Bidirectional Reflection Functions from Surface Bump Maps

BRIAN CABRAL, NELSON MAX, and REBECCA SPRINGMEYER
UC Davis/Lawrence Livermore National Laboratory

The Torrance-Sparrow model for calculating bidirectional reflection functions contains a geometrical attenuation factor to account for shadowing and occlusions in a hypothetical distribution of grooves on a rough surface. Using an efficient table-based method for determining the shadows and occlusions, we calculate the geometric attenuation factor for surfaces defined by a specific table of bump heights.

Diffuse and glossy specular reflection of the environment can be handled in a unified manner by using an integral of the bidirectional reflection function times the environmental illumination, over the hemisphere of solid angle above a surface. We present a method of estimating the integral, by expanding the bidirectional reflection coefficient in spherical harmonics, and show how the coefficients in this expansion can be determined efficiently by reorganizing our geometric attenuation calculation.

Categories and Subject Descriptors: I.3.7 [**Computer Graphics**]: Three-Dimensional Graphics and Realism — Color, shading, shadowing, and texture

General Terms: Computer graphics, image synthesis, reflectance, shading

Additional Key Words and Phrases: Bidirectional reflection function, lighting model, horizon mapping, spherical harmonics, environmental illumination

Introduction

It is possible to model the reflection properties of a rough surface with geometric optics, by assuming that at a small scale, which is still large with respect to the wavelength of light, the surface is actually continuous, smooth, and shiny, and reflects like a curved mirror, according to Fresnel's law. The shape of a patch of such a surface can be specified as a height function of two continuous parameters. If a two parameter family of parallel rays meets such a surface, they will be reflected with a certain density per solid angle, called the *bidirectional reflection density function*, or *bidirectional reflection function* for short. This function could be estimated stochastically, by tracing many parallel, randomly positioned rays.

Authors's address:
LLNL, P.O. Box 808, Livermore, CA 94550
(415) 423-0201 cabral@lll-crg.ARPA

 ACM-0-89791-227-6/87/007/0273 $00.75

In this paper, we approximate the bumpy surface by a collection of triangular facets, interpolating between the bump heights at a discrete array of points. The rays which hit a facet without being shadowed on the way in, and leave the surface directly without hitting another facet on the way out, define a parallel beam emerging from the non-obscured part of the facet. This reflected beam contributes a "delta function" peak to the bidirectional reflection, weighted by its cross-sectional area.

The bidirectional reflection function for the faceted approximation is the sum of these weighted delta functions, and is thus not continuous. We present two methods for capturing the information in the reflected beams. The first method divides the hemisphere of reflected directions into a number of bins, and assigns each reflected beam to one of the bins. This leads to a tabulated version of the bidirectional reflection function. The second method uses spherical harmonics to approximate the density per solid angle of the weighted delta functions. The result is a series of coefficients defining a continuous bidirectional reflection function as a linear combination of basis functions, effectively condensing the table into less data, and also smoothing it.

Section 1 gives several definitions from radiometry, leading up to a physical definition of the bidirectional reflection function. Section 2 summarizes previous models for rough surfaces and methods for finding their reflection properties. Section 3 summarizes previous work on using bidirection reflection functions to calculate the effects of environmental illumination. Section 4 describes horizon mapping, an efficient way for finding the non-obscured part of each facet. Section 5 shows how to assign the beams into bins corresponding to a collection of tabulated reflection directions, and section 6 argues that a bidirectional reflection function tabulated in this way should approximately satisfy Helmholtz's Law of Reciprocity. Section 7 shows how a spherical harmonic expansion of the bidirectional reflection function can be used to efficiently compute the effects of environmental illumination, and section 8 shows how to determine the coefficients in this expansion from the collection of beams reflected by the facets. Finally, section 9 explains the construction of the bump tables, and presents images rendered with the resulting reflection functions.

Our methods generate anisotropic reflection functions. However the anisotropy can be averaged out if desired for isotropic surfaces. For simplicity we consider only isotropic reflection functions in sections 7 and 8.

1. Radiometry

There are two methods of quantifying light measurements: 1) by radiant energy, which measures physical quantities, and 2) by luminous energy, which measures the response of the eye. In this paper we discuss only radiant energy, integrated over all wavelengths. However, the quantities defined below can also be considered as spectral quantities, which vary with wavelength. They can then be converted to luminous quantities by multiplying by the CIE luminous efficiency function and integrating over all wavelengths. They can similarly be multiplied by

Table of Notation

θ	polar angle in spherical coordinates
ϕ	azimuthal angle in spherical coordinates
Φ	radiant flux
D	radiant flux density
D_{in}	radiant flux density of incident beam
L	radiance
L_{in}	incident radiance
L_{ref}	reflected radiance
f_r	bidirectional reflection function
Q	unit sphere
H	unit hemisphere of directions above a surface
N	global normal to a surface
N'	local normal to bumps on a surface
W	incident direction (θ_{in}, ϕ_{in})
V	reflection and viewing direction $(\theta_{ref}, \phi_{ref})$
M	mirror reflection $(\theta_{mirror}, \phi_{mirror})$ of V
P	off specular peak direction $(\theta_{peak}, \phi_{peak})$
β	angle from surface normal to horizon
S_i	the i^{th} facet
B_i	the non-obscured area of facet S_i
N_i	the normal to facet S_i
R_i	direction of the beam reflected from S_i
G_i	"fraction" of incident flux reflected from S_i
F	Fresnel's law
δ	Kronecker delta function
V_k	direction for bin k
Φ_k	flux in bin k
Y_{ln}	real spherical harmonics basis function
P_{ln}	associated Legendre function
M_{ln}	normalization factor
$\widehat{\cos\theta}$	$max(\cos\theta, 0)$
$D^l_{n'n}$	matrix element for rotating spherical harmonic basis functions
g_n	coefficients in Fourier series for L_{ref}

the CIE color matching functions, and integrated to get CIE x, y, and z color coordinates, which can be converted to the appropriate output colors for a specific device (See Hall [15] and Cook and Torrance [8].)

To measure solid angle, we use a spherical coordinate system with the Z axis at the surface normal N. A direction V has coordinates (θ, ϕ) where θ measures the polar angle of V from the Z axis, and ϕ measures the azimuthal angle between the X axis and the projection of V in the XY plane. The element of solid angle is then $d\omega = \sin\theta d\theta d\phi$. We now state several definitions taken from ANSI Z7.1-1967 [33].

The *radiant flux*, Φ, is the time rate of flow of energy traveling through a surface in the form of electromagnetic waves. It is usually measured in watts.

The *radiant flux density*, $D = d\Phi/dA$, at an element of a surface is the ratio of the radiant flux through that surface element, to the area of the element. It can be measured in watts per square meter. When referring to radiant flux incident on a surface, the density is called *irradiance*.

The *radiance*, $L = d^2\Phi/(d\omega \cos\theta dA)$, in a direction (θ, ϕ), at a point of a surface is the following ratio. The numerator is the radiant flux leaving, passing through, or arriving at an element of the surface surrounding the point, and propagated in directions within a cone containing the given direction. The denominator is the product of the solid angle of the cone and the area of the orthogonal projection of the element of the surface on a plane perpendicular to the given direction. Radiance can be measured in watts per square meter per steradian.

Consider a beam of incident light, with radiance L_{in}, coming from a solid angle $d\omega_{in}$ about an incident direction (θ_{in}, ϕ_{in}) and reflecting from a surface. Let dL_{ref} be the radiance reflected in the direction $(\theta_{ref}, \phi_{ref})$. Then the *bidirectional reflection function* f_r is the ratio:

$$f_r(\theta_{in}, \phi_{in}; \theta_{ref}, \phi_{ref}) = \frac{dL_{ref}}{L_{in} \cos\theta_{in} d\omega_{in}}$$

(See Snell [28].) In the limit when $d\omega_{in}$ approaches zero and the incident beam becomes parallel, the denominator becomes the irradiance onto the surface from the single direction (θ_{in}, ϕ_{in}). Since f_r is a radiance divided by a radiant flux density, the units of f_r are inverse steradians.

For a radiance function $L_{in}(\theta_{in}, \phi_{in})$ which varies continuously with the incident direction, we can get the total reflected radiance in the direction $(\theta_{ref}, \phi_{ref})$ by integrating dL_{ref} over the hemisphere H of possible incident directions above a surface. Thus

$$L_{ref} = \int_H dL_{ref} = \int_H f_r(\theta_{in}, \phi_{in}; \theta_{ref}, \phi_{ref}) L_{in} \cos\theta_{in} d\omega_{in}. \quad (1)$$

2. Bidirectional Reflection Models

Torrance and Sparrow [32] developed a geometric lighting model for specular reflection, which was applied to computer graphics by Blinn [1], and modified by Cook and Torrance [8], to take better account of colored light. It has been used to create spectacularly realistic shaded renderings of surfaces such as brushed metal.

The Torrance-Sparrow model assumes that an infinitely long, symmetrical, wedge-shaped groove has been cut into the surface. The normal N' to one of the two flat sides of this groove is used, together with the lighting direction and the viewing direction, to compute the fraction of the incoming light which is neither shadowed on the way in nor obscured on the way out by the opposite side of the groove. Random grooves, with some specific distribution of normals N' about the surface normal N, generate a distribution of reflected rays from each incoming direction. One can thus compute a bidirectional reflection coefficient. However on any real surface the grooves will interfere with each other, so they cannot be both randomly oriented and infinitely long.

A real surface can be more explicitly modeled by a bump table, which gives the surface height at a two dimensional array of sample points. Blinn [2] has shown how to perturb the surface normal for use in a lighting model which takes account of the bumps. The method has become known as "bump mapping." It gives realistic light reflections for a level of detail at which the bumps are visible, but need not be modeled in full 3-D perspective. Max [21] has shown how to account for the shadows from the bumps, by using a table of horizon angles. This might appropriately be called "horizon mapping." In the present paper, we divide the surface defined by the bump table into triangular facets. For each facet and each tabulated lighting direction, we compute the reflection direction and use the horizon table to find the reflected beam, that is, the portion of the beam potentially intercepted by the facet, which is neither shadowed on the way in nor obscured on the way out.

We accumulate this reflected flux into a tabulated bidirectional reflection function, using the tabulated reflection direction closest to the actual reflected beam. Because it is derived from the bump table, the resulting lighting model will be consistent with the area average of the light from an image rendered with bump mapping and horizon mapping. This should permit smooth transitions of detail between these two methods of representing surface roughness. This is the chief justification of our method: consistency with explicitly modeled surface structure, including anisotropic reflections. Kajiya [20] has proposed smooth transitions between three roughness representations: lighting models, bump mapping, and full 3-D visible-surface calculations which include the bumps in the data base. However his method uses the wave theory of light, and is not computationally practical. If bump mapped surfaces are rendered with horizon mapping, they should in addition allow smooth transitions to full 3-D renderings including cast shadows.

Perlin [24] has also proposed a smooth transition of surface roughness components, between lighting models and normal perturbations. However his normal perturbations do not come from an explicitly specified bump table, and therefore do not permit horizon mapping, or full 3-D modeling.

Ohira [23], Takagi [30], and Kajiya [20] have discussed bidirectional reflection functions which are anisotropic. These depend on all four parameters needed to define the lighting and viewing directions, with respect to a frame specifying the orientation of two surface tangents and the surface normal. Our calculations initially generate anisotropic reflections, since the two tangent vectors are needed to orient the bump tables on the surface. If isotropic reflection functions are desired, the anisotropic tables can be averaged over the possible orientations of the bump table on the surface, so that they depend on only three parameters. This will make more reflected rays contribute to each table entry, and improve the sampling statistics.

3. Environmental Illumination

Early illumination models handled only point light sources, but there have been many attempts to include the light coming from the whole environment, based on equation (1) above. Whitted and Cook [35] analyze various specific lighting models in terms of restrictions on the full generality of equation (1). For example, Blinn and Newell [3] created mirror reflections of a room in a teapot by assuming that L_{in} is independent of position, and that f_r is a delta function. If L_{in} instead is a delta function, we have the case of a point light source. Lambert's law for diffuse reflection results when f_r is constant. Phong [26], Blinn [1], and Cook and Torrance [8] have generated glossy highlights from point sources by using more general specular reflection functions f_r, concentrated near the mirror direction. Cook, Carpenter, and Porter [7] have used distributed ray tracing to model glossy reflections, with such a concentrated f_r and an arbitrary L_{in}. Because the integral is estimated by tracing sample rays near the mirror reflection, L_{in} can vary with position and multiple reflections can be rendered.

Kajiya [19] has demonstrated that by long calculation and a little pruning of the ray tree, it is possible to include diffuse as well as specular reflections in this scheme, allowing a completely general f_r. Immel, Cohen, and Greenberg [18] have also solved the same interreflection problem for general f_r, by dividing the surface into polygons, and solving a huge system of linear equations.

In this paper we apply our bidirectional reflection distribution function to the case that L_{in} is independent of position, as if the environment were painted on a sphere at infinity. In this case $L_{in}(\theta_{in}, \phi_{in})$ can be sampled into a texture table. Greene [13] has called techniques based on such a table "environment mapping", and surveyed methods for obtaining the integral for L_{ref} efficiently from the table. For example, if f_r is constant (diffuse reflection) L_{ref} depends only on the normal N, and Miller and Hoffman [22] suggest that it be precomputed and tabulated. In the Phong reflection model [26], f_r depends only on the mirror reflection $M = (\theta_{mirror}, \phi_{mirror})$ of the viewing direction $V = (\theta_{ref}, \phi_{ref})$, so L_{ref} can again be precomputed and tabulated as a function of $(\theta_{mirror}, \phi_{mirror})$ (See [22].)

Consider a somewhat more general f_r, which, as a function of (θ_{in}, ϕ_{in}) for fixed V and fixed normal N, has elliptical symmetry about some "off specular" peak direction $P = (\theta_{peak}, \phi_{peak})$ near M. (See [1] or [32].) That is, contours of constant f_r correspond to small concentric ellipses about P on the sphere of (θ_{in}, ϕ_{in}) directions. (See figure 1.) The shape of f_r is determined by the eccentricity, which is the same for all these ellipses, and by a curve on a cross section plane through N and P, specifying f_r as a function of the ellipse size. The shape of f_r usually varies with θ_{ref}, and the orientation of the ellipses varies with N, so the integrals cannot be precomputed as a function of $(\theta_{peak}, \phi_{peak})$ alone.

The elliptical weighted average filter of Greene and Heckbert [14] can take advantage of the elliptical symmetry of f_r to estimate the integral for L_{ref} from the tabulated values of $L_{in}(\theta_{in}, \phi_{in})$. Its computation time is proportional to the number of entries in the table which contribute to the integral, so this method is most efficient for shiny surfaces, with narrow specular peaks. Greene [13] has suggested using Williams' pyramidal parametrics [36] together with elliptical weighted averages, so that the number of entries in the selected table is bounded, independent of the width of specular peak. Glassner [12] has shown how to use summed area tables (See Crow [10]) to approximate the integral of L_{in} over larger ellipses. But his method requires that f_r be constant inside the ellipse, and zero outside.

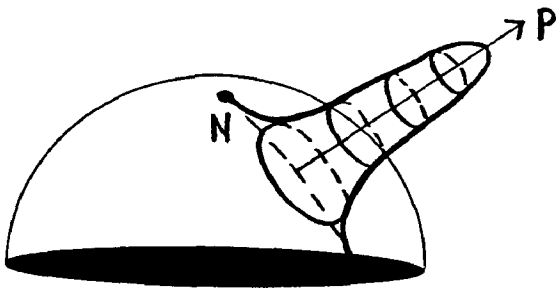

Figure 1.

Perlin [25] has generalized summed area tables to higher order, and shown how to approximate an elliptical gaussian by piecewise quadratic polynomials, whose convolutions with L_{in} can be obtained using the second order tables, in time independent of the width of the gaussian. (See Heckbert [16] for a more detailed explanation.) The ellipses must be oriented with their major and minor axes horizontal or vertical in the texture table. This can be arranged in the following case. Suppose the view of the reflecting object is from far away, so that the viewing direction V can be assumed to be constant, and that L_{in} is tabulated in spherical coordinates, with V as the north pole. An isotropic bidirectional function will have mirror symmetry with respect to a plane through V and N, so its contour ellipses must also. This guarantees the required orientation, since a plane through the pole is a line of constant azimuthal angle in the texture table.

Greene [14] also considers the problem of antialiasing reflections in a curved, perfectly shiny mirror, to account for the fact that the ray from the eye reflects in directions which vary over the area of the pixel. He proposes elliptical weighted average filtering, using values of L_{in} tabulated on the six faces of an environment cube. A gaussian weighting about the pixel center, for slowly curving mirrors, will transform to an approximate elliptical gaussian in the environment map, in spherical coordinates as well as on an environment cube. To generalize this to the case of a glossy curved mirror, the appropriate filter is the convolution of the weighting function with the specular bidirectional reflection function f_r. If f_r can be approximated by an elliptical gaussian, the convolution of these two functions is again an elliptical gaussian, which may no longer be oriented consistently with the spherical axes. The method of Perlin [24] no longer applies, but the method of Greene and Heckbert [13,14] can still be used. Miller and Hoffman [22] also include antialiasing in their technique. Of course, in distributed ray tracing [2,7], the rays can be distributed simultaneously over both the subpixel positions and the glossy reflection directions.

The method of integration we propose in section 7 below works under restricting assumptions similar to those above: a) L_{in} is independent of position, b) V is constant, and c) f_r is isotropic. However, we can handle arbitrarily shaped f_r, which can combine specular and diffuse reflection. We expand both $\cos\theta_{in} f_r$ and L_{in} in spherical harmonics, with V as the north pole. The integral for L_{ref} can then be computed efficiently as the dot product of the coefficient vectors for $\cos\theta_{in} f_r$ and L_{in}. Because the computation time is proportional to the number of coefficients, our method works best for rougher surfaces, where f_r is

wide and can be approximated by fewer terms. (This is in contrast to the methods of Greene and Heckbert [14] and Immel, Cohen, and Greenberg [18], which work better for narrow f_r.)

4. Horizon Mapping

Horizon tables were originally developed to create shadows on bump-mapped surfaces [21]. Imagine a bump height function, tabulated on an $m \times m$ grid, representing the altitude of a mountainous terrain, with $x = east$, $y = north$, and $z = up$. Suppose the sun sets to the west in a vertical plane, as it would at the equator on an equinox. Then a point on the terrain is in shadow whenever the sun is below the horizon, that is, whenever the angle θ from the Z axis to the sun is greater than the angle β from the Z axis to the western horizon. The angle to the western horizon can be computed at each tabulated position, by looking at the slopes of lines from that position to other tabulated positions to its west, and finding the maximum slope. The bump map is doubly periodic in x and y, so that the slope can be determined by the $m-1$ points to the west, even for a starting point near the western edge of the bump map. The angle of the horizon can be determined similarly for any sunset direction, that is, for any half plane through the Z axis. Max [21] proposed using only the eight principal compass directions, for which the slope to the horizon can be easily determined from the bump table. This resulted in a horizon table of size $m \times m \times 8$, and the horizon angle at other compass directions was determined by interpolation. Mountains which are missed by the sampled compass directions from a data point will fail to cast shadows on the point, so more compass directions will give more accurate results.

In this work, we have used 24 directions. Let PQ represent a ray from the data point P in one of these directions, shown intersecting the triangles in Figure 2. The bump heights at an intersection point with a triangle side are determined by interpolation from the tabulated vertices at endpoints of the side. Note that the interpolation factors, and the positions of the two endpoints relative to the start of the ray, depend only on the ray's direction. To speed up the calculation of the horizon table, this information is computed ahead of time for each ray direction [5].

Max [21] shows how to use the horizon table computed for a flat surface, to find the bump shadows when the flat surface is parametrized onto a curved surface patch. Here we will use it to estimate energy reflected from the bumpy surface.

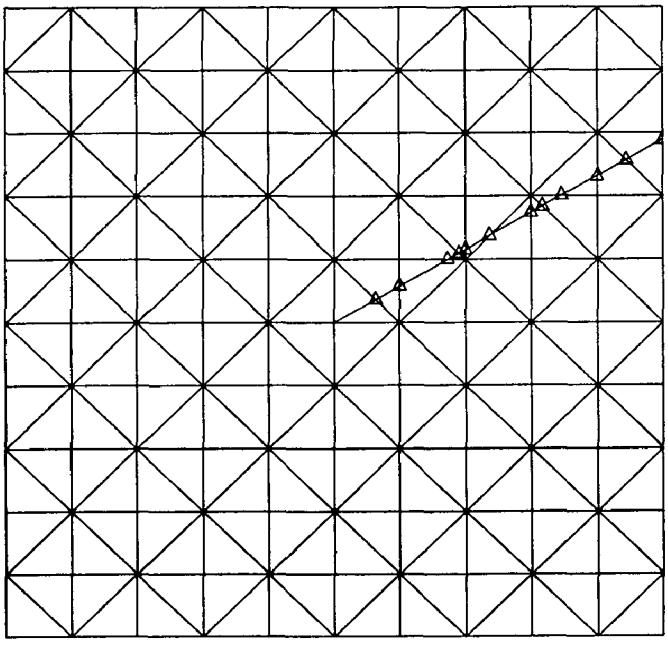

Figure 2. Ray PQ intersecting grid tessellations.

5. Reflection From Facets

Consider a collimated parallel beam of light coming to the surface S from the direction $W_{in} = (\theta_{in}, \phi_{in})$, with radiant flux density D_{in} measured per unit area normal to W_{in}. This is the limit for a source of radiance L_{in} and solid angle $d\omega_{in}$, as the source moves away to infinity, and $d\omega_{in}$ approaches zero, while the radiance increases to keep $L_{in}d\omega_{in}$ equal to D_{in}. Since W_{in} is tilted away from the normal N by an angle θ_{in}, the irradiance incident on S is $D_{in}\cos\theta_{in} = D_{in}(W_{in} \cdot N)$.

We divide S into $2m^2$ triangular facets, as shown projected in figure 2. If U_i and V_i represent the vectors along the edges of facet S_i, which project to the horizontal and vertical edges in figure 2, then the area A_i of S_i is $|U_i \times V_i|/2$, and the normal N_i to S_i is $(U_i \times V_i)/|U_i \times V_i|$. Given N_i and W_{in} we can compute the reflected direction $R_i = 2(W_{in} \cdot N_i)N_i - W_{in}$ (See Whitted [34]).

The obscuring effects of other facets can be estimated from the horizon table. We can determine from the entries for the compass direction θ_{in}, the three horizon angles β_1, β_2, and β_3 for the three vertices of S_i. If the samples for θ_{in} for the tabulated bidirectional reflection function correspond to those in the horizon table, no interpolation is required. When β_1, β_2, and β_3 are all greater than θ_{in}, the facet S_i is completely lit, and when all are less than θ_{in}, the whole facet is in shadow. When some are greater and some are less, the partial shadowing is determined as follows. Along the two sides where $\beta_j - \theta_{in}$ changes sign, we interpolate to estimate two points where $\beta = \theta_{in}$. The line joining these two points separates the triangle into a quadrilateral and a smaller triangle, one of which is illuminated and the other, shadowed.

Shadowed light contributes its flux instead to the specular reflection from the facet which intercepted it. But, as in the Torrance-Sparrow model, reflected light which is obscured on the way out is lost from the specular energy computation, and assumed to contribute instead to diffuse reflection or absorption.

The illuminated portion of the facet sends out a reflected beam in the direction $R_i = (\theta_{ref}, \phi_{ref})$. The portion of this beam which is not obscured can be found similarly, using the convention of Duff [11], Figure 1, to treat the case where $\beta - \theta_{ref}$ changes sign along all four sides of a quadrilateral. Since ϕ_{ref} can be an arbitrary azimuthal angle, interpolation in the horizon table is now required, which introduces another approximation in our calculations.

Let B_i be the area of the facet S_i which reflects the non-obscured, non-shadowed portion of the beam. This area projects to an area $B_i(N_i, W_{in})$ normal to the incident beam, and therefore intercepts a flux of $D_{in}B_i(N_i, W_{in})$. We assume the facet is perfectly smooth, so that the reflected flux is this incident flux multiplied by the Fresnel factor $F(N_i, W_{in})$. (See Born and Wolf [4].)

Now take a collection $\{V_k | k = 1, n\}$ of sampled reflection directions, and let $Nearest(V)$ be the index k of the V_k nearest to the unit vector V. (Choose the lower index in case of a tie.) Then $C_k = \{V \in D | Nearest(V) = k\}$ represents the bin on the hemisphere H assigned to V_k. Let $d\omega_k$ be the solid angle measure of C_k.

To accumulate the reflected flux into the appropriate bins, we assign the flux from facet S_i to bin $k = Nearest(R_i)$. When all $2m^2$ facets have been considered the total flux $\Phi_k(\theta_{in}, \phi_{in})$ in bin k is

$$\Phi_k(\theta_{in}, \phi_{in}) = \sum_{i=1}^{2m^2} \delta(k - Nearest(R_i))F(N_i, W_{in})D_{in}B_i(W_{in} \cdot N_i) \quad (2).$$

We find the reflected radiance $L_{ref,k}$ in the direction V_k for bin k, by dividing the flux Φ_k by the area dB of S normal to the outgoing direction V_k, and by the solid angle $d\omega_k$ of the bin. The area dA of S is taken as the area of the flat lattice, shown in figure 2, on which the bumps are based. Then $dB = dA(V_k \cdot N)$. Thus

$$L_{ref,k} = \frac{\Phi(\theta_{in}, \phi_{in})}{d\omega_k dA(V_k \cdot N)}$$

By definition, the bidirectional reflection function f_r is this reflected radiance, divided by the incoming irradiance. Since the incoming beam has a radiant flux density D_{in} measured normal to its direction W_{in}, its irradiance on the surface S is $D_{in}\cos\theta_{in} = D_{in}(W_{in} \cdot N)$, so

$$f_r(\theta_{in}, \phi_{in}; \theta_{ref}, \phi_{ref}) = \frac{L_{ref,k}}{D_{in}(W_{in} \cdot N)}$$

Putting the above three equations together, we get

$$f_r(\theta_{in}, \phi_{in}; \theta_{ref}, \phi_{ref}) = \frac{\sum_{i=1}^{2m^2} \delta(k - Nearest(R_i))F(N_i, W_{in})D_{in}B_i(W_{in} \cdot N_i)}{D_{in}(W_{in} \cdot N)d\omega_k dA(V_k \cdot N)}. \quad (3)$$

Note that D_{in} can be cancelled from both the numerator and denominator, so that $G_i = F(N_i, W_{in})B_i(W_{in} \cdot N_i)/(W_{in} \cdot N)dA$ represents the fraction of the incoming flux reflected by facet S_i, and

$$f_r(\theta_{in}, \phi_{in}; \theta_{ref}, \phi_{ref}) = \frac{\sum_{i=1}^{2m^2} \delta(k - Nearest(R_i))G_i}{d\omega_k(V_k \cdot N)}. \quad (4)$$

If this computation is repeated for each of a finite collection of incident directions W_{in}, we can build up a bidirectional reflection table $f_r(\theta_{in}, \phi_{in}; \theta_{ref}, \phi_{ref}) = f_r(W_j, V_k)$.

Torrance and Sparrow [32] consider a continuous distribution of hypothetical facets, all of the same projected area, with a density $\delta(N')$ per steradian having facet normal N'. They apply the formula $d\omega_{ref} = 4(W \cdot N')d\omega_{norm}$ from Rense [27], where $d\omega_{norm}$ is a small solid angle for normals about N', and $d\omega_{ref}$ is the solid angle of corresponding reflected directions for the incident ray from W. Because we are using an explicit collection of polygonal facets instead of a distribution, and accumulating the flux of the reflected beams into a predefined collection of bins, we do not need to include this solid angle spreading factor.

6. Reciprocity

Torrance and Sparrow also refer to the reciprocity law $f_r(V, W) = f_r(W, V)$ in Helmholtz [17], obtained by reversing the roles of the incident and reflected directions. Our calculated f_r should approximately satisfy this reciprocity law by the following reasoning.

Take the same collection of samples on the hemisphere H for both the viewing, V, and incident, W, direction indices in the tabulated bidirectional reflection. Assume that

$$C_k = \{V | Nearest(V) = k\} = \{V | Angle(V, V_k) < \gamma\}$$

for some small constant γ, so that C_k is a circular region on H, of radius γ. (This can never be exactly true, but is approximately true for a "hexagonally close packed" collection of directions based on icosahedral symmetry. The corresponding C_k's consist of 12 regular pentagonal regions and $30l^2 - 10$ regular or not-quite-regular hexagonal regions on the unit sphere. For example, a soccer ball is sewn from 12 black pentagonal and 20 white hexagonal pieces of leather, the case $l = 1$.)

Now consider a facet S_i of non-obscured area B_i, for incoming light direction W_j, and reflected beam direction R_i. Let $Nearest(R_i) = k$, so that $Angle(R_i, V_k) < \gamma$. Then the same facet will reflect V_k back to a ray within γ of W_j. The non-obscured area in the reverse situation will be close to B_i if γ is small, since all angles have changed by less than γ. Also, the angle between R_i and V_k is less than γ, so in equation (2) for incoming direction W_j, $W_j \cdot N_i = R_i \cdot N_i \approx V_k \cdot N_i$, and the Fresnel factors $F(N_i, W_j)$ and $F(N_i, V_k)$ are also close, since $F(N, W)$ depends only on the dot product of N and W. Thus the contributions of facet S_i to the sum $\Phi_k(W_j)$ in equation (2), and to the corresponding ray reversed sum $\Phi_j(V_k)$, are approximately equal. In the denominator of equation (3), we have assumed the solid angles $d\omega_k$ are all equal, and $V_k \cdot N$ and $W_j \cdot N$ both appear, so $f_r(W_j, V_k)$ and $f_r(V_k, W_j)$ are approximately equal also. We can thus decrease the size of the bidirectional reflection table by about half, and improve the sampling statistics, by setting both $f_r(W_j, V_k)$ and $f_r(V_k, W_j)$ to their average value, enforcing reciprocity.

Under the additional condition that the reflection is isotropic, $f_r(\theta_{in}, \phi_{in}; \theta_{ref}, \phi_{ref})$ depends only on $\theta_{in}, \theta_{ref}$, and $|\phi_{in} - \phi_{ref}|$. Thus, the data for all bins with the same $|\phi_{in} - \phi_{ref}|$ can be averaged, reducing the table size and improving the statistics still further. In the following two sections, we assume f_r is isotropic.

7. Spherical Harmonics in Environment Mapping

The bidirectional reflection function can be used as a part of any shading algorithm. In this section, we discuss the use of spherical harmonics in shading computations based on environmental illumination. The spherical harmonic method can be used with any bidirectional reflection function. In section 8, however, we give a particularly efficient method of determining the spherical harmonic coefficients from the reflected beams.

For representing functions defined on the unit sphere, spherical harmonics functions are the analogues of the fourier series terms used on the unit circle. They are the products of associated Legendre functions with simple periodic functions in ϕ. When used to represent wave functions in physics, they are usually complex valued, but to represent images, we need real functions. The real spherical harmonics may be expressed in terms of $\cos n\phi$ and $\sin n\phi$:

$$Y_{ln}(\theta, \phi) = \begin{cases} M_{ln} P_{ln}(\cos\theta)\cos n\phi & n \geq 0; \\ M_{ln} P_{l|n|}(\cos\theta)\sin |n|\phi & n < 0; \end{cases}$$

where

$$M_{ln} = \begin{cases} \sqrt{\frac{(2l+1)}{2\pi}\frac{(l-|n|)!}{(l+|n|)!}}, & n \neq 0; \\ \sqrt{\frac{(2l+1)}{4\pi}}, & n = 0; \end{cases}$$

and $P_{ln}(\cos\theta)$ is the associated Legendre function. (See Courant and Hilbert [9].) The spherical harmonics form an orthonormal basis for the Hilbert space of functions on the unit sphere. Thus any real function $h(\theta, \phi)$ has an expansion

$$h(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{n=-l}^{l} a_{ln} Y_{ln}(\theta, \phi).$$

We can speed up the calculation of the integral in equation (1) by approximating the factor $\cos\theta_{in} f_r(W, V)$ with spherical harmonics. We reorganize the spherical coordinates so that the pole is in the viewing direction V. The isotropic bidirectional reflection function $f_r(W, V)$ really depends on the three vectors W, V, and N. (For non-isotropic reflection functions, tangent vectors are also required.) In the previous section, N was assumed fixed, but now V is fixed. So we change notation, and let $f_{r_N}(\theta, \phi)$ denote the bidirectional reflection function of the illumination direction $W = (\theta, \phi)$, for V permanently fixed, and N temporarily fixed and indicated in the subscript. We have taken V along the Z axis, and will at first assume that N lies in the XZ plane, with spherical coordinates $(\alpha, 0)$. Since f_r is isotropic, $f_{r_N}(\theta, \phi)$ is symmetric with respect to the XZ plane, and $f_{r_N}(\theta, -\phi) = f_{r_N}(\theta, \phi)$. As a result, $\cos\theta_{in} f_{r_N}(\theta, \phi)$ can be expanded in spherical harmonics with terms in $\cos n\phi$ but none in $\sin n\phi$, that is, there are no a_{ln} terms for $n < 0$. This expansion is defined over the whole unit sphere Q, instead of just the hemisphere H, so we will replace $\cos\theta_{in}$ by $\widehat{\cos}\theta_{in} = max(\cos\theta_{in}, 0)$. Thus the expansion of $\widehat{\cos}\theta_{in} f_{r_N}(\theta, \phi)$ is:

$$\widehat{\cos}\theta_{in} f_{r_N}(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{n=0}^{l} a_{ln}(\alpha) Y_{ln}(\theta, \phi).$$

The expansion is an infinite series, but if we replace the infinite upper bound in the first summation by a finite upper bound M, we get an approximation to $\widehat{\cos}\theta_{in} f_{r_N}(\theta,\phi)$. We will calculate the coefficients $a_{ln}(\alpha_j)$ for a finite sample $\{N_j | j = 1, J\}$ of normals in the XZ plane, with spherical coordinates $(\alpha_j, 0)$, and find $a_{ln}(\alpha)$ for other normals $N = (\alpha, 0)$ by interpolation.

Now let N be an unrestricted normal, with spherical coordinates (α,β). A rotation of $-\beta$ about the Z axis takes (α,β) to $N_0 = (\alpha, 0)$, and $W = (\theta,\phi)$ to $W_0 = (\theta, \phi-\beta)$. Since the reflection function is isotropic, and only depends on the relative angles between V, N, and W, $f_{r_N}(\theta,\phi) = f_{r_{N_0}}(\theta, \phi-\beta)$. Also, $\widehat{\cos}\theta_{in}$ depends only on the angle between W and N. Therefore

$$\widehat{\cos}\theta_{in} f_{r_N}(\theta,\phi) \approx \sum_{l=0}^{M}\sum_{n=0}^{l} a_{ln}(\alpha) M_{ln} P_{ln}(\cos\theta)\cos(n(\phi-\beta)).$$

But

$$\begin{aligned}\cos n(\phi-\beta) &= \cos(n\phi - n\beta)\\ &= \cos n\phi\cos n\beta + \sin n\phi \sin n\beta.\end{aligned}$$

Therefore

$$\widehat{\cos}\theta_{in} f_{r_N}(\theta,\phi) \approx \sum_{l=0}^{M}\sum_{n=-l}^{l} \tilde{a}_{ln} M_{ln} P_{l|n|}(\cos\theta) \begin{cases}\cos n\phi, & n \geq 0;\\ \sin|n|\phi, & n < 0;\end{cases}$$

where

$$\tilde{a}_{ln} = \begin{cases} a_{ln}(\alpha)\cos n\beta, & n \geq 0;\\ a_{l|n|}(\alpha)\sin|n|\beta, & n < 0.\end{cases}$$

We now have terms in $\sin n\phi$, so there are $\tilde{a}_{ln}$ for $n < 0$. The computation of the $a_{ln}(\alpha)$ is explained in section 4.

The radiance of the environment, $L_{in}(\theta,\phi)$, may also be expanded in real spherical harmonics:

$$L_{in}(\theta,\phi) = \sum_{l=0}^{\infty}\sum_{n=-l}^{l} b_{ln} Y_{ln}(\theta,\phi).$$

Since the Y_{ln} are an orthonormal basis, the coefficients b_{ln} are computed by integrating the product of the illumination function and the spherical harmonics over the unit sphere:

$$b_{ln} = \int_Q L_{in}(\theta,\phi) Y_{ln}(\theta,\phi) d\omega \qquad (5).$$

Having expanded both the reflection function and the environmental radiance function in spherical harmonics, we can now estimate the integral of equation (1):

$$\begin{aligned}L_{ref} &= \int_H \cos\theta_{in} f_r(V,W) L_{in}(\theta,\phi) d\omega\\ &= \int_Q \widehat{\cos}\theta_{in} f_r(V,W) L_{in}(\theta,\phi) d\omega\\ &\approx \int_Q \sum_{l=0}^{M}\sum_{n=-l}^{l} \tilde{a}_{ln} Y_{ln}(\theta,\phi) \sum_{l'=0}^{\infty}\sum_{n'=-l'}^{l'} b_{l'n'} Y_{l'n'}(\theta,\phi) d\omega\\ &= \sum_{l=0}^{M}\sum_{n=-l}^{l}\sum_{l'=0}^{\infty}\sum_{n'=-l'}^{l'} \tilde{a}_{ln} b_{l'n'} \int_Q Y_{ln}(\theta,\phi) Y_{l'n'}(\theta,\phi) d\omega\\ &= \sum_{l=0}^{M}\sum_{n=-l}^{l}\sum_{l'=0}^{\infty}\sum_{n'=-l'}^{l'} \tilde{a}_{ln} b_{l'n'} \quad \delta(l,l')\delta(n,n')\\ &= \sum_{l=0}^{M}\sum_{n=-l}^{l} \tilde{a}_{ln} b_{ln}.\end{aligned}$$

Thus L_{ref} becomes the inner product of these two coefficient vectors. If L_{in} has high spatial frequencies, but f_r does not, then f_r acts like a filter in the integral for L_{ref}, averaging out the high frequencies in L_{in}. Thus when the terms up to $l = M$ are sufficient to represent $\widehat{\cos}\theta_{in} f_{r_N}(\theta,\phi)$, the higher terms in L_{in} are eliminated from the integral by the orthogonality of the Y_{ln}.

For a curved or bump-mapped surface, the bidirectional reflection function should be further widened and frequency limited, to account for the variation of surface normals within a pixel, as mentioned in the introduction above. An inexpensive way to do this is to limit the range of l indices in the expansion for $\widehat{\cos}\theta f_{r_N}$ to some maximum M' less than M. The values of a_{ln} can be gradually decreased for large l, so that the expansion varies smoothly with the range of normals within a pixel. Decreasing the weight for terms with large l also helps eliminate the "ringing", or Gibbs phenomenon, at sharp edges or in $L_{in}(\theta,\phi)$.

The M-term approximation for $\widehat{\cos}\theta_{in} f_{r_N}(\theta,\phi)$ cannot be exactly zero in the hemisphere $Q - H$ behind the surface, where $\widehat{\cos}\theta_{in}$ is zero. This can cause slight reflections of objects behind the surface, but such artifacts do not appear disturbing in the images shown here.

If $L_{in}(\theta,\phi)$ is represented as a cube whose six faces are texture maps, as in Greene [13], the integral in equation (5) can be estimated by summation over the texture pixels. This integral must be computed only once, for the first viewing direction V. For each new viewing direction the set of coefficients b_{ln} may be transformed by taking a linear combination of the original coefficients:

$$\tilde{b}_{ln} = \sum_{n'=-l}^{l} b_{ln'} D^l_{n'n}(\alpha,\beta,\gamma)$$

where the matrices $D^l_{n'n}$ describe how the spherical harmonics transform under rotation through the Eulerian angles α,β, and γ. (See Tinkham [31].) Note that each D matrix will be of size $2l+1$, with the largest matrix requiring $(2M+1)\times(2M+1)$ elements. The use of these rotation matrices is clearly more efficient than recomputing the coefficients for each new viewing direction.

The expansion

$$L_{ref} \approx \sum_{l=0}^{M}\sum_{n=-l}^{l} \tilde{a}_{ln}(\alpha) b_{ln} \qquad (6)$$

has $(M+1)^2$ terms, each each requiring interpolation and rotation, and it must be computed once per pixel. This is better than doing an integral per pixel, but it can be simplified still further by precomputing some of the interpolation, as follows.

Once per surface texture, for $j = 1, ..., J$, for $l = 0, ..., M$, and for $n = 0, ..., l$, compute

$$a_{ln}(\alpha_j) = \int_Q \widehat{\cos}\theta_{in} f_{r_{N_j}}(\theta,\phi) Y_{ln}(\theta,\phi)\sin\theta d\theta d\phi \qquad (7)$$

for $N_j = (\alpha_j, 0)$. Once per viewing direction/surface-texture combination, and for $j = 1, ...J$, compute

$$g_n(\alpha_j) = \sum_{l=|n|}^{M} a_{l,n}(\alpha_j) b_{l,n}. \qquad (8)$$

There may be only one constant viewing direction, in the approximation of distant viewing discussed above. But suppose the scene contains several objects, each subtending a small angle of view, even if the whole image covers a large angle. Then there may be a separate viewing direction for each object. There may even be a separate environment for each object, if the hidden surface computation for the environment texture cube is recomputed from the point of view of each object. This allows one level of glossy interreflection between objects to be approximated. There may also be different viewing directions for different frames of an animation.

Now suppose we are at a specific pixel on a specific surface, where the normal N has spherical coordinates (α, β). For each $n = -M, ..., M$, we interpolate $g_n(\alpha)$ from the $g_n(\alpha_j)$ for the appropriate surface texture. In order not to introduce derivative discontinuities in the shading, which could be visible as Mach bands, some form of C^1 interpolation, such as Catmull-Rom splines [6], should be used. We can precompute ahead of time the coefficients of the interpolation polynomials for each interval $\alpha_j \leq \alpha \leq \alpha_j + 1$, for $n = 0, ...M$, and for each surface texture. Once the $g_n(\alpha)$ are found, the radiance L_{ref} becomes

$$L_{ref} = \sum_{n=0}^{M} (g_n(\alpha) \cos n\beta + g_{-n}(\alpha) \sin n\beta)$$

which results when equation (8) and the definitions for $\tilde{a}_{ln}$ are combined with equation (6). This equation represents L_{ref} as a Fourier series in β for fixed α, which is a useful way of approximating a smooth function on a circle. It is only necessary to evaluate $\cos\beta$ and $\sin\beta$, as $\cos n\beta$ and $\sin n\beta$ can be obtained from these inductively, with two multiplies and one add each. The cubic interpolations take three multiplies and three adds each. Therefore, the total cost per pixel is two trigonometric function evaluations, about 12M multiplies, and about 9M adds. On a vector computer like the Cray-1, it is easier to do 2M trig evaluations, 8M multiplies and 7M adds, because everything can then be vectorized.

8. Computing the coefficients $a_{ln}(\alpha_j)$

The bidirectional reflection function enters into the shading computations of the previous section, in the integral of equation (7) for the coefficients $a_{ln}(\alpha_j)$. We wish to apply the method of section 5 for the bidirectional reflection function from an explicit surface model. We could estimate this integral by assuming $f_{r_N}(\theta, \phi)$ was constant over each of the bins in the table. (Note that these bins are in spherical coordinates with N as axis, not V.) But it is possible to get a better estimate, taking advantage of the exact flux and direction of the beam reflected from each facet.

For each of the sampled normals $N_j = (\alpha_j, 0)$ in the XZ plane, place the bump texture in an orientation perpendicular to N_j. The texture can be rotated about N_j, and the calculations below can be summed over all the rotated positions, to get better statistics for an isotropic bidirectional reflection function. Neglecting the $\sin n\phi$ terms in the expansion automatically enforces averaging over the reflection of the texture in the XZ plane, by projecting onto the subspace of functions f_r with $f_r(\theta, \phi) = f_r(\theta, -\phi)$.

By reciprocity, we can trace the beams from the fixed viewing direction V, and accumulate the flux at the different lighting directions W. Recall from Section 2 that B_i represents the non-obscured area of the i^{th} facet, N_i its normal, and $F(N_i, V)$ the Fresnel law, so that $G_i = B_i F(N_i, V)(N_i \cdot V)/(V \cdot N)dA$ is the fraction of the incoming flux reflected by the facet, in the direction $W_i = 2(V \cdot N_i)N_i - V$. Let $h_{W_i}(W)$ be a smoothing function of W which is maximum at W_i, zero far away from W_i and outside the hemisphere H, and has integral over the unit sphere equal to one. Then $G(W) = \sum_{i=1}^{2m^2} h_{W_i}(W) G_i$ is an approximation to the flux fraction per solid angle reflected from V to the unit sphere Q, by the faceted surface, smoothed so as to be integrable.

In Equation (4), we divided this flux fraction per solid angle by $V_k \cdot N$, which corresponds to $W_i \cdot N$ here, to get the bidirectional reflection function. We need not divide now by $W_i \cdot N = \cos\theta_{in} = \widehat{\cos}\theta_{in}$, since we want to estimate $\widehat{\cos}\theta_{in} f_{r_N}$ in the integral (7). Thus

$$\begin{aligned} a_{ln}(\alpha_j) &= \int_Q \widehat{\cos}\theta_{in} f_r(\theta, \phi) Y_{ln}(\theta, \phi) d\omega \\ &\approx \int_Q G(W) Y_{ln}(\theta, \phi) d\omega \\ &= \sum_{i=1}^{2m^2} G_i \int_Q h_{W_i}(W) Y_{ln}(\theta, \phi) d\omega \end{aligned}$$

Denote the spherical coordinates of W_i by (θ_i, ϕ_i), and take the limit as $h_{W_i}(W)$ approaches the delta function. We get

$$a_{ln}(\alpha_j) \approx \sum_{i=1}^{2m^2} G_i Y_{ln}(\theta_i, \phi_i)$$

Therefore we can estimate the $a_{ln}(\alpha_j)$ by computing G_i for each facet, and then for each l and n, adding $G_i Y_{ln}(\theta_i, \phi_i)$ to the appropriate sum. The resulting spherical harmonic expansion defines a smooth function which is a least squares approximation to the density distribution of the flux in the various reflected beams.

9. Bump Maps and Results

The bidirectional reflection algorithm presented in section 5 relies on precomputed height or bump maps. We generated several bidirectional reflection function tables based on several different 256 x 256 bump maps. These bump maps represented different types of surface micro geometries.

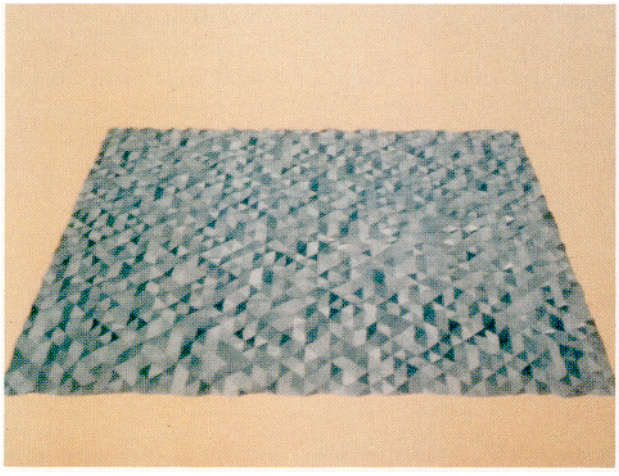

Figure 3a. Sample white noise surface.

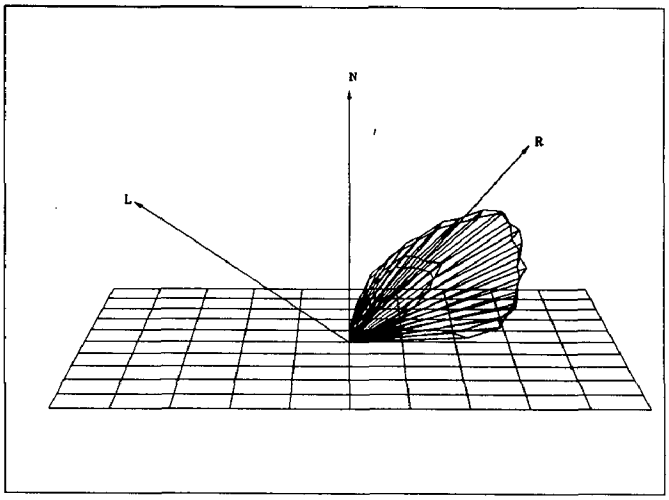

Figure 3b. BDRF with $\theta_{in} = 0.81$ and $\phi_{in} = 0.0$

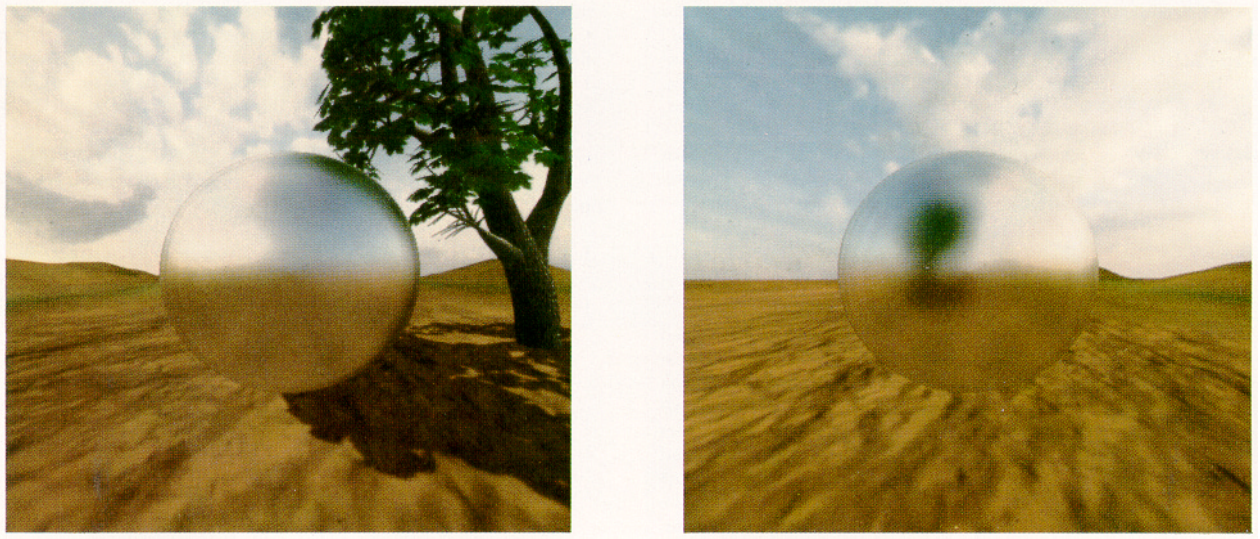

Figure 4a. Sphere reflecting environment map.

The random bump map was made up of scaled white noise. This surface was supposed to model isotropic diffuse surfaces much like the models used by Torrence and Sparrow [32]. The scale of the noise was chosen to be such that ratio of the heights to the inter-height spacing was similar to real surfaces. According to [29] this ratio is on the order of $0.01 - 0.1$. We used a of scales 0.3 for our images. A scale of 0.3 results in surface with an RMS slope of $m \approx 0.218$ [8]. Figure 3a is a sample of this type of random surface. Figure 3b is a bidirectional reflection funtion of the surface in figure 3a for the incident angles $\theta = 0.81$ and $\phi = 0.0$. We built a 24 x 24 x 24 x 24 bidirectional reflection function table. The took 12 hours of VAX 11/785 cpu time to build the table..

Figure 4a and 4b are front and back renderings of an environment map (from NYIT) as reflected off a sphere. The sphere surface's BDRF was generated with the method discussed in section 5. The BDRF was combined with the spherical harmonics with the method described in section 8.

Conclusion

Table generated bidirectional reflection functions based on bump mapped micro surfaces can be used with spherical harmonics to render environments reflected off of various surfaces. We presented methods to approximate the micro surfaces, light reflecting off the surfaces, and the environment-reflection integral (eq. 1).

There are a number of future directions that need to be explored. Much more study into the micro surface properties of real surfaces needs to be done in order to more accurately model various micro surfaces. A bidirectional reflection function algorithm based on ray tracing could produce a more accurate horizon approximation and model higher order reflections. This ray tracer could be designed to exploit the surface geometry and thus be made to run faster than an ordinary ray tracer. Lastly a technique could be developed for taking the bidirectional reflection tables and blending them into texture maps for specific viewing and lighting angles.

Acknowledgements

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract number W-7405-ENG-48. We wish to thank Don Vickers, Ken Joy, Chuck Grant, Michael Gwilliam, Jeff Kallman, and Dave Temple for helpful discussions and suggestions, and Ned Greene and the New York Institute of Technology for the environment map from [19].

Disclaimer

References

1. Blinn, James. *Models of Light Reflection for Computer Synthesized Pictures.* Proceedings of SIGGRAPH '77 (San Jose, California, July 20-22, 1977). In Computer Graphics 11, 2, (July 1977), 192-198.

2. Blinn, James. *Simulation of Wrinkled Surfaces.* Proceedings of SIGGRAPH '78 (Atlanta, Georgia, August 23-25, 1978). In Computer Graphics 12, 3, (August 1978), 286-292.

3. Blinn, James, and Newell, Martin. *Texture and Reflection in Computer Synthesized Pictures.* Comm. ACM. Vol. 19, No. 10, pp. 542-547. 1976.

4. Born, Max, and Wolf, Emil. *Principles of Optics.* Pergamon Press. Oxford 1980

5. Bresenham, J. E. *Algorithm for computer control of a digital plotter.* IBM Systems Journal. Vol 4. pp. 25-30. 1965.

6. Catmull, Edwin, and Rom, Raphael. *A Class of Interpolating Splines,* in Computer Aided Geometric Design. Barnhill and Riesenfeld, Eds. Academic Press, pp. 317-326. New York.1974.

7. Cook, Robert, Porter, Thomas, and Carpenter, Loren. *Distributed Ray Tracing.* Proceedings of SIGGRAPH '84 (Minneapolis, Minnesota, July 23-27, 1984). In Computer Graphics 18, 3, (July 1984), 137-145.

8. Cook, Robert and Torrance, Kenneth. *A Reflectance Model for Computer Graphics.* Proceedings of SIGGRAPH '81 (Dallas, Texas, August 3-7, 1981). In Computer Graphics 15, 3, (August 1981), 307-316.

9. Courant, R., and Hilbert, D. *Methods of Mathematical Physics.* Interscience Publishers, Inc., p. 513. New York. 1953.

10. Crow, Franklin. *Summed-Area Tables for Texture Mapping.* Proceedings of SIGGRAPH '84 (Minneapolis, Minnesota, July 23-27, 1984). In Computer Graphics 18, 3, (July 1984), 207-212.

11. Duff, Tom. *Compositing 3-D Rendered Images.* Proceedings of SIGGRAPH '85 (San Francisco, California, July 22-26, 1985). In Computer Graphics 19, 3, (July 1985), 41-44.

12. Glassner, Andrew. *Adaptive Precision in Texture Mapping.* Proceedings of SIGGRAPH '86 (Dallas, Texas, August 18-22, 1986). In Computer Graphics 20, 4, (August 1986), 297-306.

13. Greene, Ned. *Applications of World Projections.* IEEE CG&A. Vol. 6, No. 11. pp. 21-29.1986.

14. Greene, Ned, and Heckbert, Paul. *Creating Raster Omnimax Images from Multiple Perspective Views using the Elliptical Weighted Average Filter.* IEEE CG&A. Vol. 6, No. 6, pp. 21-27. 1986.

15. Hall, Roy, and Greenberg, Donald. *A Testbed for Realistic Image Synthesis* IEEE Computer Graphics and Applications Volume 3, No. 8, pp. 10-20 1984

16. Heckbert, Paul. *Filtering by Repeated Integration* Proceedings of SIGGRAPH '86 (Dallas, Texas, August 18-22, 1986). In Computer Graphics 20, 4, (August 1986), 315-321.

17. Helmholtz, H.v. *Helmoltz's Treatise on Physiological Optics.* Optical Society Of America. Washington, D.C. 1924.

18. Immel, David, Cohen, Michael, and Greenberg, Donald. *A Radiosity Method for Non-Diffuse Environments.* Proceedings of SIGGRAPH '86 (Dallas, Texas, August 18-22, 1986). In Computer Graphics 20, 4, (August 1986), 133-142.

19. Kajiya, James. *The Rendering Equation.* Proceedings of SIGGRAPH '86 (Dallas, Texas, August 18-22, 1986). In Computer Graphics 20, 4, (August 1986), 143-150.

20. Kajiya, James, and Von Herzen, Brian. *Ray Tracing Volume Densities.* Proceedings of SIGGRAPH '84 (Minneapolis, Minnesota, July 23-27, 1984). In Computer Graphics 18, 3, (July 1984), 165-174.

21. Max, Nelson. *Shadows for Bump Mapped Surfaces.* Advanced Computer Graphics, T. L. Kunii, Ed. Springer Verlag, Tokyo. pp. 145-156. 1986.

22. Miller, Gene, and Hoffman, Robert. *Illumination and Reflection Maps: Simulated Objects in Simulated and Real Environments.* Advanced Image Synthesis Course Notes. Siggraph Conference. 1984.

23. Ohira, Tomohiro. *A Shading Model for Anisotropic Reflection.* Technical Report of The Institute of Electronics and Communication Engineers of Japan, in Japanese. Vol. 82, No. 235, pp. 47-54. 1983.

24. Perlin, Kenneth. Course Notes. Siggraph Conference. 1984.

25. Perlin, Kenneth. *Course Notes.* SIGGRAPH '85 State of the Art in Image Synthesis Seminar Notes. submitted to IEEE CG&A, and personal communication. 1986.

26. Phong, Bui-Tuong. *Illumination for Computer Generated Images.* Comm. ACM. Vol. 18, No. 6, pp. 311-17. 1975.

27. Rense, W. A. *Polarization Studies of Light Diffusely Reflected from Ground and Etched Glass Surfaces.* J. Opt. Soc. Am. Vol 40, No. 1. pp. 55-59. 1950.

28. Snell, Jay. *Radiometery and Photometery.* Handbook of Optics. Driscoll, W. G. and Vaughen. W., Eds. McGraw-Hill. 1978.

29. Spangenberg, D. B., Strang, A. G., and Chamberlin, J. L. *Surface Texture Measurments of Metal Surfaces* National Bureau Standards, Special Publication 300 Vol 7. Washington, D.C. 1971.

30. Takagi, J., Yokoi, S., and Tsurwoka, S. *Comment on the Anisotropic Reflection Model.* Bulletin of SIG. Graphics and CAD, Information Processing Society of Japan, in Japanese. Vol. 11. No. 1 pp. 1-9. 1983.

31. Tinkham, M. *Group Theory and Quantum Mechanics.* McGraw Hill pp.101-115. New York. 1964.

32. Torrance, Kenneth, and Sparrow, Ephraim. *Theory for Off-Specular Reflection from Roughened Surface.* Journal of the Optical Society of America. Volume 57 No. 9 1967.

33. *USA Standard Nomenclature and Definitions for Illuminating Engineering.* USAS Z7.1-1967. 1967.

34. Whitted, Turner. *An Improved Illumination Model for Shaded Display.* Comm. ACM. Vol. 23, No. 6, pp. 343-349. 1980.

35. Whitted, Turner, and Cook, Robert. *A Comprehensive Shading Model.* Image Rendering Tricks Course Notes. Siggraph Conference. 1985.

36. Williams, Lance. *Pyramidal Parametrics.* Proceedings of SIGGRAPH '83 (Detroit, Michigan, July 25-29, 1983). In Computer Graphics 17, 3, (July 1983), 1-11.