

Wavelength Dependent Reflectance Functions

Jay S. Gondek, Gary W. Meyer, Jonathan G. Newman

Department of Computer and Information Science
University of Oregon

Abstract

A wavelength based bidirectional reflectance function is developed for use in realistic image synthesis. A geodesic sphere is employed to represent the BRDF, and a novel data structure is used to store this description and to recall it for rendering purposes. A virtual goniospectrophotometer is implemented by using a Monte Carlo ray tracer to cast rays into a surface. An optics model that incorporates phase is used in the ray tracer to simulate interference effects. An adaptive subdivision technique is applied to elaborate the data structure from rays scattered into the hemisphere above the surface. The wavelength based BRDF and virtual goniospectrophotometer are utilized to analyze and make pictures of thin films, idealized pigmented materials, and pearlescent paints.

Categories and Subject Descriptors: I.3.7 [Computer Graphics]: *Three-Dimensional Graphics and Realism*.

Additional Key Words and Phrases: BRDF, full spectral rendering, Monte Carlo.

1 Introduction

The appearance of an object is determined by both the spatial distribution and the wavelength composition of the light that is reflected from the object's surface. These geometric and optical properties of the scattered light are what an observer uses to determine what a material looks like. Variation in the spatial distribution of the reflected light causes changes in appearance characteristics such as *gloss*, *haze*, *luster*, and *translucency*. Changes in the wavelength composition of the reflected light can alter the *hue*, *saturation*, and *lightness* that are seen by the observer. Spectrophotometric measurements can be taken to determine the spectral energy distribution of the reflected light and goniospectrometric measurements can be made to find the spatial distribution of that light. The *measurement of appearance* is the term that has been coined to identify the family of measurements that are necessary to characterize both the color and the surface finish of an object [13].

While it is possible to separately measure the spatial and spectral distribution of the reflected light, these two dimensions work together to establish the overall appearance of an object. For example, the spectral distribution of the light reflected from a paint or a plastic is not completely determined by the absorptive properties of the pigment particles below the surface of the material. Light

that travels into the substance interacts with the pigment particles and this interplay does change the wavelength composition of the light that eventually leaves the material. However, roughening the surface will increase the amount of spectrally nonselective light that is reflected in all directions from the object's surface. This desaturates the color of the object and thus alters its appearance even though the pigment has not been changed. On the other hand, the spatial distribution of the light reflected from multilayer systems (such as an iridescent paint) is only partially determined by the surface roughness of the material. Light that reflects from the topmost surface may be scattered uniformly in all directions. However, light that enters the surface, interacts with the layers below, and emerges from the surface, may have certain wavelengths reinforced while the rest of the spectrum is cancelled. When the reinforcement takes place in the mirror direction, the amount of specularly reflected light is increased and the specular reflection assumes a color that is different than the incident light. As a result, a metallic appearance is produced that might not have been predicted from surface roughness alone.

There are two basic types of local illumination models that are used in computer graphics to control the appearance of an object. In the most widely used approach, separate reflection mechanisms are independently modeled and then linearly combined [1, 5, 11, 16]. The most elaborate of these models, the He model [11], includes directional diffuse and ideal specular terms to account for surface reflection and an ideal diffuse component to accommodate subsurface scattering. The directional diffuse and ideal specular terms in the He model permit spectral reflectance to vary with incident angle, but the wavelength distribution of the ideal diffuse component is modeled as being constant across the scattering hemisphere. Also important amongst the linear combination models are those that include anisotropy in the spatial distribution [14, 17, 22]. The second and most recent type of local illumination model employs a bidirectional reflectance function (BRDF). Data for BRDFs has been generated by performing simulations [4, 23, 10] and by taking measurements [21]. Imaging systems that use BRDFs as their principle local illumination model have been developed [15, 19]. The only BRDF representation scheme proposed thus far for rendering has been the spherical harmonic approach suggested by Cabral et al. [4]. The use of spherical harmonics makes it possible to create reflectance distributions that have arbitrary spatial complexity. However, the use of a separate spherical harmonic representation for each spectral sample is not very efficient.

This paper presents a local illumination model that increases the generality with which an object's appearance can be specified. This is accomplished by adopting a BRDF representation scheme that allows a flexible and efficient representation of spectral as well as spatial information. This spectral BRDF is created by model-

Permission to make digital or hard copies of part or all of this work or personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

SIGGRAPH '94, July 24-29, Orlando, Florida
© ACM 1994 ISBN: 0-89791-667-0 ...\$5.00

ing surface and subsurface microstructures and then casting rays at the surface from all positions in the hemisphere above the surface. The rays are traced into the surface where they interact with the subsurface microstructures. An optics model is used that includes phase information. This allows simulation of interference effects due to path length differences and phase changes. A unique data structure is employed to record the rays that are reflected from each incident direction in the hemisphere. The data structure is created by using an adaptive scheme to elaborate the structure for those regions into which the most rays were reflected. This data structure is subsequently used in a shader to create realistic images.

This paper is divided into three additional major sections. In the following section the optics model that incorporates phase information and that was used to create the spectral BRDFs is described. Next, the data structure that was employed to represent the spectral BRDF is discussed. Finally, the simulations that were done to generate the spectral BRDFs are covered. The paper closes with examples that show how the spectral BRDF can be used to capture subtle appearance variations that are the result of both spatial and spectral variation in the light reflected from a surface.

2 The Ray Optics Model

In computer graphics, the ray model of light has been successfully used for over a decade to produce realistic images. In ray optics, each ray represents the propagation of a point sample on a wavefront of light. Light waves are generally described in terms of a number of physical properties, including wavelength, amplitude, speed, state of polarization, and phase. Polarization effects can become visible after multiple reflections from dielectrics and metals, as demonstrated in computer graphics by Wolff and Kurlander [24]. The phase of the light wave is needed to compute interference effects. Previous work in rendering the phenomenon of thin film interference has relied on analytical solutions for determining the reflectance from a thin film surface [20, 6]. By introducing the attribute of phase to the current ray model in computer graphics, we are able to predict interference results by simulation rather than relying on analytical solutions. Furthermore, this model allows the computation of reflectance functions from complex interference systems, such as iridescent paint, that defy an analytical solution.

Light is an electromagnetic wave, and thus has an electric and a magnetic field associated with it. Because the electric and magnetic fields are simply related in nonbirefringent dielectric media, the magnetic field can be disregarded and the behavior of light can be described in terms of just the electric field. Natural light can be represented as two arbitrary, incoherent, orthogonal, linearly polarized waves of equal amplitude [12]. Thus, a simulation of natural light is implemented by initially assigning two fields, E_1 and E_2 , to each ray incident from the light source; these fields are plane polarized, orthogonal to each other, and each has an initial amplitude of one. Because the two fields are incoherent with respect to each other, they are treated independently in all subsequent calculations. The propagation of light through the surface model will therefore be described in terms of a single electric field E .

2.1 The Propagation of Light

The polarization state of E can be modeled by two orthogonal vector components, E_x and E_y , that are also both orthogonal to the direction of propagation. Each component may have a unique amplitude and phase. An electric field for a ray that has propagated some distance through the environment is described by the functions

$$E_x = A_x \sin(\omega(t + P/c) + \delta_x) \quad (1)$$

$$E_y = A_y \sin(\omega(t + P/c) + \delta_y) \quad (2)$$

where A is the vector component amplitude, t is time, ω is the angular frequency, P is the optical path length that the ray has traversed from the light source to its current position, c is the speed of light in vacuum, and δ is a phase offset that results from the




Figure 1: Polarization geometry of incident light.

accumulation of phase jumps that may occur upon reflection, as detailed below.

When a ray strikes a transparent surface, a transmitted ray and a reflected ray are produced, as shown in Figure 1. The components of E for the incident ray (E_x and E_y) are re-cast in terms of a new orthogonal basis that coincides with the plane of incidence and the perpendicular to the plane of incidence (x' and y'). The amplitudes of the resulting field components are multiplied by the Fresnel amplitude coefficients for reflection and transmission (r_{\perp} , r_{\parallel} , t_{\perp} , and t_{\parallel} in Figure 1) to give the reflected and transmitted amplitudes [12].

The field components of a ray may undergo independent changes in phase upon reflection. In the case where total internal reflection does not occur, this change in phase is predicted by the signs of the solutions to the Fresnel equations. Otherwise the phase jump is given by the following equations

$$\tan \frac{\delta_{\perp}}{2} = -\frac{\sqrt{\sin^2 \theta_i - n^2}}{\cos \theta_i} \quad (3)$$

$$\tan \frac{\delta_{\parallel}}{2} = -\frac{\sqrt{\sin^2 \theta_i - n^2}}{n^2 \cos \theta_i} \quad (4)$$

where δ_{\perp} and δ_{\parallel} are the phase jumps for the reflected fields parallel or perpendicular to the plane of incidence, θ_i is the incident ray angle, and n is the transmitted index of refraction divided by the incident index of refraction [3]. The new phase offset for each reflected or transmitted component is equal to the sum of the incident phase offset and the phase jump calculated at the time of reflection.

2.2 Interference

After a single ray enters the geometric model, the phases associated with the field components of the multiple exiting rays are used to compute interference. This model predicts thin film interference, or, more generally, amplitude splitting interference. Amplitude splitting interference occurs when a wave is divided through reflection and transmission, and later recombines. We do not attempt to model wavefront splitting interference, which is commonly the result of diffraction.

Exiting rays that share the same direction of propagation to within a small tolerance have the potential to interfere. These rays represent partially overlapping wavefronts of light that are focused by the eye or camera lens at an arbitrary distance. According to the Fresnel-Arago laws, interference may be produced in the case where components of electric fields share the same plane of polarization, with the caveat that electric fields originating from different incoherent fields of natural light never produce visible interference and must be treated independently [12]. The coherent field components of parallel rays are summed to produce amplitudes at each




Figure 2: Angles of light scattering.

wavelength that are the result of interference. This summation is achieved by using a trigonometric identity that is described in optics texts (for example, see [12]).

3 Representing the Scattering Function

The ratio of exitance to incident irradiance is computed over the spectrum for all reflected and transmitted rays. Irradiance, D , is defined as $D = c\epsilon_0 A^2/2$ for an electric field vector with amplitude A where ϵ_0 is the electric permittivity of free space [12]. Since the two initial fields of white light each had an amplitude of one, the initial spectral irradiance, $D_{\lambda,in}$, is simply $c\epsilon_0$. If n rays are used to represent the incident wavefront of light, then the spectral exitance for each ray is

$$D_{\lambda,out} = \left(\frac{c\epsilon_0}{2n}\right) (A_{1,x}^2 + A_{1,y}^2 + A_{2,x}^2 + A_{2,y}^2) \quad (5)$$

where $A_{1,x}$, $A_{1,y}$, $A_{2,x}$, and $A_{2,y}$ are the spectral amplitudes of the electric field components of the two incoherent fields associated with the ray. Rays from all sampled incident directions contribute to the final BRDF representation.

The bidirectional spectral reflectivity of a material ρ'' is the ratio of reflected radiance L_{out} to incident flux density D_{in} , defined for all incident and reflected directions over the hemisphere enclosing a surface element [18]. A BRDF is a function of five variables

$$\rho''(\lambda, \theta_r, \phi_r, \theta_i, \phi_i) = \frac{L_{out}(\lambda, \theta_r, \phi_r, \theta_i, \phi_i)}{D_{in}(\lambda, \theta_i, \phi_i)} \quad (6)$$

for light incident at an elevation angle of θ_i and an azimuth angle of ϕ_i , reflected in the direction of θ_r, ϕ_r (see Figure 2). In choosing this formulation of ρ'' , we assume that incident light is unpolarized, and that the sampled surface does not contain fluorescent material. L_{out} is the radiance of outgoing light, and is expressed in terms of exitance as

$$L_{out} = \frac{D_{out}}{d\omega_{out} \cos \theta_{out}} \quad (7)$$

for flux density propagating through the solid angle $d\omega_{out}$. The bidirectional spectral transmissivity function, or BTDF, is similarly defined for the transmission of light.

Cabral, Max, and Springmeyer [4] used an array as one representation for a BRDF. In their approach, this array essentially represents a set of discrete buckets that cover the hemisphere above a sample element. These buckets are used to capture the scattering of light from each surface element for a particular incident angle. We use an improved data structure that is related to the Cabral et al. method to give a compact and accurate representation of the BRDF and BTDF. Furthermore, this data structure is used directly by a Monte Carlo ray tracer to render images.




Figure 3: Levels of Geodesic Sphere Construction

3.1 The Capture Sphere

An adaptively-built geodesic sphere of unit radius is employed to capture the scattered rays, where facets of this sphere serve as the capture buckets. A single sphere captures the reflectance and transmittance from all incident light angles. Furthermore, this sphere is used to represent the light scattering function upon completion of the simulation. A full sphere is required in the simulation not only to characterize the transmission function, but also because the direction of exiting rays will undergo a transformation, detailed in section 3.3, that may change the orientation of the exiting hemisphere depending on the angle of incidence. Figure 3 shows an example of the recursive subdivision technique that is used to build the sphere geometry. In Figure 3, each facet of the sphere is divided into four smaller facets with each increasing level of recursion. The facets of the sphere serve as buckets to tabulate the ratio of exitant to incident flux density. The ratio $D_{\lambda,k}/D_{\lambda,i}$ for sphere facet k is given by

$$\frac{D_{\lambda,k}}{D_{\lambda,i}} = \frac{1}{D_{\lambda,i}} \sum_R D_{\lambda,R} \quad (8)$$

for all exiting rays R that have direction vectors passing through cell k . By substituting Equation 7 into Equation 6, the BRDF for cell k can be expressed as

$$\rho''(\lambda, \theta_k, \phi_k, \theta_i, \phi_i) = \frac{D_{\lambda,k}}{D_{\lambda,i} d\omega_k \cos \theta_k} \quad (9)$$

where $d\omega_k$ is the solid angle of cell k .

3.2 Adaptive Subdivision

Eight quad-trees are used to represent the triangles that form the geodesic sphere. Each root node denotes a basis triangle (four of these triangles are visible in the left polyhedron of Figure 3). A node at level i represents a triangle that is the result of i subdivisions. An adaptive approach is used to subdivide facets independently, thus providing a variable sampling resolution to capture features in the exitance distribution. Figure 4 shows the steps in adaptive subdivision that occurred when capturing the reflected flux scattered from a Gaussian surface with light incident at 25 degrees.

During subdivision, a facet k is potentially divided into facets k_1, k_2, k_3 , and k_4 . To decide if subdivision is necessary, the root mean squared deviation (RMSD) is computed at each wavelength λ for the spectral flux density propagating through cells k_1, \dots, k_4 . If the average of the spectral RMSD values is above a small tolerance, then subdivision occurs, and the new cells are recursively tested to see if further subdivision is required. Otherwise, a record containing the incident direction (θ_{in}, ϕ_{in}) , spectral exitance ratios $(D_{\lambda,k}/D_{\lambda,i})$, and a reflectance or transmittance flag is inserted at the current node (for cell k) in the tree structure. The need for identifying the exitance




Figure 4: Subdivision for Capturing a Scattering Distribution

as being either reflected or transmitted will be detailed in the next section. Clearly, only the leaf nodes will contain data after the first incident angle is sampled. As more incident angles are sampled, cells represented by interior nodes may obtain data; furthermore, new nodes may be required to capture the additional exitance.

3.3 Transforming the Direction Vectors

Consider the problem of accurately characterizing a mirror reflection with this adaptive subdivision method. In this case, all of the rays incident from θ_i, ϕ_i will be captured by a small bucket aligned with the $\theta_i, \phi_i + \pi$ direction. This implies a high degree of subdivision with every incident angle, impacting storage space and computation time during the simulation. Furthermore, such a representation is difficult to accurately interpolate and creates added time and space expense during rendering.

We define a transformation T_r that is dependent on the angle of incidence, such that T_r is a rotation of $-(\phi_i + \pi)$ about the normal, followed by a rotation of $-\theta_i$ about the y_s axis of the surface (see Figure 2). For the example of capturing a mirror reflection, this transformation rotates every reflected ray into the surface normal, and a single small bucket at the top of the sphere captures the mirror-reflected rays for all incident directions. Transmitted rays are rotated with transformation T_t in a similar way so that the ideal transmission direction is also aligned with the top of the sphere. Thus, for light incident on a smooth pane of glass in air, all reflected and transmitted rays are captured by the same small bucket. Note that multiple records may be associated with each bucket to account for different sample directions.

This method of concisely capturing ideal reflected and transmitted rays extends to any distribution that shows a directional bias for scattering in the mirror or ideal transmission directions. For example, the spatial reflectance distribution from many rough surfaces exhibit a specular lobe that is approximately aligned with the mirror reflectance direction. In such a case, the entire specular lobe is rotated into the direction normal to the surface, substantially localizing the area of high subdivision. For a given bucket and an associated data record the inverse transformations T_r^{-1} and T_t^{-1} can be applied to the vertices of the triangular bucket to recover the cell corresponding to the direction of the unrotated rays.

3.4 Interpolating the Function

Because the surface is sampled in a number of discrete incident directions, the tree structure for the capture sphere provides a scattering function that can be evaluated only at these specific directions. From this, we move to a representation of the BRDF and BTDF that is valid over continuous ranges of incident directions. Suppose that we have samples at incident directions $(\theta_1, \phi_1), (\theta_1, \phi_2), (\theta_2, \phi_1),$ and (θ_2, ϕ_2) . Given an incident direction θ_{in}, ϕ_{in} such that $\theta_1 \leq \theta_{in} \leq \theta_2$ and $\phi_1 \leq \phi_{in} \leq \phi_2$, an interpolated result

can be found with a two-step process. First, for any cells that are larger for one sample direction than another, a common group of cells is computed by subdividing the cells associated with the four surrounding incident angles. This new group of cells becomes the set of cells associated with (θ_{in}, ϕ_{in}) . Second, the spectral flux density ratios for these cells are computed with bi-linear interpolation of the values from $(\theta_1, \phi_1), (\theta_1, \phi_2), (\theta_2, \phi_1),$ and (θ_2, ϕ_2) . Interpolation of the direction of reflection or transmission simply falls out as a result of assigning (θ_{in}, ϕ_{in}) to this group of cells before applying the T_r^{-1} or T_t^{-1} transform. This interpolation method gives a scattering function that can be evaluated for any direction inside the range of sampled incident directions.

3.5 Rendering the Data

The above interpolation method is used within the context of a raytracer to render objects constructed from materials that have had their light scattering properties characterized by simulation. Let (θ_r, ϕ_r) be the reflected direction from a surface in the backward raytracing paradigm. Because of reciprocity, the raytracer can “look out” into the scene through the cells that were used to capture the exitance from light incident at (θ_r, ϕ_r) . The reflected spectral radiance, L_λ , for (θ_r, ϕ_r) is given by

$$L_\lambda(\theta_r, \phi_r) = \sum_k \rho''(\lambda, \theta_k, \phi_k, \theta_r, \phi_r) L(\theta_k, \phi_k) \cos \theta_k d\omega_k. \quad (10)$$

Equation 9 is substituted into Equation 10 to give

$$L_\lambda(\theta_r, \phi_r) = \sum_k \frac{D_{\lambda,k}}{D_{\lambda,i}} L(\theta_k, \phi_k). \quad (11)$$

In this formulation, both the cosine and solid angle terms cancel. This is one reason why we chose to capture the flux density ratios rather than the actual BRDF in the data structure representation, although the BRDF can easily be computed from the data structure by Equation 9. The flux density ratio for any cell is then used as a measure of the importance of that cell for sampling. To implement this, cells with higher flux ratios are sampled more densely than those with low flux ratios. This gives importance-based sampling of the local scattering function.

4 Light Scattering Simulations

The optics model and capture dome described in the previous two sections are used in conjunction with the geometric modeling of surface microstructure to produce a virtual goniospectrophotometer. During the simulation, rays are cast into layered surfaces that may have specific spectral absorption properties. Furthermore, structural colors produced by interference are accounted for, as detailed




Figure 5: Simple paint microgeometry.

in Section 2. The number of rays per incident angle and the number of sample directions are specified before the simulation begins (although it is conceivable that adaptive sampling techniques could be employed to make these decisions during execution). Twenty wavelength samples spaced over the visible spectrum were used for the simulations and renderings described in the following subsections. All reflection and refraction of rays occur in the ideal mirror and ideal refracted directions, respectively. Any diffuse appearance that is produced from rendering the resulting BRDF and BTDF is fundamentally the accumulation of these reflections and transmissions.

Cabral et al. [4] did some of the early work in characterizing the reflectance functions of modeled surface geometry. Westin et al. [23] introduced the approach of casting rays at complex models of surface microstructure. Hanrahan and Krueger [10] proposed an analytical model of subsurface scattering that more accurately defines the diffuse reflection component, but lacks the generality of simulating actual subsurface microstructure.

In the following simulations, we considered the index of refraction to be constant across the visible spectrum. With this assumption, an entire spectrum can be assigned to each ray. Furthermore, with a wavelength independent index of refraction, Fresnel's equations predict the same amplitude coefficients based on the po-




Figure 6: The effect of increasing surface roughness in a simple pigment model.

larization state of the ray, regardless of the wavelength. Thus the electric fields associated with a ray describe the polarization state for the entire spectrum that the ray represents. This simplification was made to decrease the computation time of the simulation. Another option, which would allow for wavelength dependent indices of refraction, could be implemented by storing field representations at each wavelength. More generally, a ray could be cast for each wavelength, facilitating both dispersion and wavelength dependent indices of refraction. To allow for pigments and dyes, spectrally dependent absorption is implemented by maintaining a list of wavelength dependent amplitude coefficients with each ray.

4.1 Plastics and Paints

Plastics and paints consist of pigment particles suspended in a transparent matrix. Light scattering qualities of the particles are determined by their geometry and composition. The aggregate result of scattering from numerous particles contributes to the spatial and spectral distribution of reflected light. In many paints, the reflection of light from the surface of each pigment particle is responsible for the change in wavelength composition of incident light. In other paints, the transmission of light through each pigment particle causes spectrally dependent absorption of incoming light; reflection from the surface of the particles simply contributes to the diffusion of light within the substrate [7, 13]. Pigmented surfaces are common in most settings, so it is natural that the problem of rendering such surfaces has received attention in computer graphics [21, 9].

We present a simple model of a paint coating by defining a substrate filled with small pigment-particle microspheres, as shown in Figure 5. These particles are modeled as dielectric spectral filters; attenuation of transmitted light is governed by Bouguer's Law [7]. Although this model is not an attempt to rigorously characterize the geometry and material attributes of pigment surfaces, it serves to demonstrate some of the reflective properties that real paints exhibit.

Surface roughness is a significant contributing factor to the appearance of paints and plastics. Smooth surface finishes give a spectrally non-selective mirror reflection for a portion of the incident light. A surface with this property is described as *shiny* or *glossy*. As the texture of the upper surface becomes rough, the spectrally non-selective reflected light assumes a diffuse distribution. In such a case, spectrally dependent scattering from the interior of the surface is combined with the reflectance from the air-surface interface to produce a less saturated, *matte* appearance.

The effect of surface roughness is demonstrated in Figure 6. The




Figure 7: The effect of decreasing particle size in a simple pigment model.




Figure 8: Sunglasses with a thin film coating.

surface microstructures for each of the tori in Figure 6 were modeled as 2 mm wide and 1.2 mm deep sections of paint. Each microsphere had an index of refraction of 2.0 and was $80\ \mu\text{m}$ in diameter. The surrounding matrix had an index of refraction of 1.3. The geometry of the air-surface interface for each paint microstructure was defined as a Gaussian surface with an autocorrelation length of $20\ \mu\text{m}$ and RMS surface heights of 0.0, 1.5, 3.0, and $6.0\ \mu\text{m}$. The undercoating was an ideal black absorber. Incident rays that happened to exit from the sides of the microstructure instead of the top were not considered valid, and thus did not contribute to the computed reflectance function. The tori in Figure 6 were illuminated with an area light source in the shape of a four-paned window. As the surface roughness increases, the specular highlight makes a transition from glossy to diffuse. Because of the increase in diffuse scattering of white light, there is an observable decrease in the color saturation of the lower-right torus compared to the upper-left torus.

Another interesting property of pigments is the relationship between particle size, absorption, and scattering. Interior scattering events generally occur at the interface between the pigment particle and the surrounding matrix. The nature of the effect depends upon whether the particles are opaque or transparent. In the case where the particles are transparent, the amount of absorption that occurs depends on the path length between the entry and exit points for light passing through each particle. Thus, the ratio of absorption to scattering decreases with a decrease in particle size [7, 13]. The tori in Figure 7 illustrate this phenomenon. In this figure, the surface microstructure for each torus has an identical shape, but the scale of the microstructure decreases from left to right and top to bottom, producing microsphere diameters of 80, 40, 20, and $10\ \mu\text{m}$. Because of this shift in scale, the resulting change in the reflectance function is exclusively the effect of the decreasing microsphere diameters. The decline in the ratio of relative absorption to scattering is evidenced by the steps of desaturation that are visible in Figure 7.

4.2 Interference Structures

Absorptive pigmentary coloration is but a single class of color producing phenomenon. Colors that are produced entirely by the optical effects of surface geometry are termed *structural colors*. One example of a structural color is the color exhibited by thin film

layers. Section 2 describes the optics model that was used to predict thin film interference during the light scattering simulation.

Figure 9 shows the result of rendering data from the simulation of light incident on eight simple thin films. Each sphere in Figure 9 was illuminated by two window-shaped area light sources. The films have an index of refraction of 2.3 and are surrounded by air. The index of refraction was chosen to approximate that of titanium dioxide, an interference producing agent in iridescent paints. Transmitted rays are absorbed by an ideal black backing so that only the reflected color is observable. This figure depicts the two most brilliant series of interference color. Optical thicknesses for the top row of films, from left to right, are 218.5, 264.5, 310.5, and $345.0\ \text{nm}$. The bottom row of films have optical thicknesses of 448.5, 517.5, 586.5, and $667.0\ \text{nm}$. Note the change in color as the incident direction goes from normal to oblique. This shift in appearance is due to the change in the optical path of interfering rays as well as the change in Fresnel amplitude coefficients. Figure 8 shows a pair of sunglasses with a thin film coated lens in a texture-mapped environment. In this image, both reflection and transmission are characterized with a full bidirectional scattering function.

Figure 10 compares the analytical solution for thin film reflectance at normal incidence with results from simulation. These




Figure 9: Thin film surfaces showing two orders of interference colors.




Figure 10: Spectral reflectance for the first and second order blue interference colors. Smooth curves correspond to the analytical solution, discrete points correspond to the results of simulation.

curves correspond to the blue interference films shown in Figure 9 (third column). The discrete points on this graph are data produced by ray sampling; the smooth curves are the result of evaluating the analytical formula for single layer thin film reflectance [18]. The results from the simulation clearly match the analytical solution.

Titanium dioxide (TiO_2) coated mica is often used as the color-producing pigment in iridescent and pearlescent paints [2, 8]. The optical thickness of the TiO_2 coating is controlled and varies from about 40 to 400 nm for different paints; the thickness of the mica is not controlled and randomly varies from particle to particle in the same paint. For different paints, the TiO_2 coated platelets can vary in width from 2 to 50 μm . As the width increases, the interference paint tends to produce more of a *sparkling* reflection. The platelets in the paint align themselves in a somewhat parallel fashion upon application, and produce a semi-specular metallic appearance when applied to dark surfaces. If this paint is applied to a light surface, then complementary light that is transmitted through the paint and reflected from the undercoating has a tendency to reduce the saturation of the overall spectral scattering distribution.

Figure 12 is an image of an iridescent blue teapot that was rendered with data computed by simulation. Figure 11 shows the geometry of the modeled surface structure. For this simulation, the optical thickness of the TiO_2 coating was 310 nm , the thickness of the mica enveloped in TiO_2 varied from 10 to 300 nm , and the width of each platelet was 20 μm . The index of refraction for the TiO_2 , mica, and surrounding substrate were 2.3, 1.58, and 1.3, respectively. The platelets were randomly distributed in the medium and randomly rotated slightly out of parallel. The combined effect of the placement and rotation gives the paint a more subtle change in hue with the change in incident angle, and creates a semi-specular reflection.

When the optical thickness of the TiO_2 coating is in the range of 80 to 150 nm , constructive interference occurs across most of the incident spectrum, reflecting approximately 30 to 40 percent of the light in the bluish white to yellowish white regime of the incident spectrum [8]. Interference paints with this characteristic are known as pearlescent pigments. Figure 13 shows an image of an artificial pearl that was created by modeling interference pigments. The optical thickness of the TiO_2 coating for this example was 122 nm . The index of refraction for the matrix surrounding the platelets was set to 1.1, providing a minimal specular contribution and enhancing the luster of the directional diffuse scattering from the platelets. A similar effect could be produced by allowing the platelets to extend out of the matrix, as often occurs in real paints.




Figure 11: Microstructure of interference paint.

5 Summary and Conclusions

In this paper we have presented a new computer graphics reflection model that has sufficient generality to represent both light scattering dimensions considered important by those in industry who make careful measurements of appearance. The model accomplishes this by using a wavelength based BRDF to describe the spatial *and* the spectral distribution of the light reflected from a surface. This is a step beyond current BRDF based reflection models in which it is difficult to efficiently incorporate wavelength information. A new method has been introduced to represent this extended BRDF. It employs a data structure to describe a geodesic sphere that has been adaptively subdivided to capture the important features of the BRDF.

We have also described a virtual goniospectrophotometer that was developed to measure these wavelength based BRDFs. The Monte Carlo ray tracer used to simulate this device adopted a ray optics model that kept track of phase information and could perform interference calculations when necessary. Models were constructed of subsurface microstructures and rays were traced beneath an object's surface. The virtual goniospectrophotometer was employed to show how the extended BRDF can be used to model both a spectral change in reflected light caused by a modification to surface roughness and a spatial change in reflected light produced by wavelength based interference below the object's surface. These changes in appearance could not be accurately modeled without including the dimension of wavelength in the BRDF.

This work represents a first step in developing a powerful tool for those interested in exploring how surface and subsurface microstructures affect the appearance of an object. However, to turn it into a true design tool, several important steps must be taken. Comparisons must be performed against goniospectrophotometric data taken from real surfaces for which the surface and subsurface geometries are known. The optical model must be extended to include other effects such as diffraction and dispersion. Improved techniques are necessary to sample the spectrum, control the geodesic sphere subdivision, and interpolate the information stored in the data structure. Even without these improvements, however, the system can generate physically consistent data for realistic imaging applications. This is still quite useful, because in many cases measured appearance data is simply not available.

6 Acknowledgements

This work was supported by the ATG Graphics Software Group at Apple Computer and the Oregon Advanced Computing Institute. Additional support was provided by the Hewlett-Packard Company which donated the HP Apollo workstations on which most of the large scale simulations were performed. The authors acknowledge Rob Wicke who wrote a renderer that was used to make some of



Figure 12: Iridescent teapot rendered by simulating interference paint.

the first pictures with the new reflection model. Finally, the authors thank Professor Michael Raymer of the University of Oregon Physics Department for several helpful discussions and for his careful review of Section 2. His comments helped to refine and sharpen several of the concepts presented in this paper.

7 References

- [1] Blinn, James F. Models of Light Reflection for Computer Synthesized Pictures. Proceedings of SIGGRAPH '77. In *Computer Graphics 11*, 2, (July 1977) 192-198.
- [2] Bolomey, R. A. and L. M. Greenstein. Optical Characteristics Of Iridescent And Interference Pigments. *Journal of Paint Technology 44*, 566, (Mar. 1972), 39-50.
- [3] Born, Max and Emil Wolf. *Principles of Optics*. 6th Ed. Pergamon Press, Oxford, 1980.
- [4] Cabral, Brian, Nelson Max, and Rebecca Springmeyer. Bidirectional Reflection Functions from Surface Bump Maps. Proceedings of SIGGRAPH '87. In *Computer Graphics 21*, 4, (July 1987) 273-281.
- [5] Cook, Robert L. and Kenneth E. Torrance. A Reflectance Model for Computer Graphics. *ACM Trans. Graph.* 1, 1, (Jan. 1982), 7-24.
- [6] Dias, Maria L. Ray Tracing Interference Color. *IEEE Comput. Graph. Appl.* 11, 2, (Mar. 1991), 54-60.
- [7] Evans, R. M. *An Introduction to Color*. John Wiley & Sons, New York, 1948.
- [8] Greenstein, L. M. Pearlescence: The Optical Behavior of Nacreous and Interference Pigments. In *Pigment Handbook, Volume III*, T. C. Patton, Ed. John Wiley & Sons, New York, 1973, 357-390.
- [9] Haase, Chet S. and Gary W. Meyer. Modeling Pigmented Materials for Realistic Image Synthesis. *ACM Trans. Graph.* 11, 4 (Oct. 1992), 305-335.
- [10] Hanrahan, Pat and Wolfgang Krueger. Reflection from Layered Surfaces due to Subsurface Scattering. Proceedings of SIGGRAPH 93 In *Computer Graphics, Annual Conference Series*, 1993 165-174.
- [11] He, Xiao D., Kenneth E. Torrance, François X. Sillion, and Donald P. Greenberg. A Comprehensive Physical Model for Light Reflection. Proceedings of SIGGRAPH '91 In *Computer Graphics 25*, 4 (July 1991), 175-186.
- [12] Hecht, Eugene. *Optics*. 2nd Ed. Addison-Wesley Publishing Co. Reading, 1987.
- [13] Hunter, Richard S. and Richard W. Harold. *The Measurement of Appearance*. 2nd Ed. John Wiley & Sons, New York, 1987.
- [14] Kajiya, James T. Anisotropic Reflection Models. Proceedings of SIGGRAPH '85. In *Computer Graphics 19*, 3, (July 1985), 15-21.
- [15] Kajiya, James T. The Rendering Equation. Proceedings of SIGGRAPH '86. In *Computer Graphics 20*, 4, (Aug. 1986), 143-150.
- [16] Phong, Bui-Tuong. Illumination for Computer Generated Pictures. *Commun. ACM 18*, 6 (June 1975), 311-317.
- [17] Poulin, Pierre and Alain Fournier. A Model for Anisotropic Reflection. Proceedings of SIGGRAPH '90. In *Computer Graphics 24*, 4, (Aug. 1990), 273-282.
- [18] Siegel, Robert and John R. Howell. *Thermal Radiation Heat Transfer*. McGraw-Hill, New York, 1981.
- [19] Sillion, François X., James R. Arvo, Stephen H. Westin, Donald P. Greenberg. Proceedings of SIGGRAPH '91. In *Computer Graphics 25*, 4, (July 1991), 187-196.
- [20] Smits, Brian E. and Gary W. Meyer. Newton's Colors: Simulating Interference Phenomena in Realistic Image Synthesis. *Eurographics Workshop on Photosimulation, Realism, and Physics in Computer Graphics Conference Proceedings, 1990*, 185-194.
- [21] Takagi, Atsushi, Hitoshi Takaoka, Tetsuya Oshima, and Yoshinori Ogata. Accurate Rendering Technique Based on Colorimetric Conception. Proceedings of SIGGRAPH '90. In *Computer Graphics 24*, 4, (Aug. 1990), 263-272.
- [22] Ward, Gregory J. Measuring and Modeling Anisotropic Reflection. Proceedings of SIGGRAPH '92. In *Computer Graphics 26*, 2, (July 1992), 265-272.
- [23] Westin, Stephen H., James R. Arvo, and Kenneth E. Torrance. Predicting Reflectance Functions from Complex Surfaces. Proceedings of SIGGRAPH '92. In *Computer Graphics 26*, 2, (July 1992), 255-264.
- [24] Wolff, Lawrence B. and David J. Kurlander. Ray Tracing with Polarization Parameters. *IEEE Comput. Graph. Appl.* 10, 6, (Nov. 1990), 44-55.



Figure 13: An artificial pearl produced from a model of pearlescent paint.