able to provide the radiance of the patch when viewed from a certain direction.

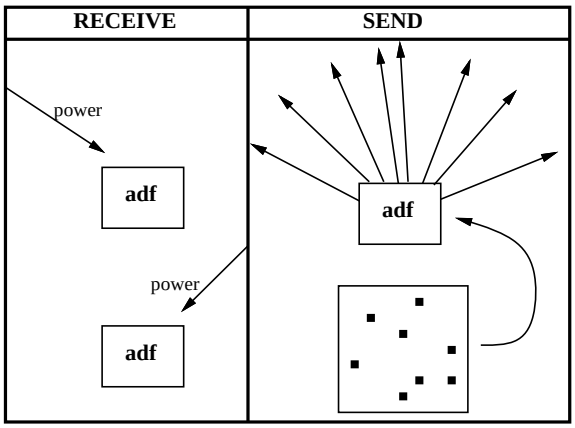

Figure 5: The two crucial methods of and adf (angular distribution function): receive power, and send power according to some set of sample points.

For a lambertian zone, the zonal-data module is easy to implement because there is no dependancy on the incoming direction of intensity. We need to store the total power, Φ, and the unsent accumulated power, Φ_u. Each new incoming ray carrying power Φ_i will imply $\Phi = \Phi + k_d\Phi_i$ and $\Phi_u = \Phi_u + k_d\Phi_i$. When it is time for the zone to emit, it will send N rays each carrying power Φ_u/N. These rays will be sent in a cosine distribution. Just as done with pixel sampling in Section 2.2, we can derive N (u, v) pairs and then transform these to the appropriate (θ, ϕ) pairs. The radiance of the zone will just be $\Phi/(\pi A)$, where A is the area of the zone.

For a zone with directionally dependent reflection behavior, such as brushed steel, we must maintain the total and unsent power as some kind of directional table [18, 39, 36]. A simple way to do this is a simple spherical coordinate array of bins, with the total power going through each bin. The unsent and total power of the diffuse case must be generaized to a new black box, the angular distribution function (adf). This function maintains whatever information is necessary to receive power, and later send power as a set of rays (Figure 5).

The receiving method of a directionally dependent adf can be implemented by using the ray-material interaction module of Section 2.4. Simply reflect the incoming ray using the ray-material interaction module as a black box, and add the attenuated reflected power to whichever angular bin(s) the reflected ray(s) land in. The sending stage can be implemented using the warping methods of Section 2.2.2, or by independantly sending a pattern of N rays from each angular bin. Because the adf actually stores spectral values, it can only be converted to a probability density by converting the entries to scalars. We use luminance to do this.

It would be very wasteful of storage to store an explicit directional table for diffuse surfaces, though it would still work. We therefore implement a lambertian adf by storing only the total and accumulated power. The black box

interface still looks the same to the zonal module. To accomplish this in an extendable way, we add an access function to the ray-material interaction class which tells whether the reflection behavior is directionally dependent or independent, and whether the surface is reflective or reflective and transmissive. If the material is reflective and directionally independent (e.g. lambertian), it will use a adf module that internally only stores total and unsent power. If it is reflective and transmissive and directionally independent (e.g. translucent), then these quantities will be maintained both above and below the surface. If directionally dependent, the directional tables will be maintained either for $(0 < \theta < \pi/2)$ or $(0 < \theta < \pi)$ depending on whether the material is transmissive. If the material can provide an estimate of specularity (e.g. the phong exponent), then this can be used to choose the resolution of the table.

Once the adf modules have been initialized in each zonal module, their internal representation in invisible. This allows the programmer to detatch the local lighting models from the global light transport.

The zonal-data abstract class accomplishes two important functions. First, it removes any reference to surface reflection type from the light transport code. This makes the code more readable and allows new reflection types to be added in a modular fashion. The second important function is that it allows variable storage for the zonal-data of different reflection types. Thus, adding a surface with a large directional table does not force the lambertian surfaces to use more memory.

If the zone is textured, its average reflection properties must be found. To avoid aliasing problems, the methods of Section 2.2 are used to point sample the zone to estimate the average properties.

3.1 Reducing Precision Problems

One advantage of ray tracing is that the programmer does not have to be concerned with maintaining the 'correct' outward facing surface normals, because we know whether a ray is 'inside' or 'outside' simply by counting surface crossings. For zonal methods, however, we need to emit power carrying rays toward the 'outside', so we need to maintain an outward facing normal. This is unfortunate, but it can be used to our advantage by eliminating the infamous ϵ test.

When generating a parametrically defined ray $(\mathbf{o} + t\vec{\mathbf{v}})$ on a surface, as is done with reflected rays, power carrying rays, and shadow rays, we will always have an intersection at point $\mathbf{o}$ at approximately $t = 0$. To avoid complications that arise when imprecision causes the hit to be represented as a slightly positive number, we look for the first hit where $t > \epsilon$, where ϵ is a small positive number that bounds the possible roundoff error. Amanitides and Mitchell[4] showed that problems can arise when there are real surfaces closer than $t = \epsilon$ and provided strategies to use in these cases.

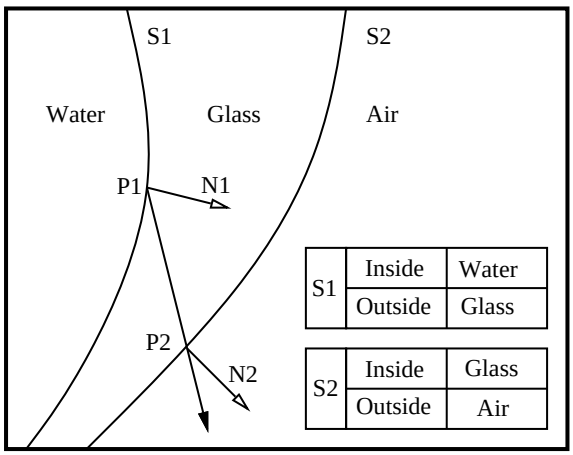

Figure 6: Elimination of ϵ-test.

We can eliminate the ϵ test by using the surface normal information. Each surface can be viewed as an interface between two materials[5]. The first material can be viewed as the inside material (facing away from the normal), and the second is the outside material. As mentioned in Section 2.1, each ray stores the material it is currently in. This means a dot product of the ray direction with the normal of a surface will tell the ray which material is in front of the surface from the point of view of the ray. If this material is not the same as the material the ray is traveling through, then the ray cannot hit that surface. This idea is illustrated in Figure 6, where a ray originates at point P1 and propagates into glass (the outside material). When the ray is tested for intersection against S1, the dot product between the ray direction and N1 indicate that the ray is going inside to outside and is thus approaching from the water side. Since the air is in glass, we know there can be no intersection, so no ϵ test is needed. For S2, the inside material is glass, so the intersection point P2 is taken to be valid.

4 Example Framework Usage

We have implemented a simple global illumination system based on the framwork components and components described in Sections 2 and 3. The first phase of the system reads in the geometrical data and other data such as antialiasing parameters. It also specifies which surfaces will be treated as zones, and whether direct lighting will be handled by the zonal phase or the viewing phase. Those surfaces that are not to be treated as zones are not assigned zonal components. A global octree is used as an efficiency structure housing all geometrical primatives.

The zonal comonent was modified to separately store the total power adf in a direct and indirect lighting component,

[5] Our translucent surface does not fit this definition, so some care must be taken when assuming such a model

and to also have an adf to represent emitted power. Those zones that are luminaires have their emitted power adf initialized, and non-luminaires have null emitted power adfs. The zonal phase of the system then begins cycling through the zones, having each zone send N power carrying rays, where N is set to be proportional to the power stored in the unsent adf of the sending zone. If a surface without a zonal comonent is hit by a power carrying ray, that surfaces material comonent is used to scatter the ray from the surface. After some user set number of full interreflections (usually 5-10), this phase ends and each zone's radiosity comonent is sent a message that it can now perform bookkeeping that will allow queries about radiance values of the zone. The radiosity comonent will pass this message on to the direct and indirect adfs, which will adjust themselves accordingly. For directionally dependent adfs, this simply converts the power through each bin, to radiance through each bin. Interpolation to bin vertices could then be performed.

Once the zonal phase has placed radiance values at each zone, a distributed ray tracing phase creates an image. The ray tracer uses the sampling comonent to create a (u, v) pattern of n samples, and then uses a filter comonent to warp these sample points to approximate a user-specified weighing function. This generates a ray that we need to find a radiance value for. A simple function that returns a color seen along the ray finds the object intersected, recursively calls itself for any reflected rays specified by the material function, and uses zonal information for the indirect lighting. If the user specifies direct lighting to be zonal, the direct radiance is also taken from the zonal component, and we the direct lighting plus any indirect lighting from the reflected rays is returned. If the direct lighting is to be calculated in the viewing stage, a standard direct lighting calculation is added.

To ensure even sampling of reflected ray direction, light source area, and camera lens area, we generate an additional three sets of n (u, v) pairs for each pixel, and pass one set of these pairs along so that each dimension can be appropriately sampled in a manner similar to uncorrelated jittering.

Figure 7 shows a picture produced using this example system. The lampshade is translucent, the walls are lambertian, and the floor is polished and has a board texture. Zonal transport is used for the indirect lighting, and the direct light is calculated in the viewing phase.

Figure 7: A lamp.

5 Conclusion

While experimenting with different ways of combining ray tracing and zonal methods, it became clear to us that it is essential to be able to assemble a global illumination renderer from some basic components. With such a system, at a global level, the graphics programmer will be able to experiment with different ways of assembling the renderer. While at a local level, different implementations of the same component can be exchanged invisibly. In this way, it would be possible to isolate the testing of different algorithms.

Two basic programming strategies have been used to achieve our goal. The first was the creation of toolbox classes of points, vectors, colors, rays, orthonormal bases, noise generators, texture coordinates, and textures. These toolbox classes were used as primitive types, much like integers and floats in numerical codes, in the creation of the image generation code. The second basic strategy was the identification of key components in a global illumination system and the creation of abstract classes to support these components.

We identified the key components of a stochastic ray tracer to be: sample point generator, ray-object intersector, and ray-material interactor. Corresponding abstract classes are defined to support these components: sample-generator allows different sampling strategies and distributions to be plugged into the basic ray tracing module, geom-object allows the addition of new object types and collection structures, and ray-material interaction class allows the defintion of new materials with distinct light intereation characteristics.

When identifying the key components for supporting zonal calculations, we realize that most of the ray trace components are reusable through subclassing. For example, the reflected energy ray distribution is a subclass of the sample point generator, and the surface energy reflection behavior is implemented by subclassing the ray-material interaction class. The zonal-data class hides the mechanics of power accumulation and re-distribution, and thus allows efficient storage systems to be implemented for each surface type.

One issue that remains unresolved is how to handle surfaces with complex material properties. For example, we could define a floor surface with alternating tiles made of marble and steel. The steel would respond to light as a metal, and the marble as a polished surface. It would also be desirable to be able to add a layer of dust (probably using a pro-

cedural texture), that would cover the marble and steel in a nonuniform manner, reducing the specularity of both. If the 'material' were to handle all shading in this situation, it would need access to steel, marble, and dust reflectance behavior, as well as the procedural texture describing the dust. This could be accomplished in a manner similar to a Renderman shader [44, 21], where the shading routine has access to the internals of reflection models and textures. Unfortunately, such a shader does not hide much information, and can become quite unwieldy. It would be very desirable to put the capabilities of a general shader in a class structure that preserves modularity and data hiding, but exactly how to create such a class structure is still a research topic.

Our discussion is summarized by the description of a global illumination system assembled from the defined components. Without altering of other components in the system, we have implemented different octree traversal algorithms and storage strategies. The different implementations of the component were plugged in and tested invisibly from the rest of the system. In this way, some observations that contradicts to common beliefs were made. For example, we notice that the tree walking algorithm is not always the better octree traversal choice; jittering may outperform poisson disk sampling and for a radiosity renderer, mail box may not be a worthwhile effort. During the course of the system development, we have changed our color models and sampling strategies several times. These were accomplished with only local code alterations. We believe all these indicate that we have identify a *useful* set of components for the construction of global illumintation system. One should note that there is no *correct final* set of components for any software system. As new algorithms are invented, the component design should be *refined* to support the implementation of the new algorithms.

6 Acknowledgements

Thanks to Greg Rogers for general help with object-oriented design, Claude Puech for discussion about the ray-object intersection, John Airey, Holly Rushmeier, and John Wallace for numerous radiosity discussions, Jim Arvo for providing insight and information about his ray tracing work, Don Mitchell for some expert guidance in sampling theory, Ralph Johnson, who motivated a design based on something besides pure hacking, and Special Thanks to William Kubitz, who was the advisor for this project.

References

[1] John M. Airey and Ming Ouh-young. Two adaptive techniques let progressive radiosity outperform the traditional radiosity algorithm. Technical Report TR89-20, University of North Carolina at Chapel Hill, August 1989.

[2] John M. Airey, John H. Rohlf, , and Frederick P. Brooks. Towards image realism with interactive update rates in complex virtual building environments. *Computer Graphics*, 24(1):41–50, 1990. ACM Workshop on Interactive Graphics Proceedings.

[3] Kansei Iwata Akira Fujimoto, Takayu Tanaka. Arts: Accelerated ray-tracing system. *IEEE Computer Graphics and Applications*, pages 16–26, April 1986.

[4] John Amanatides and Don P. Mitchell. Antialiasing of interlaced video animation. *Computer Graphics*, 24(3):77–86, August 1990. ACM Siggraph '90 Conference Proceedings.

[5] John Amanatides and Andrew Woo. A fast voxel traversal algorithm for ray tracing. In *Eurographics '87*, 1987.

[6] Bruno Arnaldi, Thierry Priol, and Kadi Bouatouch. A new space subdivision method for ray tracing csg modelled scenes. *Visual Computer*, 3:98–107, 1987.

[7] James Arvo. Backward ray tracing. *Developments in Ray Tracing*, pages 259–263, 1985. ACM Siggraph '85 Course Notes.

[8] James Arvo and David Kirk. A survey of ray tracing acceleration techniques. In Andrew S. Glassner, editor, *An Introduction to Ray Tracing*. Academic Press, San Diego, CA, 1989.

[9] Norman F. Barnes. Color characteristics of artists' pigments. *Journal of the Optical Society of America*, May 1939.

[10] Michael F. Cohen, Shenchang Eric Chen, John R. Wallace, and Donald P. Greenberg. A progressive refinement approach to fast radiosity image generation. *Computer Graphics*, 22(4):75–84, August 1988. ACM Siggraph '88 Conference Proceedings.

[11] Michael F. Cohen and Donald P. Greenberg. The hemicube: a radiosity solution for complex environments. *Computer Graphics*, 19(3):31–40, July 1985. ACM Siggraph '85 Conference Proceedings.

[12] Robert L. Cook. Stochastic sampling in computer graphics. *ACM Transactions on Graphics*, 5(1):51–72, January 1986.

[13] Robert L. Cook, Thomas Porter, and Loren Carpenter. Distributed ray tracing. *Computer Graphics*, 18(4):165–174, July 1984. ACM Siggraph '84 Conference Proceedings.

[14] James D. Foley, Andries van Dam, Steven K. Feiner, and John F. Hughes. *Computer Graphics: Principles and Practice*. Addison-Wesley, Reading, MA, second edition, 1990.

[15] Andrew Glassner. Implementation notes for ray tracers. *Advanced Topics in Ray Tracing*, 1990. ACM Siggraph '90 Course 24 Notes.

[16] Andrew S. Glassner. Space subdivision for fast ray tracing. *IEEE Computer Graphics and Applications*, 4(10):15–22, 1984.

[17] Andrew S. Glassner. How to derive a spectrum from an rgb triplet. *IEEE Computer Graphics and Applications*, 9(7):95–99, 1989.

[18] David Edward Hall. An analysis and modification of shao's radiosity method for computer graphics image synthesis. Master's thesis, Department of Mechanical Engineering, Georgia Institute of Technology, March 1990.

[19] Roy Hall. *Illumination and Color in Computer Generated Imagery*. Springer-Verlag, New York, N.Y., 1988.

[20] Roy Hall and Donald P. Greenberg. A testbed for realistic image synthesis. *IEEE Computer Graphics and Applications*, 3(8):10–20, 1983.

[21] Pat Hanrahan and Jim Lawson. A language for shading and lighting calculations. *Computer Graphics*, 24(3):289–298, August 1990. ACM Siggraph '90 Conference Proceedings.

[22] Paul S. Heckbert. Writing a ray tracer. In Andrew S. Glassner, editor, *An Introduction to Ray Tracing*. Academic Press, San Diego, CA, 1989.

[23] Paul S. Heckbert. Adaptive radiosity textures for bidirectional ray tracing. *Computer Graphics*, 24(3):145–154, August 1990. ACM Siggraph '90 Conference Proceedings.

[24] Frederik W. Jansen. Data structures for ray tracing. In L. R. A. Kessener, F. J. Peters, and M. L. P. van Lierop, editors, *Data Sturctures for Raster Graphics*, pages 57–373. Springer-Verlag, Netherlands, 1986.

[25] Ralph E. Johnson and Brian Foote. Designing reusable classes. *Journal of Object Oriented Programming*, pages 22–35, June 1988.

[26] James T. Kajiya. The rendering equation. *Computer Graphics*, 20(4):143–150, August 1986. ACM Siggraph '86 Conference Proceedings.

[27] Timothy L. Kay and James T. Kajiya. Ray tracing complex scenes. *Computer Graphics*, 20(4):269–278, August 1986. ACM Siggraph '86 Conference Proceedings.

[28] David Kirk and James Arvo. The ray tracing kernal. In *Proceedins of Ausgraph*, pages 75–82, July 1988.

[29] Donald Knuth. *The Art of Computer Programming, Volume 3*. Addison-Wesley, New York, N.Y., 1981.

[30] Gary W. Meyer, Holly E. Rushmeyer, Michael F. Cohen, Donald P. Greenberg, and Kenneth E. Torrance. An experimental evaluation of computer graphics imagery. *ACM Transactions on Graphics*, 5(1):30–50, January 1986.

[31] Don P. Mitchell. Generating antialiased images at low sampling densities. *Computer Graphics*, 21(4):65–72, July 1987. ACM Siggraph '87 Conference Proceedings.

[32] Masataka Ohta and Mamoru Maekawa. Ray coherence theorom and constant time ray tracing algorithm. In Toshiyasu Kunii, editor, *Computer Graphics 1987*, pages 303–314. Springer-Verlag, Tokyo, Japan, 1987.

[33] Edward D. Palik. *Handbook of Optical Constants of Solids*. Academic Press, New York, N.Y., 1985.

[34] Ken Perlin. An image synthesizer. *Computer Graphics*, 19(3):287–296, July 1985. ACM Siggraph '85 Conference Proceedings.

[35] Ken Perlin and Eric M. Hoffert. Hypertexture. *Computer Graphics*, 23(3):253–262, July 1989. ACM Siggraph '89 Conference Proceedings.

[36] Bertrand Le Saec and Christophe Schlick. A progressive ray-tracing-based radiosity with general reflectance functions. In *Proceedings of the Eurographics Workshop on Photosimulation, Realism and Physics in Computer Graphics*, pages 103–116, June 1990.

[37] Y. A. Screider. *The Monte Carlo Method*. Pergamon Press, New York, N.Y., 1966.

[38] Peter Shirley. *Physically Based Lighting Calculations for Computer Graphics*. PhD thesis, University of Illinois at Urbana-Champaign, November 1990.

[39] Peter Shirley. Physically based lighting calculations for computer graphics: A modern perspective. In *Proceedings of the Eurographics Workshop on Photosimulation, Realism and Physics in Computer Graphics*, pages 67–81, June 1990.

[40] Peter Shirley. A ray tracing algorithm for global illumination. *Graphics Interface '90*, May 1990.

[41] Peter Shirley. Discrepancy: A goodness measure for sampling distributions. In *Submitted to Eurographics '91*, 1991.

[42] Robert Siegel and John R. Howell. *Thermal Radiation Heat Transfer*. McGraw-Hill, New York, N.Y., 1981.

[43] Kelvin Sung. Dda traversal of octree structures. In *Submitted to Eurographics '91*, 1991.

[44] Steve Upstill. *The Renderman Companion*. Addison-Wesley, Reading, MA, 1990.

[45] John R. Wallace, Kells A. Elmquist, and Eric A. Haines. A ray tracing algorithm for progressive radiosity. *Computer Graphics*, 23(3):335–344, July 1989. ACM Siggraph '89 Conference Proceedings.

[46] Gregory J. Ward, Francis M. Rubinstein, and Robert D. Clear. A ray tracing solution for diffuse interreflection. *Computer Graphics*, 22(4):85–92, August 1988. ACM Siggraph '88 Conference Proceedings.

[47] Turner Whitted. An improved illumination model for shaded display. *Communications of the ACM*, 23(6):343–349, June 1980.