Reprinted from *IEEE Computer Graphics and Applications*, 10, 6 (November 1990). pp. 44-55.

Ray Tracing

Ray Tracing with Polarization Parameters

Lawrence B. Wolff and David J. Kurlander
Columbia University

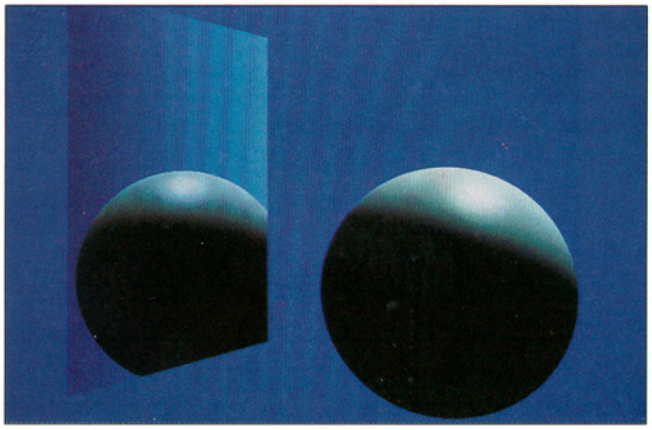

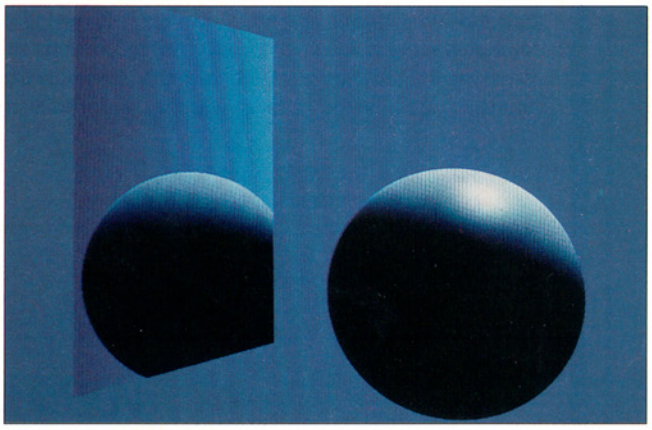

We demonstrate that incorporating polarization parameters into the lighting model can enhance the physical realism of images rendered with a ray tracer. Polarization effects can be important in certain scenes, and the difference in rendering even simple scenes with and without proper treatment of polarization can be rather striking. All light waves possess a state of polarization, which changes almost every time light reflects off a material surface. A single reflection partially polarizes and may even completely polarize previously unpolarized light. Polarization influences the rendering of a scene because the reflected radiant intensity depends largely on the incident light wave's polarization state.

We have incorporated Emil Wolf's coherence matrix formalism for polarization into the Torrance-Sparrow reflectance model. This combination enables elegant quantitative derivations of the altered polarization state of light upon reflection in a ray tracer. Comparisons of identical scenes rendered with a conventional ray tracer and our ray tracer incorporating a polarization model show that our method renders specular interobject reflections more accurately with respect to reflected radiance and color.

Often the realistic rendering of visual scenes demands a theory of reflection that accurately models the polarizing properties of material surfaces. After the first specular reflection, a light wave almost always has a polarization state different from what it originally had after emanating from a light source. Computation of the reflected radiant intensity for a second reflection must be based on the newly acquired polarization state. After two reflections, even initially unpolarized light can have an intensity different from the intensity predicted by lighting models that do not account for polarization.

0272-17-16/90/1100-0044$01.00

In this article we incorporate into a ray tracer a comprehensive representation formalism for light polarization. Previously, ray tracers have ignored the physical phenomenon of polarization and therefore the effects of light polarization on reflected radiance.

A comprehensive formalism of polarization applicable to the rendering of a wide variety of scenes must go beyond modeling completely polarized light. Natural sunlight and lamp light are by and large completely unpolarized and can become partially polarized or even completely polarized when specularly reflected off material surfaces. Emil Wolf's polarization formalism for light is sufficient for describing unpolarized, partially polarized, and completely polarized light.[1,2] The polarization state of light is represented by a 2×2 coherence matrix, which is a Hermitian matrix of special form defined in the section "Polarization and the coherence matrix." Elsewhere,[3] this polarization formalism has been incorporated into the Torrance-Sparrow reflectance model.[4] The Torrance-Sparrow model already inherently accounts for polarization effects.[5] However, incorporating the coherence matrix formalism into this model affords a conceptually elegant methodology for implementing a ray tracer with polarization parameters.

Polarization and the coherence matrix

All light waves can be decomposed into two wave components mutually orthogonal in the plane perpendicular to the direction of propagation. The magnitudes of these wave components can be empirically resolved by transmitting the light through a linear polarizer oriented in the directions of the wave components. A complete description of the polarization state of a light wave consists of the magnitude of its mutually orthogonal wave components along with the correlation of these wave components with respect to one another.

Completely polarized light waves have wave components that may differ in magnitude and also in phase. The wave components are completely correlated in the sense that given the value of the electric field of one wave component at a given instant in time, you can deterministically derive the value of the electric field for the other wave component at the same moment. The difference in phase between two mutually orthogonal wave components refers to the relative translation of one component with respect to the other, expressed in terms of the wavelength of the harmonic oscillations of the two component waves. A phase difference of zero or one half wavelength gives rise to the phenomenon of *linear polarization*, while all other phase differences give rise to *elliptical polarization*.

Almost all naturally occurring light is either partially polarized or completely unpolarized. The wave components of completely unpolarized light have equal magnitudes and are completely uncorrelated, so that at any instant information about the electric field of one wave component gives no information whatsoever about the electric field of the other. Partial correlation between wave components exists for partially polarized light. Except at normal incidence, completely unpolarized light such as that produced by the sun or an incandescent lamp becomes partially polarized when reflected from a material surface. Dielectric surfaces (that is, surfaces that do not conduct electricity) have an angle of incidence—the angle between the incident direction of the light wave and the normal to the surface—at which specularly reflected light becomes linearly polarized. This angle of incidence, called the Brewster angle, is determined by the index of refraction for the dielectric material.

Emil Wolf proposed an elegant mathematical description for all the above light-wave polarization phenomena.[1,2] Selecting the axes x and y to be aligned with the mutually orthogonal component wave (complex) functions E_x and E_y respectively, define the *coherence matrix* $\mathbf{J}$ as

$$\mathbf{J} = \begin{pmatrix} \langle E_x E_x^{\star} \rangle & \langle E_x E_y^{\star} \rangle \\ \langle E_y E_x^{\star} \rangle & \langle E_y E_y^{\star} \rangle \end{pmatrix} \quad (1)$$

where the angle brackets denote values averaged over time and the starred superscript denotes the complex conjugate. The diagonal terms of the coherence matrix denote the x and y time-averaged electric-field intensities, respectively, and are always nonnegatively real valued. The off-diagonal terms denote the time-averaged product of the x and y electric-field components and are in general complex valued.

Clearly the off-diagonal terms of the coherence matrix are complex conjugates of one another and therefore the coherence matrix is Hermitian. Intuitively, the magnitudes of the off-diagonal terms, relative to the product of the diagonal terms, represent the correlation between the x and y electric-field components. Denote the four terms of the coherence matrix by J_{xx}, J_{xy}, J_{yx}, and J_{yy} respectively. The trace of the coherence matrix $Tr\mathbf{J} = J_{xx} + J_{yy}$ is the total intensity of the light radiation.

The representation of the coherence matrix[1] as the polarization state for a light wave depends on the

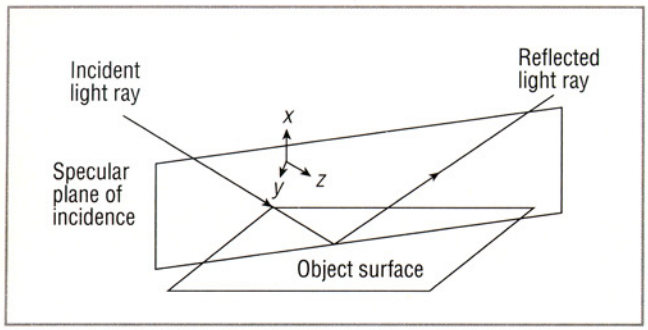

Figure 1. A specular plane of incidence and an associated right-handed 3D coordinate system.

assumption that the width of the spectral distribution for the light wave is very narrow compared with the mean wavelength of the spectral distribution. We call such a narrow spectral distribution *quasi-monochromatic*. A light wave can actually be a packet of waves consisting of diverse wavelengths. If we want to represent the polarization state of a wide spectral distribution (as we do in our ray tracer), we must approximate the spectrum by quasi-monochromatic distributions and maintain a coherence matrix for each. We can validly represent these distributions in our calculations by a single wavelength. As a consequence of the linear superposition principle for light waves, we represent the total polarization state of independently superimposed light waves that have the same quasi-monochromatic spectral distribution by the sum of their individual coherence matrices.

In summary, we perform two different approximations. We approximate quasi-monochromatic distributions by their peak wavelengths, which is accurate because the distributions are so narrow; and we approximate the visible spectrum by several wavelength samples. The issues related to the latter approximation are exactly the same as those encountered by Cook and Torrance,[6] and we can deal with aliasing in color space by increasing the number of samples.

For a light wave reflecting off a surface, define the *specular plane of incidence* to be the plane determined by the vectors in the directions of the incident and reflected light waves, as Figure 1 shows. We can simplify the computation of the reflected polarization state of a light wave specularly reflected from a surface: We choose a coordinate system for the coherence matrices so x is parallel to the specular plane of incidence, y is perpendicular to the specular plane of incidence, and positive z is in the direction of light propagation. Suppose that after the first reflection of a light wave, a subsequent reflection occurs so the specular planes of incidence of the two reflections do not coincide. The coherence matrix representing the polarization state of the light wave after its first reflection now needs to be represented with respect to new x-y coordinate axes that are parallel and perpendicular to the second specular plane of incidence.

The well-defined algebraic properties of the coherence matrix enable us to represent very elegantly in our ray tracer the polarization state of a light wave with respect to different specular planes of incidence. Assume that the second specular plane of incidence is rotated counterclockwise by the angle θ from the first specular plane of incidence (looking from the positive z direction)—rotated about the axis representing the incident direction of the light wave just before its second reflection (equivalently, the axis in the direction of the light wave just after the first reflection). According to Wolf and Born,[1,2] the transformation between the representation of the coherence matrix from the first specular plane of incidence to the second is given by

$$\begin{pmatrix} J'_{xx} & J'_{xy} \\ J'_{yx} & J'_{yy} \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} J_{xx} & J_{xy} \\ J_{yx} & J_{yy} \end{pmatrix} \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \quad (2)$$

where the primed components are for the coherence matrix with respect to the second specular plane of incidence.

The Torrance-Sparrow reflectance model

After its first use in computer graphics by Blinn,[7] Cook and Torrance[6] summarized and demonstrated in more detail the Torrance-Sparrow reflectance model. Here we summarize only the model's salient features. The model assumes that every surface has a microscopic level of detail that consists of a statistically large number of perfectly smooth planar microfacets. These microfacets are oriented according to a probability distribution function that depends only on the angle between the normal to each microfacet and the normal to the surface. Thus the surface is assumed to be isotropically rough.

As Figure 2 shows, according to the Torrance-Sparrow model the total reflected radiance is produced from the sum of a specular component and a diffuse component. The specular component consists of light waves that reflect off of a single, planar, surface microfacet. The diffuse component consists of light waves that have had multiple reflections between microfacets and/or penetrate into the outermost layer

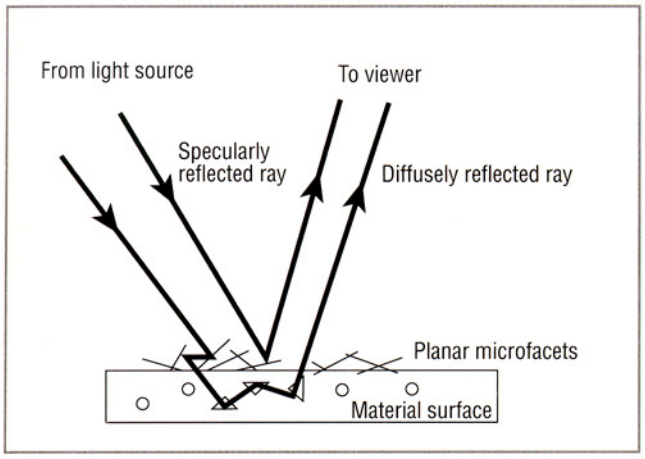

Figure 2. Specular and diffuse reflection from a microscopically rough material surface.

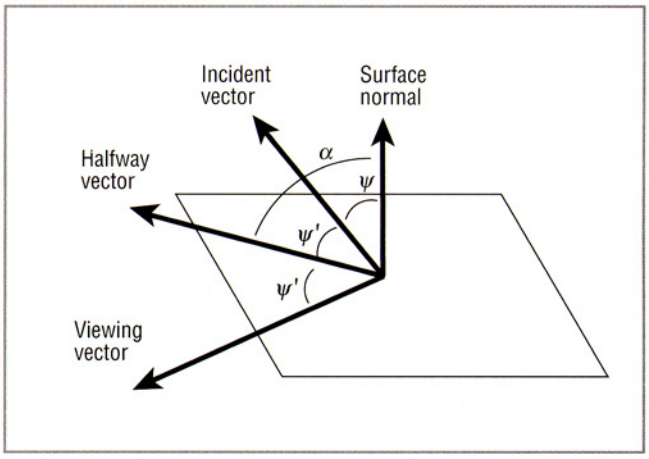

Figure 3. Angular arguments for the Torrance-Sparrow model.

of the material surface and then emanate again. Using Cook and Torrance's notation,[6] the total reflected radiance is directly proportional to

$$sR_s + dR_d, \quad s + d = 1$$

where R_s and R_d are functions for the specular and diffuse reflected radiant intensities, respectively, and s and d are real and nonnegative. The specular and diffuse component radiant intensity functions are given by

$$R_s = \frac{F(\Psi', \eta)G}{\cos \Theta} P(\alpha), \quad R_d = \cos \Psi$$

where the angular arguments are depicted in Figure 3, and G and $P(\alpha)$ are given by

$$G = \min\left\{1, \frac{2 \cos \alpha \cos \Theta}{\cos \Psi'}, \frac{2 \cos \alpha \cos \Psi}{\cos \Psi'}\right\}$$

$$P(\alpha) \propto \frac{1}{m^2 \cos^4 \alpha} \exp\left(-\frac{\tan^2 \alpha}{m^2}\right)$$

Here G represents the geometric attenuation factor. The function $P(\alpha)$ is the probability distribution function for the orientation of the planar microfacets, for which we use the Beckmann distribution.[8] (Cook and Torrance[6] were the first to apply the Beckmann distribution to the Torrance-Sparrow model.) The term m is the root-mean-square slope value for the microfacet surface normals. The roughness of the surface increases with m.

The Fresnel reflection coefficients

The function $F(\Psi', \eta)$ in the specular reflectance component for the Torrance-Sparrow model is the *Fresnel reflection coefficient* and represents the proportional attenuation of the intensity of an incident light wave with respect to specular reflection off a perfectly smooth surface such as a planar microfacet. The parameter $\eta = n - \kappa i$ is the complex index of refraction for the material surface. The coefficient of extinction, κ, is zero for dielectrics and nonzero for conductors. The term n is the simple index of refraction. Siegel and Howell gave a detailed derivation of the Fresnel reflection coefficient, representing it as the linear superposition of the Fresnel coefficients for perpendicular and parallel components of the electric field with respect to the specular plane of incidence.[9] That is,

$$F(\Psi', \eta) = uF_\perp(\Psi', \eta) + vF_\parallel(\Psi', \eta) \quad u, v \geq 0,\ u + v = 1$$

$$F_\perp(\Psi', \eta) = \frac{a^2 + b^2 - 2a \cos \Psi' + \cos^2 \Psi'}{a^2 + b^2 + 2a \cos \Psi' + \cos^2 \Psi'} \tag{3}$$

$$F_\parallel(\Psi', \eta) = \frac{a^2 + b^2 - 2a \sin \Psi' \tan \Psi' + \sin^2 \Psi' \tan^2 \Psi'}{a^2 + b^2 + 2a \sin \Psi' \tan \Psi' + \sin^2 \Psi' \tan^2 \Psi'} F_\perp(\Psi', \eta)$$

where

$$2a^2 = [(n^2 - \kappa^2 - \sin^2\Psi')^2 + 4n^2k^2]^{1/2} + n^2 - \kappa^2 - \sin^2\Psi' \quad (4)$$

$$2b^2 = [(n^2 - \kappa^2 - \sin^2\Psi')^2 + 4n^2k^2]^{1/2} - (n^2 - \kappa^2 - \sin^2\Psi')$$

Through the Fresnel reflection coefficients, different polarization states of incident light influence the magnitude of the specularly reflected radiance. The Fresnel coefficients also dictate how most materials polarize light upon reflection. For dielectrics, the relative magnitude of the perpendicular Fresnel coefficient to the parallel Fresnel coefficient is usually very high, thus making dielectrics fairly good polarizers of initially unpolarized light. At the Brewster angle, the parallel Fresnel coefficient is zero for a dielectric, so light reflected at this angle of incidence is linearly polarized.

Metals do not polarize light nearly as well upon reflection. However, for some common metals (for example, copper) the Fresnel coefficients are relatively sensitive to changes in incident wavelength, making specularly reflected light for different wavelengths significantly different in polarization state. This accounts for observed color variations of highlights from these metals when reflected against dielectrics. The polarizing properties of dielectrics redistribute the intrinsic color distribution of the light reflected from metals.

Torrance-Sparrow model with coherence matrices

Coherence matrices were first incorporated into the Torrance-Sparrow reflection model to compute surface orientation from pairs of images taken with different orientations of a polarizing filter.[3] Such incorporation is based on two fundamental assumptions about the polarization states of the individual diffuse and specular components of reflection:

1. Because many random microfacet and/or internal reflections of light waves make up the diffusely reflected light wave, the diffusely reflected light wave is always completely unpolarized, regardless of the incident polarization state.
2. The mutually orthogonal component wave functions that make up the specular component of reflection are proportionally attenuated according to the Fresnel reflection coefficients from the original component wave functions making up the incident light wave.

These physical assumptions are consistent with Torrance's assumptions.[5]

Completely polarized light is characterized by a coherence matrix with a vanishing determinant.[1,2] On the other hand, completely unpolarized light is characterized by a coherence matrix of the form

$$\frac{1}{2}I_0\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad (5)$$

where I_0 is the radiant intensity. The normalization constant of 1/2 allows the trace of the matrix to add to I_0.

Assumption 1 above implies that the coherence matrix for the diffuse component of reflection is proportional to the matrix of Expression 5. Assumption 2 leads to the derivation of the coherence matrix for the specular component of reflection. The derivation of the specular coherence matrix begins with a theorem stating that any coherence matrix can be represented as the unique sum of an unpolarized coherence matrix and a completely polarized coherence matrix. That is, for the incident coherence matrix

$$\begin{pmatrix} J_{xx}^{(i)} & J_{xy}^{(i)} \\ J_{yx}^{(i)} & J_{yy}^{(i)} \end{pmatrix} = \begin{pmatrix} A^{(i)} & 0 \\ 0 & A^{(i)} \end{pmatrix} + \begin{pmatrix} B^{(i)} & D^{(i)} \\ D^{(i)\star} & C^{(i)} \end{pmatrix} \quad (6)$$

Here $B^{(i)}C^{(i)} - D^{(i)}D^{(i)\star} = 0$, meaning that the determinant of the second additive coherence matrix is 0 and therefore represents a completely polarized light wave. The (i) superscript denotes that this is the incident polarization state. From assumption 2 above, the diagonal terms for the resulting reflected coherence matrix are attenuated by the corresponding Fresnel reflection coefficients parallel and perpendicular respective to the specular plane of incidence. The part of the specular reflected coherence matrix resulting from the unpolarized component of the incident coherence matrix is therefore given by

$$\begin{pmatrix} F_{\parallel}A^{(i)} & 0 \\ 0 & F_{\perp}A^{(i)} \end{pmatrix} \quad (7)$$

The specular reflection of a completely polarized light wave from a perfectly smooth microfacet should have no depolarization effect, so the determinant of the coherence matrix of the reflected completely polarized component should also be zero. Because of the attenuation produced from the Fresnel coefficients, the diagonal components for the part of the specular

component arising from the completely polarized portion of the incident light wave will be $F_{\|} B^{(i)}$ and $F_{\perp} C^{(i)}$ respectively. We now have both $B^{(i)}C^{(i)} = D^{(i)}D^{(i)\star}$ and $F_{\|} F_{\perp} B^{(i)}C^{(i)} = D^{(r)}D^{(r)\star}$, where the (r) superscript denotes the components for the reflected wave. The general complex solution for $D^{(r)}$ follows as

$$D^{(r)} = \sqrt{F_{\|} F_{\perp}} D^{(i)} e^{i\delta}$$

for any real δ. The resulting coherence matrix for the part of the specular component arising from the completely polarized portion of the incident light wave is

$$\begin{pmatrix} F_{\|} B^{(i)} & \sqrt{F_{\|} F_{\perp}} D^{(i)} e^{i\delta} \\ \sqrt{F_{\|} F_{\perp}} D^{(i)\star} e^{-i\delta} & F_{\perp} C^{(i)} \end{pmatrix} \quad (8)$$

The phase difference δ is in fact physically significant because there is a retarding action on the specular component of reflection. Siegel and Howell[9] showed that there is a phase retardance of $\delta_{\|}$ and $\delta_{\perp}$ of the parallel and perpendicular wave components, with respect to the specular plane of incidence, such that

$$\tan\delta_{\|} = \frac{2\cos\Psi'[(n^2 - k^2)b - (2nk)a]}{(n^2 + k^2)^2\cos^2\Psi' - (a^2 + b^2)}$$

$$\tan\delta_{\perp} = \frac{2b\cos\Psi'}{\cos^2\Psi' - a^2 - b^2}$$

where n and k are the components of the complex index of refraction, and a and b are as in Equations 4. The net phase retardation is $\delta = \delta_{\|} - \delta_{\perp}$.

Adding together the coherence matrices in Equations 7 and 8, and recalling the initial decomposition of the incident coherence matrix in Equation 6, the coherence matrix for the specular component of reflection is

$$\begin{pmatrix} F_{\|} J^{(i)}_{xx} & \sqrt{F_{\|} F_{\perp}} J^{(i)}_{xy} e^{i\delta} \\ \sqrt{F_{\|} F_{\perp}} J^{(i)}_{yx} e^{-i\delta} & F_{\perp} J^{(i)}_{yy} \end{pmatrix}$$

From the additivity of coherence matrices for independent light waves, the polarization state of the reflected light wave is the sum of the coherence matrices for the diffusely and specularly reflected light waves, respectively. This is

$$\frac{sR_s}{F_{\|} J^{(i)}_{xx} + F_{\perp} J^{(i)}_{yy}} \begin{pmatrix} F_{\|} J^{(i)}_{xx} & \sqrt{F_{\|} F_{\perp}} J^{(i)}_{xy} e^{i\delta} \\ \sqrt{F_{\|} F_{\perp}} J^{(i)}_{yx} e^{-i\delta} & F_{\perp} J^{(i)}_{yy} \end{pmatrix} \quad (9)$$

where the traces of the individual reflection component coherence matrices have been normalized to the magnitudes of the diffuse and specular components, respectively.

The parameters u and v used to determine $F(\Psi', \eta)$ in the function R_s are determined by

$$u = \frac{J^{(i)}_{xx}}{J^{(i)}_{xx} + J^{(i)}_{yy}}, \quad v = \frac{J^{(i)}_{yy}}{J^{(i)}_{xx} + J^{(i)}_{yy}}$$

Expression 9 is valid in the coordinate system defined in the section "Polarization and the coherence matrix" with respect to the specular plane of incidence. Observe that for initially unpolarized light, the polarization state of the reflected light is significantly altered when the magnitude of the specular component of reflection is in turn significant.

Implementation and performance

In this section we assume that you are familiar with Whitted's classical ray-tracing algorithm,[10] which Glassner described in detail.[11]

The intensity and polarization of each light ray is represented by a set of coherence matrices, one for each wavelength sampled. The ray representation also includes a *reference normal*, perpendicular to the ray, which together with the light-wave propagation direction defines the coordinate system of each of the coherence matrices with respect to the current specular plane of incidence. We choose a right-handed coordinate system, with the positive z direction being the direction of light propagation, and the positive y axis lying along the reference normal.

To determine the color values for an individual pixel, our algorithm traces a ray backward from the eye through the pixel, in the opposite direction of actual light propagation. The ray's coherence matrices are initially set to zero, and the reference normal is initialized to an arbitrary direction orthogonal to the light ray. Call the ray that we are tracing the primary ray. For each unoccluded light source in the scene, we consider the ray cast from that light source, incident upon the point of intersection, and reflected along the direction of the primary ray. The coherence matrices for the incident light ray and the primary ray are then transformed with Equation 2 so that their reference normals lie perpendicular to the specular plane of incidence.

After these transformations, we can use Equation 9 to find the coherence matrices for the reflected ray. Since we transformed the primary ray's coordinate system to the coordinate system of the reflected ray, their coherence matrices can now be summed to-

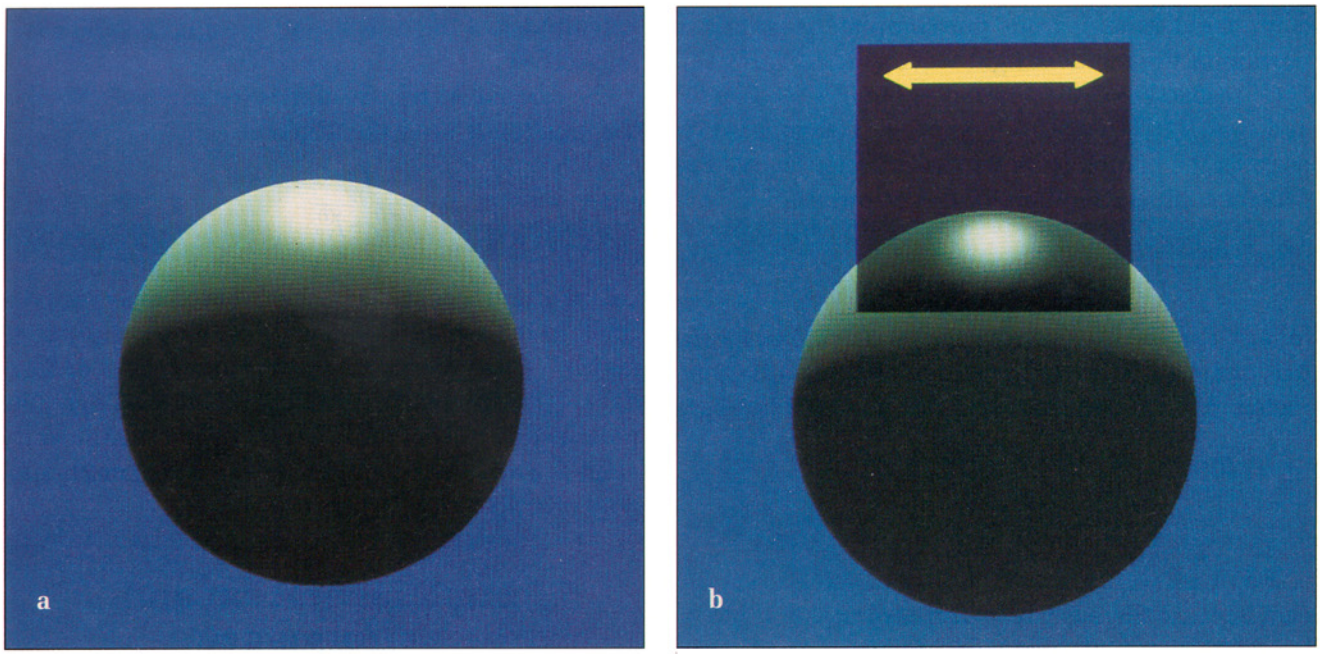

Figure 4. Computer-generated images of a sphere and a polarizing filter.

gether to find the new coherence matrices of the primary ray, including this reflected component.

Our algorithm repeats the steps above—including the transformation steps—for each light source, summing the contributions of the light sources directly visible at the intersection into the primary ray's coherence matrices.

As in conventional ray tracing, to render mirror reflections of objects in the scene we recursively cast a ray from the point of intersection, whose angle of incidence equals its angle of reflection with respect to the surface normal. This newly cast ray becomes the primary ray for a new level of ray-tracing recursion, and we calculate its coherence matrices. When done, we return from this level of recursion and treat the ray (no longer the primary ray) as a ray cast directly from a light source. That is, we transform the coherence matrices of this incident ray and those of the primary ray to the specular plane of incidence, find the coherence matrices of the reflected ray, and add these matrices to those of the primary ray.

After completing the recursion and calculating the coherence matrices for a ray cast from the eye through a pixel, we use these matrices to compute the pixel's color. The trace of the coherence matrix provides the intensity of a single quasi-monochromatic band. Since we sample at red, green, and blue wavelengths, we choose an RGB triplet for the pixel by taking values proportional to the traces of coherence matrices for quasi-monochromatic distributions centered at the wavelengths of our display device's red, green, and blue phosphors. As described in the section on polarization and the coherence matrix, we can very accurately approximate each quasi-monochromatic distribution by its peak wavelength, and we approximate the visible spectrum with a set of wavelength samples. For better color accuracy, we can sample more wavelengths in the visible spectrum and, using techniques described by Cook and Torrance[6] and Meyer and Greenberg,[12] transform the spectral distribution into an RGB value.

The calculation of $F_\perp$, $F_\parallel$, and δ in Equation 9 is potentially quite expensive, perhaps requiring many computations for each ray-object intersection. To make the calculations more tractable, we build lookup tables for these values during a preprocessing phase. Since each of these functions depends on Ψ', n, and κ, a single linear lookup table for each will not suffice. To complicate matters further, n and κ depend on wavelength as well as parameters of the material. For each combination of material and wavelength that we sample, a lookup table is constructed for $F_\perp$, $F_\parallel$, and δ. During the ray-tracing lighting calculations, we use table lookup and interpolation to find values for these functions.

Incorporating polarity calculations into a ray-tracing lighting model introduces extra computational overhead. However, the major bottleneck in ray-trac-

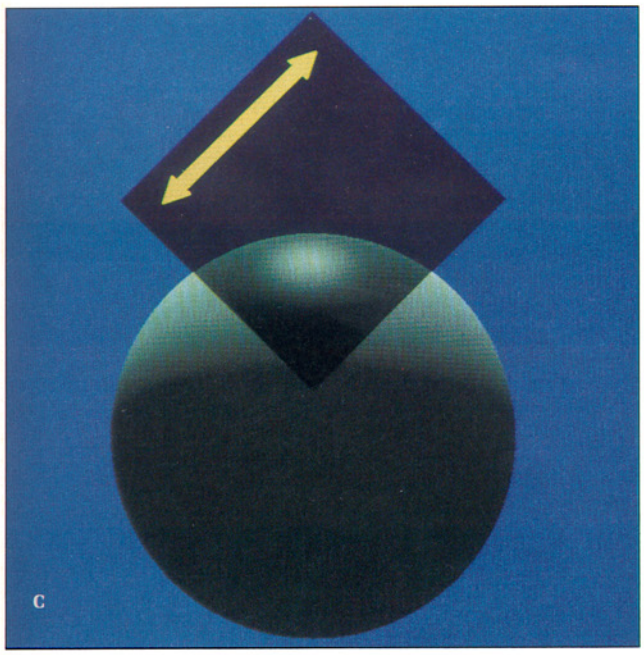

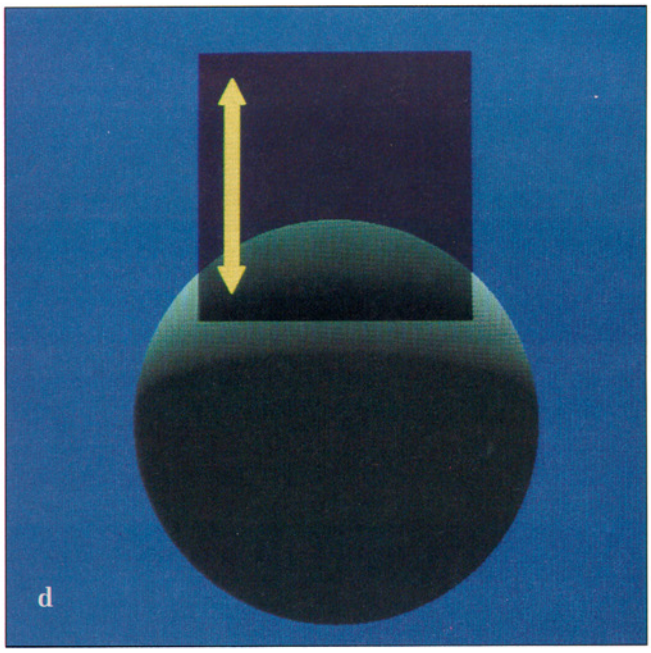

ing implementations usually is the ray-object intersection calculations. As scenes become more complex, ray-object intersections become exceedingly more costly in relation to the time spent calculating scene lighting. We have found that a ray tracer with a polarization lighting model takes approximately twice the time to render simple scenes with few objects as a ray tracer with the Torrance-Sparrow model alone. As scenes increase in complexity, the execution time of the polarization model approaches that of the unmodified Torrance-Sparrow algorithm.

Simple scenes with a polarizing filter

Figure 4a shows a computer-generated image of a dielectric sphere illuminated by a single light source with reflected radiance values predicted by the Torrance-Sparrow model. Since this image was generated using ray tracing with polarization parameters, the reflected light rays contain polarization state information as well as radiant intensity. This is shown by simulating the effect of a polarizing filter placed between the viewer and the specular highlight at different orientations. The polarizing filter passes only the portion of the electric-field component of light waves that is parallel to the arrow. As the sequence in Figures 4b, 4c, and 4d shows, the specular highlight disappears as the arrow tends toward a vertical position. Except for uniform attenuation through the polarizing filter, the diffuse component of reflected light remains invariant.

The center of the specular highlight completely disappears because at this point on the sphere light cast from an unpolarized light source, which is placed above and a little behind the sphere, is incident at the Brewster angle (defined in the section "Polarization and the coherence matrix"). Thus the only component of the electric field that is reflected is normal to the specular plane of incidence, and the reflected light wave is linearly polarized in this direction. This complete polarization of the specular component of a reflected light wave occurs regardless of the incident polarization state. The simulated polarizing filter in Figure 4d cancels out the center of the specular highlight because at this location the specular highlight is composed only of reflected light linearly polarized perpendicular to the arrow. Most of the rest of the specular highlight is also canceled out because it is almost completely linearly polarized in the same direction.

Observe the difference in contrast between the specular highlight and the background of diffuse reflectance between Figures 4a and 4b. Since the diffuse component is unpolarized, one half of the diffuse component is canceled by the linear polarizer, whereas almost none of the specular highlight is canceled while the polarizer is aligned with the electric

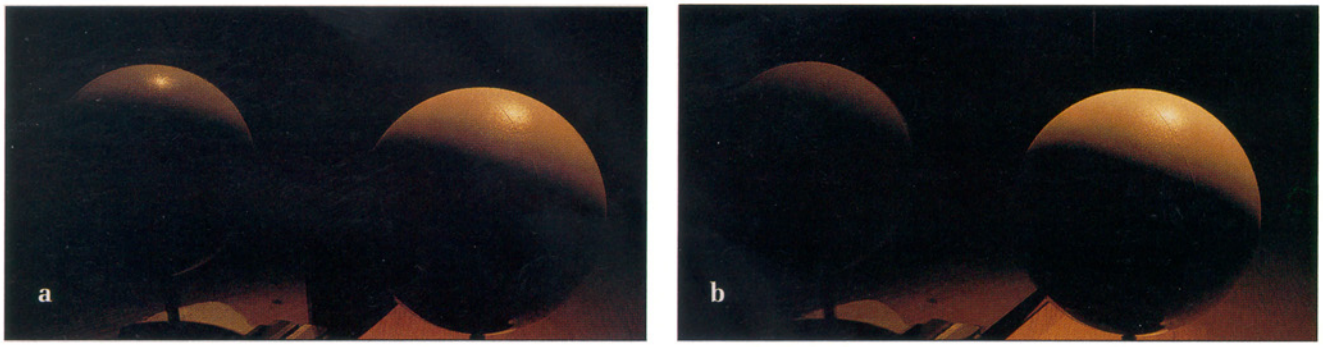

Figure 5. Pictures of a real dielectric sphere being reflected in (a) a metal mirror and (b) a dielectric mirror. These are *not* computer-generated images.

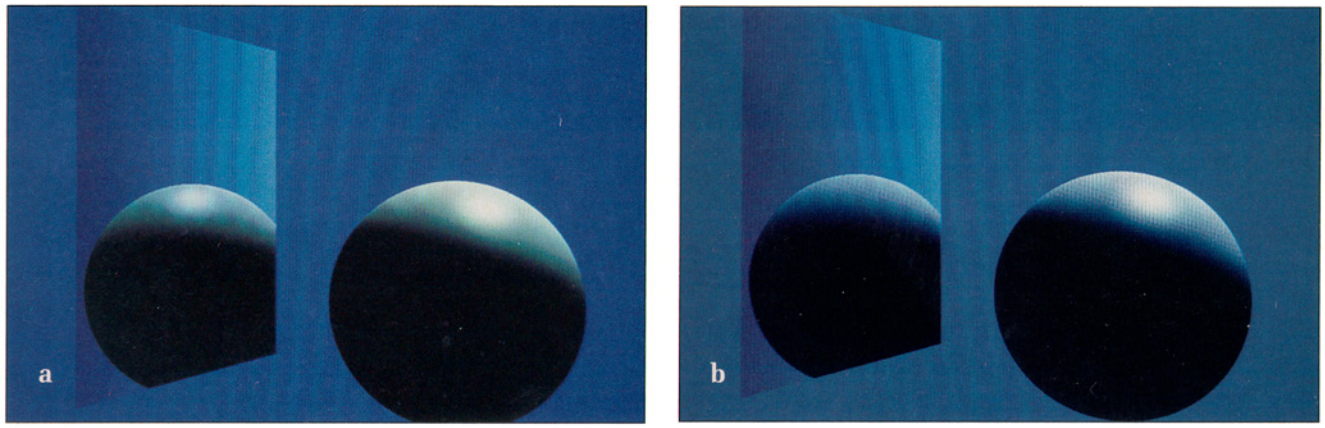

Figure 6. Computer-generated images simulating the scene in Figure 5b, generated with (a) the traditional Torrance-Sparrow model and (b) the polarization model.

field of the specular component. Because it is completely unpolarized, the diffuse reflected component remains invariant with respect to the polarizing filter's orientation.

Scenes with multiple reflections

As described in the section on the Fresnel reflection coefficients, the reflected radiant intensity from dielectrics is more sensitive than conductors to the incident polarization state of light. Thus, polarization effects are usually more pronounced in pictures containing multiple reflections off of dielectrics. However, we will show later that pictures containing metallic specular reflections can also have interesting polarization effects.

Figures 5a and 5b are photographs of a real dielectric sphere being reflected in (a) a metal mirror and (b) a dielectric mirror. The only evidence of the mirrors is in the reflections that they cast. The geometry of the two scenes is identical; only the materials are different. In both cases, the spheres are illuminated by an unpolarized light source. Note, however, that the metallic mirror's reflection of the sphere contains a distinct specular highlight not visible in the reflection in the dielectric mirror. This is an effect of polarization. The configuration of the scene is such that a light wave reflecting off the center point of the specular highlight on the sphere reflected in the glass is incident on the sphere at its Brewster angle. After reflecting off the sphere, the light wave is incident on the mirror at the Brewster angle with respect to the dielectric mirror's index of refraction. The light then strikes the image plane. Also, the specular planes of incidence for the two reflections (one off the sphere, the other off of the mirror) are perpendicular.

In terms of the discussion above, the specular component of the reflected light wave off the sphere—at the center of the specular highlight—is linearly polarized normal to the first plane of incidence. After re-

Figure 7. Metallic and dielectric vases resting on a reflecting dielectric surface, generated with (a) the traditional Torrance-Sparrow model and (b) the polarization model.

flecting off the glass in Figure 5b, the specular component of the reflected light wave must be linearly polarized normal to the second plane of incidence. But the first and second planes of incidence are perpendicular, causing the specular component to be completely annihilated after the second reflection. Reflected radiant intensity of light off of metals is less sensitive to incident polarization, and metallic materials do not have a Brewster angle that completely linearly polarizes reflected light. Therefore, unlike the glass mirror, the metallic mirror in Figure 5a reflects a specular highlight of the sphere.

As polarization does not have a significant effect on the reflected intensities of Figure 5a, you would expect traditional ray-tracer lighting models to be able to render this scene accurately. Indeed, we simulated this scene with ray tracers using both the traditional Torrance-Sparrow model and the polarization model, and they yielded similar results. We also rendered a simulation of Figure 5b with both ray tracers. Figure 6a is a computer simulation of Figure 5b using the Torrance-Sparrow lighting model, and Figure 6b is a simulation of the same figure using the polarization model. Note that Figure 6b accurately shows the extinction of the sphere's specular highlight reflected in the mirror. Figure 6a inaccurately contains a reflection of the specular highlight, showing that ray tracing without polarization parameters can be unrealistic.

Figures 7a and 7b show two vases—one metal (left), the other dielectric (right)—reflected in a glass table using an unpolarized light source. Both images are computer generated. Figure 7a is rendered without polarization and Figure 7b is rendered with polarization. The difference is in the specular highlights reflected by the table. The reflected highlight for the dielectric vase in Figure 7b, taking polarization parameters into account, is brighter because the plane of incidence for reflection off the vase at the specular

Figure 8. Multiple reflections of a dollar bill in a dielectric slab and a dielectric cube, generated with (a) the traditional Torrance-Sparrow model and (b) the polarization model.

highlight is parallel to the plane of incidence for the subsequent reflection off the glass table. In Figure 7a the light reflected off the dielectric vase at the specular highlight is assumed to be still unpolarized. Therefore, more light is absorbed into the glass than in Figure 7b, where the light reflected off the vase has a larger component perpendicular to the plane of incidence.

The reflection of the specular highlight off the metal vase in the glass appears in noticeably different hues in Figures 7a and 7b. The Fresnel coefficients for conductors are in general significantly sensitive to changes in wavelength. Therefore, the specularly reflected coherence matrices should significantly differ between the discrete wavelengths selected to approximate the polychromatic distribution of the light source.

Figures 8a and 8b show the multiple reflections of the completely diffuse surface of a dollar bill in a dielectric cube mounted on a horizontal dielectric slab. Figure 8a does not account for polarization, while Figure 8b does. The relative darkening in Figure 8b of the multiple reflections of the dollar bill is due again to double reflections near the Brewster angle for the dielectric surfaces with perpendicular planes of incidence, similar to those for the specular highlight in Figures 5b and 6b.

Conclusion

Recall the approximation made in the section "Torrance-Sparrow model with coherence matrices" that the diffuse component of reflection is completely unpolarized. As Wolff discussed,[13] this approximation breaks down for inhomogeneous dielectrics at surface points where viewing is nearly orthogonal to the surface normal (e.g., at occluding contours). In fact, the diffuse reflection component can be significantly partially polarized from these points. This polarization reflection theory[13] explains this phenomenon using the Fresnel reflection coefficients. We would like to examine the consequences this may have on the realistic rendering of certain scenes.

We would also like to examine polarization's effects under extended light sources, color shifts in the reflections of metallic objects, and the polarizing effects of transparent media. Our initial investigation of transparent media suggests that transmission has a significantly smaller effect on polarization than does reflection. ■

Acknowledgments

Lawrence Wolff was supported in part by an IBM graduate fellowship. Thanks to Kicha Ganapathy for

originally suggesting this application to computer graphics, and thanks to Steve Feiner and Terry Boult for reviewing drafts and giving helpful suggestions. Hewlett-Packard generously provided us with computers to generate the images.

References

1. E. Wolf, "Coherence Properties of Partially Polarized Electromagnetic Radiation," *Il Nuovo Cimento*, Vol. 13, No. 6, Sept. 1959, pp. 1,165-1,181.
2. M. Born and E. Wolf, *Principles of Optics*, Sixth Ed., Pergamon Press, Oxford, UK, 1981.
3. L.B. Wolff, "Surface Orientation from Polarization Images," *Optics, Illumination, and Image Sensing for Machine Vision II*, SPIE vol. 850, SPIE, Bellingham, Wash., 1987, pp. 110-121.
4. K.E. Torrance and E.M. Sparrow, "Theory for Off-Specular Reflection from Roughened Surfaces," *J. Optical Soc. Am.*, Vol. 57, No. 9, Sept. 1967, pp. 1105-1114.
5. K.E. Torrance, "Theoretical Polarization of Off-Specular Reflection Peaks," *ASME J. Heat Transfer*, Vol. 91, No. 2, May 1969, pp. 287-290.
6. R.L. Cook and K.E. Torrance, "A Reflectance Model For Computer Graphics," *Computer Graphics* (Proc. SIGGRAPH), Vol. 15, No. 3, Aug. 1981, pp. 307-316.
7. J.F. Blinn, "Models of Light Reflection For Computer Synthesized Pictures," *Computer Graphics* (Proc. SIGGRAPH), Vol. 11, No. 2, July 1977, pp. 192-198.
8. P. Beckmann and A. Spizzichino, *The Scattering of Electromagnetic Waves From Rough Surfaces*, Macmillan, New York, 1963.
9. R. Siegel and J.R. Howell, *Thermal Radiation Heat Transfer*, McGraw-Hill, New York, 1981.
10. T. Whitted, "An Improved Illumination Model For Shaded Display," *CACM*, Vol. 23, No. 6, June 1980, pp. 343-349.
11. A. Glassner, ed., *An Introduction to Ray Tracing*, Academic Press, London, 1989.
12. G.W. Meyer and D.P. Greenberg, "Perceptual Color Spaces for Computer Graphics," *Computer Graphics* (Proc. SIGGRAPH), Vol. 14, No. 3, July 1980, pp. 254-261
13. L.B. Wolff, *Polarization Methods in Computer Vision*, PhD thesis, Columbia University, New York, in preparation.

Lawrence B. Wolff is a doctoral student in computer science at Columbia University. His research interests are in computer vision, machine vision, computer graphics, and differential geometry. Before beginning his graduate studies, he worked in industry for five years in image processing, pattern recognition, and computer graphics.

Wolff received his BS in mathematics and physics from Yale University in 1981 and his MS in computer science from Columbia University in 1987. He is a member of IEEE and AAAI.

David J. Kurlander is a doctoral student in computer science at Columbia University. His research interests include graphical editing, realistic image synthesis, and user interfaces. He has spent summer months working on graphics-related projects at the Pentagon, Hewlett-Packard Labs, and Xerox PARC.

Kurlander received his BA in applied mathematics from Harvard College in 1985 and his MS in computer science from Columbia University in 1987.

The authors can be reached at the Computer Science Department, Columbia University, New York, NY 10027.