UCRL-JC-126062
PREPRINT

Plane Parallel Radiance Transport for Global Illumination in Vegetation

N. Max, C. Mobley, B. Keating, E.-H. Wu

This paper was prepared for submittal to the
Special Interest Group on Computer Graphics
Los Angeles, CA
August 3-8, 1997

January 5, 1997

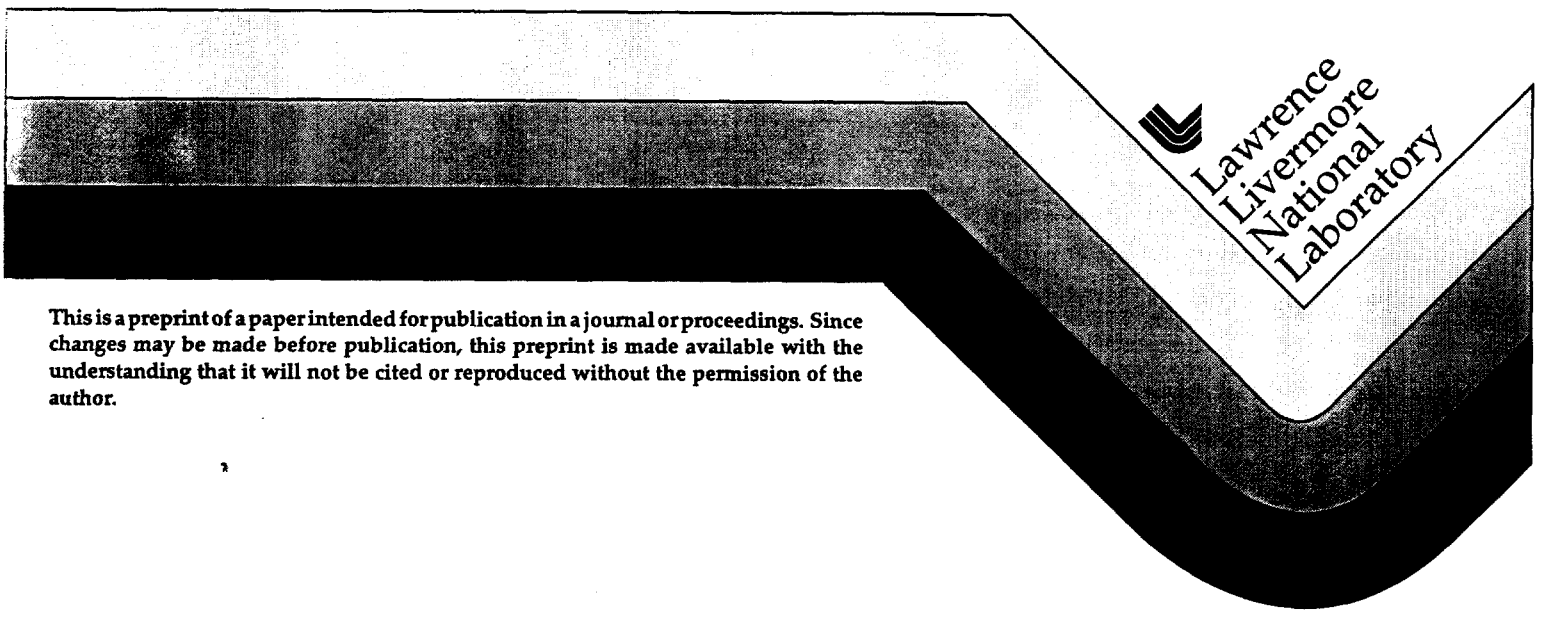

Plane Parallel Radiance Transport for Global Illumination in Vegetation

Nelson Max

Curtis Mobley

Brett Keating

En-Hua Wu

Abstract

This paper applies plane parallel radiance transport techniques to scattering from vegetation. The leaves, stems, and branches are represented as a volume density of scattering surfaces, depending only on height and the vertical component of the surface normal. Ordinary differential equations are written for the multiply scattered radiance as a function of the height above the ground, with the sky radiance and ground reflectance as boundary conditions. They are solved using a two-pass integration scheme to unify the two-point boundary conditions, and Fourier series for the dependence on the azimuthal angle. The resulting radiance distribution is used to precompute diffuse and specular "ambient" shading tables, as a function of height and surface normal, to be used in rendering, together with a z-buffer shadow algorithm for direct solar illumination.

1. Introduction

Radiosity [Cohen] and radiance methods [Ward] were developed to accurately account for global illumination in realistic rendering. The radiosity at a surface or volume exits diffusely, to produce equal intensity in all directions, while radiance varies with viewing angle. Computing the radiosity at all positions necessary for realistic shading usually involves the iterative solution of a large system of linear equations, and for radiance, the number of variables is even larger [Imm]. Alternative Monte Carlo ray tracing methods [Ward, Kaji] need a large number of rays to converge to an accurate solution. For very complex geometries, such as a forest with thousands of different leaf surfaces, both methods become impractical.

Sillion [Sill95a, Sill95b] has suggested grouping objects hierarchically into volumes, and applying hierarchical volume radiosity methods [Bhate]. Volume radiance methods also exist [Sieg, Max95a], but the storage and computation become unwieldy if one must determine the radiance in each direction at each position in 3D space.

One can reduce the problem to 1D by assuming that the radiance depends angularly on the

light flow direction on the unit sphere, but positionally only on z, the height above the ground, and not on x or y. If the directions are discretely sampled, this plane parallel assumption reduces the partial differential equations (PDEs) for radiance transport [Sieg, Max95b] to ordinary differential equations (ODEs) in z alone. As we show below, this permits direct solution by the Runge-Kutta method, instead of the iterative methods usually needed to solve the general PDEs. The volume scattering must also be plane parallel, so we assume a random collection of vegetation, whose leaf area and surface normal densities also depend only on z. This is a significant limitation, prohibiting isolated trees, but is a good approximation for dense vegetation.

Radiation transport in leaf canopies has been studied for the purposes of calculating the possibilities for photosynthesis, and also for interpreting remote sensing data. Myneni, Ross, and Asrar [Myn] present a detailed survey of this field, reviewing the measurement of leaf area, surface normal densities, BRDFs, BRTFs, and the various mathematical methods of solving for the radiance.

For the plane parallel problem of scattering within a leaf, Hanrahan and Krueger [Hanr] have applied Monte-Carlo techniques, and Ma *et al.* [Ma] have applied discrete ordinate techniques, to find the bidirectional reflection (or transmission) distribution function (BRDF or BTDF) of the leaf, and Ma *et al.* [Ma] and Woolley [Wool] have measured these functions experimentally. Their results could be used as input to our calculations, but at present, we use only perfectly diffuse reflection and transmission, and Phong [Bui] specular reflection.

As pointed out by Li and Strahler [Li] and Gerstl and Borel [Gerstl], the ODE solution simulates light attenuation through a continuous distribution of infinitesimal scatterers, and cannot correctly model the "hot spot", when the viewing direction is near the solar illumination direction, so that finite-sized trees and leaves hide their own shadows. This hot spot is well known in remote sensing, and is also visible around the shadow of an airplane, but will not be present in our images because we use a geometric shadow calculation for greater realism. We subtract the attenuated direct solar illumination from our ODE solution, and in our rendering, we use a z-buffer shadow algorithm [Will] for the direct solar illumination. The ODE solution is used only for the attenuated direct sky illumination, and for indirect illumination that has been scattered at least once. This additional illumination is best appreciated on surfaces in shadow from direct sunlight.

The next section gives the background and derivation of the radiance transport equations, and the following one describes the table-based shading methods to efficiently use its radiance output. Section 4 explaines the modeling and rendering system, and sections 5 and 6 give results and conclusions, respectively. The Appendix describes the computation of the necessary scattering phase functions.

2. Plane parallel radiance transport equations.

The plane parallel radiance transport equations were originally developed for planetary and

stellar atmospheres [Amb, Chan] and have more recently been applied to underwater illumination [Mobl94]. In these applications, the scattering is off randomly located and randomly oriented molecules, phytoplankton, dust particles, droplets, ... etc., so the phase function, giving the angular distribution of scattered light, depends only on the angle between the incident and scattered rays. However, the basic mathematics applies to any phase function. We have therefore modified the Hydrolight code [Mobl95] for the surface reflection case where the distribution of surface normals is rotationally symmetric around the vertical axis, which occurs for rotationally symmetric plants, or random rotational placement of arbitrary plants. As described below, this permits certain efficiencies in the solution.

The simplest plane parallel solutions [Milne, Edd] found the light flow in only two representative directions, up and down. This "two-stream" approximation was applied to scattering from paint, modeled as pigment particles in a transparent carrier, by Kubelka and Munk [Kub] and has also been applied to scattering in plant canopies [Nils, Allen]. More powerful computers now allow more accurate approximation of the angular variation of the radiance, either as (*a*) spherical harmonics, (*b*) a combination of azimuthal Fourier series and discrete sampling of the polar angle, or (*c*) division of the unit sphere into angular bins which are discrete in both directions. All these methods lead to PDEs in x, y, and z for the discretised radiance, or its basis expansion coefficients. Discussions of these methods, as well as equation (1) below, can be found in [Sieg], [Myn], and [Max95b]. Here, as in [Mobl94], we use method (*c*), but eventually switch to azimuthal Fourier series. A complete discussion of the mathematics behind the Hydrolight code can be found in [Mob94], of which the following is a short summary.

The general equation of radiance transport is

$$\frac{dI(X, \omega)}{ds} = -\kappa(X, \omega)I(X, \omega) + \kappa(X, \omega)\int_{4\pi} ar(X, \omega', \omega)I(X, \omega')d\omega' \qquad (1)$$

where X is 3D position, ω and ω' are ray directions on the unit sphere, $I(X, \omega)$ is the radiance at position X flowing in direction ω, $r(X, \omega', \omega)$ is the scattering phase function from direction ω' to direction ω (representing in our case both reflection from and transmission through the leaves), the integral is over the 4π solid angle of incoming directions ω' at X, κ is the extinction or beam attenuation coefficient representing absorption plus scattering, and a is the albedo, the fraction of the attenuation that is scattered rather than absorbed. Usually κ is assumed to be isotropic, that is, independent of ω, but in our application, it represents geometric occlusion, from the polygon surface area, projected normal to ω, per unit volume at X. This depends on ω unless the distribution of polygon normals is completely random.

In spherical coordinates, θ is the polar angle, measured from the vertical z axis, and φ is the azimuthal angle, measured from the x axis in the x-y plane. We allow our distribution of surface normals to depend on θ, but not on φ. This means that if $\omega = (\theta, \varphi)$ and $\omega' = (\theta', \varphi')$, then $r(X, \omega', \omega)$ depends only on X, θ, θ', and $|\varphi - \varphi'|$, as long as the BRDF and BTDF of the leaves satisfy these same "isotropy" assumptions in a coordinate system with the z axis along the surface normal.

We divide the unit sphere into a number of discrete direction bins, $2m$ in θ and n in φ. However we do not subdivide by φ the first and last θ bins, but instead keep them joined into a pair of polar caps. This gives a total of $(2m - 2)\ n + 2$ bins. We precompute a square matrix of this size, for each θ_N orientation of the normal to an opaque surface, and each wavelength band, by averaging the volume scattering phase function r over sampled pairs of rays, one from the bin for ω and one for ω', and over the range of φ_N angles for the surface normal. Actually, by the "isotropy" assumption above, many of the matrix entries are identical, and need only be computed and stored once. The analytic integration over the 2π range of φ_N variation for the surface normal, necessary to determine κ for occluding surfaces, and r for diffusely reflecting or transmitting surfaces, or specularly reflecting surfaces, is given in the Appendix. We make separate r matrices for diffuse reflection, diffuse transmission, and specular reflection.

We also tabulate, for each height z, the density of surface area per unit volume with normals in each θ_N bin (but collecting all φ_N). We tabulate the leaf surface density separately from that of the trunk, branch, stem, or twig polygons, since the leaf color is different, the leaves transmit light as well as reflect it, and the leaf polygons are two sided, reflecting light both from their front and back surfaces. We can then approximate $\kappa(X, \omega)$ and the product $\kappa(X, \omega)\ a\ r(X, \omega, \omega')$ at any position by combining the precomputed quantities or matrices, weighted by the tabulated area densities, and the albedos for each color band.

In our plane parallel case, the intensity I depends positionally only on z, and we represent it by a row vector of intensities $I_i(z)$, where i varies over the direction bins. Let μ_i be the cosine of the representative θ for the i^{th} direction bin, so that the ds in equation (1) becomes $dz\ /\ \mu_i$. Let $r_{ij}(z)$ be the bin to bin scattering matrix $\kappa(X, \omega)\ a\ r(X, \omega, \omega')$, whose computation was just described, modified so that $r_{ii}(z)$ contains the extinction term $-\kappa(X, w)$. Then the integral-differential equation (1) reduces to a linear system of ordinary differential equations

$$\mu_i \frac{dI_i(z)}{dz} = \sum_j I_j(z) r_{ji}(z) \tag{2}$$

It is convenient to separate the vector of intensities into two smaller row vectors, $I_u(z)$, representing the upwards flow, for direction bins in the upper hemisphere, and $I_d(z)$ representing the downwards flow. (We choose $2m$ bins, so that no bins straddle the equator.) Then we replace equation (2) by the two coupled vector differential equations

$$\begin{aligned} \frac{dI_u(z)}{dz} &= I_u(z)\tau_{uu}(z) + I_d(z)\rho_{du}(z) \\ -\frac{dI_d(z)}{dz} &= I_u(z)\rho_{ud}(z) + I_d(z)\tau_{dd}(z)\ . \end{aligned} \tag{3}$$

The τ's and ρ's are matrices which are multiplied by the row vectors I_u and I_d. For example, the entries in ρ_{du} are the terms $r_{ji}\ /\ |\mu_i|$, for directions j in the lower hemisphere and i in the upper. (Note that the signs on the derivatives are reversed in [Mobl94], which measures z downwards from the water surface, instead of upwards from the ground.)

Let h be the height above ground of the top of the vegetation layer. Then we know one boundary condition: $I_d(h)$ is determined by the illumination flowing from the sun and sky. The semiempirical model of Harrison and Coombes [Harr] is used to compute the angular distribution of the sky radiance. The clear-sky irradiance model of Gregg and Carder [Gregg] is used to compute the direct and diffuse spectral irradiances (on a horizontal plane) for clear skies. The model of Kasten and Czeplak [Kast] is used to adjust the Gregg and Carder irradiances when the cloud fraction is greater than 0.25. The (adjusted) irradiances are used to set the magnitude of the spectral sky radiance.

This is not enough information to specify initial conditions at $z = h$, however, because we do not know $I_u(h)$. Our second boundary condition is given by the BRDF of the ground at $z = 0$. We have used a fallen-leaf texture map to determine this BRDF, but any other method, such as the model and empirical measurements of [Irons] would do as well. The bin-to-bin matrix F approximation to this BRDF gives the boundary condition $I_u(0) = I_d(0)\, F$.

Thus we have "two-point" boundary conditions, one at $z = h$ and one at $z = 0$. The shooting and relaxation methods recommended for this situation [Press] both require complicated iterations. However, it is possible to determine $I_u(h)$ by solving an auxiliary differential equation, and then solve equation (3) as an initial value problem.

Let $F(z)$ be the BRDF matrix for reflection from the ground and all the vegetation in the vertical range $[0, z]$, including the effects of all multiple scattering (reflection, transmission, and occlusion). Then

$$I_u(z) = I_d(z)\, F(z)\,. \tag{4}$$

We can derive a differential equation for $F(z)$ by taking the derivative of equation (4)

$$\frac{dI_u(z)}{dz} = \frac{dI_d(z)}{dz}F(z) + I_d(z)\frac{dF(z)}{dz}$$

and substituting the right hand sides of equations (3) to get

$$I_u(z)\tau_{uu} + I_d(z)\rho_{du} = (\, - I_u(z)\rho_{ud} - I_d(z)\tau_{dd})\, F(z) + I_d(z)\frac{dF(z)}{dz}\,.$$

Substituting $I_d(z)F(z)$ for $I_u(z)$ and rearranging terms, we get

$$I_d(z)\frac{dF(z)}{dz} = I_d(z)\,(\,\rho_{du} + F(z)\tau_{uu} + \tau_{dd}F(z) + F(z)\rho_{ud}F(z)\,)\,.$$

Since the row vector $I_d(z)$ is arbitrary, we must have

$$\frac{dF(z)}{dz} = \rho_{du} + F(z)\tau_{uu} + \tau_{dd}F(z) + F(z)\rho_{ud}F(z)\,. \tag{5}$$

This is a non-linear ordinary differential equation, called a Riccati equation, for the matrix $F(z)$, with initial conditions $F(0) = F$, the known BRDF of the ground. Thus it can be integrated incrementally upwards from $z = 0$ to $z = h$. Once $F(h)$ is known, we can determine $I_u(h)$ from equation (4), and then integrate equations (3) incrementally downwards from $z = h$ to $z = 0$. For our numer-

ical integration, we used the adaptive Runge-Kutta routines ODEINT, RKQS, and other auxiliary routines from the Numerical Recipes package [Press]. For equation (5), this involves multiplication of large square matrices, of size equal to the number of direction bins $N = (m - 1)\, n + 1$ in a hemisphere, taking time $O(N^3)$. This can become very slow for a fine subdivision of the direction sphere.

Therefore we expand the row vectors $I_u(z)$ and $I_d(z)$ as a finite Fourier series (*i. e.*, a discrete Fourier transform) in ϕ. The resulting Fourier coefficients depend continuously on z and discretely on θ. The matrices ρ, τ, and F are also represented by their Fourier coefficients in ϕ. As described above, the scattering terms in the coefficients of equations (2) and (3) are assumed to depend directionally only on θ, θ', and $|\phi - \phi'|$. Under this symmetry assumption, it can be shown [Mobl94] that the system (2) or (3) of $2N$ coupled differential equations splits into $n+1$ independent groups of $2m$ equations. These are much easier to solve, in time $O(nm^3)$. The solution Fourier coefficients are converted back to the functions $I_u(z)$ and $I_d(z)$, which are tabulated at selected values of z, and interpolated in z to find the radiance for shading purposes.

To decrease the data storage and the integration over the angle bins necessary to get the entries in the τ and ρ matrices, we currently assume that $\tau_{uu} = \tau_{dd}$ and $\rho_{du} = \rho_{ud}$. This will be the case if the two-sided translucent leaves look the same from either side (under equivalent illumination), and if the stems, branches, and trunks, bounded by the one-sided opaque polygons, are cylindrical, so that polygons with one normal are balanced by nearby polygons with the opposite normal.

3. Shading

In the simulations reported here, we used $2m = 20$ latitude bins and $n = 24$ longitude bins, for a total of $2N = 434$ bins, 217 of which will be in the hemisphere above a surface to be shaded. If we performed shading by multiplying each of these 217 radiance terms by the appropriately evaluated BRDF, and summing these contributions, the shading would be very slow. Instead, we use a Phong BRDF, with perfectly diffuse reflection plus a specular term which depends only on the angle between the viewing direction and the perfect mirror reflection of the incident ray. Special precomputation speedups then apply.

The Phong shading model [Bui] takes the form

$$f(L, V, N, \lambda) \cos\theta = d(\lambda) \cos\theta + s\, \{R \cdot V\}^c,$$

where f is the BRDF, L is the direction "to the light source" opposite to the incoming ray direction ω, V is the outgoing ray direction, N is the surface normal, λ is the wavelength of the light, θ is the angle between L and N, d is a wavelength dependent diffuse reflection coefficient, s is a wavelength independent specular reflection coefficient, R is the mirror reflection direction, the 180° rotation of L about N, c is an arbitrary exponent, and the special dot product $\{R \cdot V\}$ is $R \cdot V$ if $R \cdot V$ is

positive, and 0 otherwise.

To determine the surface intensity $S(V, \lambda)$ in the viewing direction V we must integrate the contribution of the incoming radiance $I(\omega, \lambda) = I(-L, \lambda)$ for L in the 2π hemisphere above the surface, and thus ω in the hemisphere below it.

$$S(V, \lambda) = \int_{2\pi} I(\omega, \lambda) f(-\omega, V, N, \lambda) \cos\theta d\omega$$

$$= d(\lambda) \int_{2\pi} I(\omega, \lambda) \cos\theta d\omega + s \int_{2\pi} I(\omega, \lambda) \{R \cdot V\}^c d\omega$$

Note that the first integral, for perfectly diffuse reflection, depends only on the surface normal N. As pointed out by Miller and Hoffman [Mill], it can be precomputed and tabulated for a set of sampled directions N, and interpolated during shading. The same tables can be used with the opposite normal to determine the diffuse transmission through translucent leaves.

The second integral can be treated by a method proposed by Cabral, Max, and Springmeyer [Cabr]. Let Q be the 180° rotation of V about N. Then $\{R \cdot V\} = \{L \cdot Q\} = \{-\omega \cdot Q\}$, so the second integral can be replaced by

$$\int_{2\pi} I(\omega, \lambda) \{-\omega \cdot Q\}^c d\omega$$

If we imagine for a moment that light can be spwcularly reflected even though it comes from the back of the surface, the 2π unit hemisphere in this integral is replaced by the 4π unit sphere. Then the integral depends only on Q, and can also be precomputed. For the large exponents c appropriate for glossy surfaces, the error introduced by extending the integral to the whole unit sphere will only be significant for glancing viewing rays, when Q is near the boundary of the hemisphere of integration. In this case brighter highlights are expected anyway from the Fresnel reflection law. In fact, there is no reason why s should be constant. It could easily depend on $N \cdot L$ according to Fresnel's law, adjusted to compensate for the brightening mentioned above, and also on the wavelength λ. However, we have kept it independent of λ, because the uncolored highlights seem appropriate for the waxy outer coating of leaves.

As mentioned in the introduction, we have subtracted the contribution of the direct sunlight, as continuously attenuated by the effect of the extinction $\kappa(z, \omega)$ on its direct path through the vegetation, before storing the solution $I(z, \omega)$. The shading process adds back the effect of the unattenuated direct sunlight, using the Phong BRDF in the usual way for a discrete parallel light source, only if the surface point is not in shadow.

4. Modeling and Rendering.

For modeling the vegetation, we used the hierarchical modeling and rendering system described in [Max96]. This image-based rendering system is able to resample precomputed multi-layer z-buffer images of hierarchical subparts of a plant, and combine them onto a multi-layer z-

buffer for a new viewpoint. The size of the subparts to be resampled depends on the viewpoint location. Parts closer to the viewpoint are taken from higher-detail images farther down in the hierarchical model, and only the closest parts are scan converted from the original texture/transparency mapped polygons. This allows extremely complex and extensive vegetation to be rendered in a reasonable time.

Each z-buffer layer has a surface normal, a surface diffuse reflectivity color, a subpixel coverage mask, a z value, and a leaf bit. This is enough information to perform as a post-process the shading described in the previous section, including a z-buffer shadow algorithm, at each of the multiple layers. (Details of the shadow algorithm will be given elsewhere.) The subpixel masks are used, as in the A-buffer algorithm [Carp] to composite these shading results into an antialiased color value for the pixel. If the leaf bit is set, the shading will take into account the diffusely transmitted light from the back of the leaf.

The geometry from the vegetation model is used to create the input for the radiance transport phase. In order to compute the extinction coefficient $\kappa(z, \omega)$, we must know at each z and each θ_N the surface area density per unit volume with normal of polar angle θ_N. We divide the (z, θ_N) ranges into a two dimensional array of bins, and make an initial traversal of the hierarchical model, using the detailed polygonal description throughout. Each polygon is clipped to the (z, θ_N) ranges it overlaps, and the clipped areas are summed into the appropriate (z, θ_N) bins. In the Appendix we show how to compute the contribution of the area in these bins to the extinction coefficient $\kappa(z, \omega)$ and scattering coefficient $r(z, \omega, \omega')$.

5. Results

The maple tree of [Max96] was processed to give the density of both leaf and non leaf surfaces per unit volume at 10 different normal orientation bins, and 10 different height intervals. The sky radiance model was for a clear sky, with the sun at 30° from the vertical, and 35° to the right of the viewing directions of figures 1 and 2. The radiance was calculated for a hemisphere divided at 10° latitude and 15° longitude intervals, *i. e.*, $m = 20$ and $n = 24$. The radiance transport calculation took 7 hours to find the radiance in each of the $N = 217$ direction bins at each of the 10 heights, and for each of three (red, green, and blue) wavelength bands. Computing the shading by the methods of section 3 added only one minute to the post-processing of a 640 by 480 image.

Figure 1 shows a maple grove, with 28 rotated and translated copies of the basic maple tree. The view is near the bottom of the tree canopy, where little direct sunlight has penetrated. The shading on the diffuse grey test sphere shows a greenish tint from the radiance which has filtered through and multiply scattered in the canopy. In contrast, the shading on the test sphere in figure 2, which is near the top of the canopy, shows the actual grey color. Note the sun flecks from the shadow computation, visible on the tree branches, and the yellow-green glow of the diffuse transmission through the back-lit leaves, which was included in the radiance transport calculations.

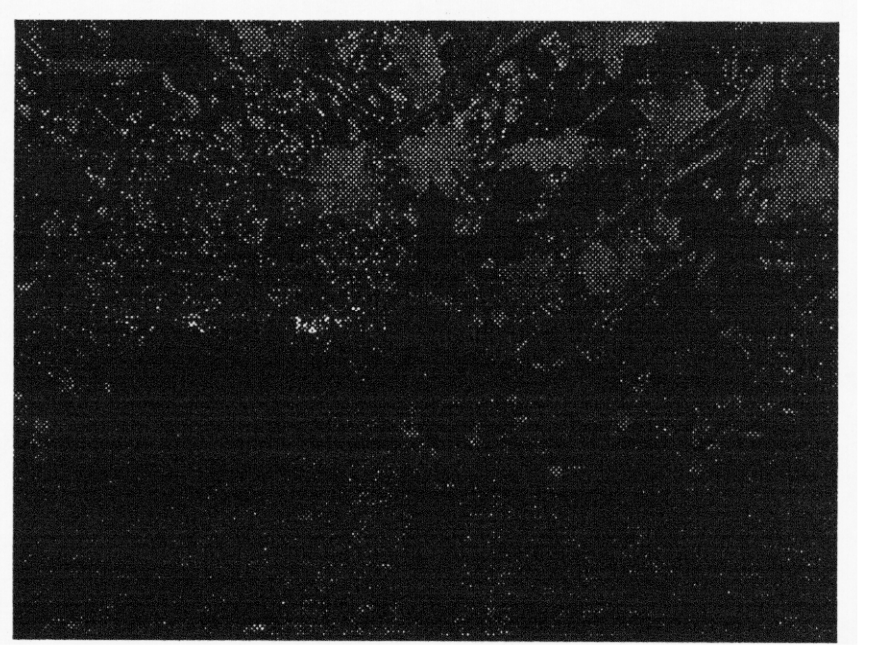

Figure 1. View of maple grove from near bottom of tree canopy.

Figure 2. View of maple grove from near top of tree canopy.

6. Summary and conclusions

The calculations described here are efficient for the following reasons.

(*a*) The computation of the scattering phase functions $r(z, \omega, \omega')$, described in the Appendix, is done once per θ_N bin, and integrated over each pair of ω and ω' bins. Because $r(z, \omega, \omega')$ is rotationally symmetric, depending only on θ, θ' and $|\phi - \phi'|$, only a fraction of these coefficients need be computed and stored; the rest are accessed by an appropriate indexing scheme. These coefficients are found once and for all, even before any vegetation model is specified.

(*b*) The plane parallel assumption reduces the dimensionality of the radiance transfer problem, replacing PDEs by ODEs. The ODE solution is further simplified by replacing a two-point boundary value problem by two initial value problems.

(*c*) The first of these initial value problems, equation (5) for the matrix F, needs to be solved only when the vegetation model changes, and therefore the leaf area and normal densities change. The easier equations (3) for the vectors I_u and I_d must be solved each time the sky illumination changes. Once tabulated, the radiance solution can be reused if the viewpoint moves, or if small objects not contributing to the radiance transport simulation move.

(*d*) Again due to the rotational symmetry of $r(z, \omega, \omega')$, the large coupled system of ODEs separates into several smaller mutually independent systems of equations, after a discrete Fourier transform in ϕ.

(*e*) The shading during rendering is based on tables precomputed from the radiance solution, which also do not change when the viewpoint moves.

(*f*) The hierarchical model lets very complex vegetation to be rendered efficiently from precomputed z-buffer images. Since these precomputed images contain surface color and normal information, and are shaded in a post-process after resampling, they are unaffected when the radiance solution changes.

7. Future work

We hope to extend these calculations to underwater vegetation, taking into account the refraction by the waves on the water surface, as in [Mobl94].

It would be interesting to animate vegetation blowing in the wind. For simple waving of leaves, we could assume that the area density distribution of surface normals did not change. Then the radiance solution could be reused, although the hierarchical system of precomputed images could not. However, for more violent motion, such as grass or grain bending in wind gusts, our rotational symmetry assumptions are violated; when grass stems are coherently blown away from the vertical, the distribution of surface normals is no longer independent of ϕ. In this case, the effi-

ciencies listed in points (*a*) and (*d*) of the previous section no longer apply. The full matrices r, τ, and ρ must be calculated, and used in solving one large system of coupled differential equations. This more general case is also necessary for vertical vegetation growing on a slanted hill, because the parallel planes along which the radiance is constant are no longer horizontal. This general case could also handle multiple scattering in hair and fur, when the hairs are not perpendicular to the skin.

Acknowledgments

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. It originated during the visit of author N. M. to the Institute of Software, Academia Sinica, Beijing, supported by the Natural Science Foundation of China. We wish to thank Pat Hanrahan and Peter Shirley for helpful conversations, and Wenchang Wang, Jing Ren, and Honju Li for debugging, system, and logistic help in Beijing. The development of the Hydrolight 3.0 was supported by the Environmental Optics Program of the Office of Naval Research under a contract to author C. M.

Appendix: Calculating the scattering phase function.

Suppose we have a random collection of small polygons, of total area 1, with normals $N = (\theta_N, \phi_N)$, where θ_N is fixed, and ϕ_N is uniformly distributed in $[0, 2\pi]$. In order to find the contribution of these polygons to $\kappa(z, \omega)$, we need to determine their projected area normal to a viewing direction $\omega = V = (\phi_V, \theta_V)$. Also, in order to find the contribution to $r(z, -L, V)$, we need to know the average intensity of this projected area, when illuminated by a light source in direction $L = (\theta_L, \phi_L)$, with unit flux per unit area normal to the beam, assuming the surface diffuse reflectivity k_d is also 1. In the integro-differential equation (1), the light scattering takes place in an infinitesimal slab of thickness ds, where the probability of random polygon overlap approaches zero, so we can neglect mutual occlusion or shadowing of the polygons.

In Cartesian coordinates,

$$N = (\sin\theta_N \cos\phi_N, \sin\theta_N \sin\phi_N, \cos\theta_N)$$

and

$$V = (\sin\theta_V \cos\phi_V, \sin\theta_V \sin\phi_V, \cos\theta_V).$$

Since the total polygon area is 1, the area with normals between ϕ_N and $\phi_N + d\phi$ is $d\phi / (2\pi)$, so we integrate $(N \cdot V)\, d\phi / (2\pi)$ to find the projected area perpendicular to ω. There are two cases to be considered. For one-sided polygons which form the boundaries of solid structures, we integrate $N \cdot V$ only over the range R of ϕ_N where $N \cdot V > 0$. For two-sided leaves, we must add the integral of $-N \cdot V$ over the complementary range $[0, 2\pi] - R$ where $N \cdot V < 0$. These ranges are bounded by

the ϕ_N values where $N \cdot V = 0$. Solving this equation for ϕ_N,

$$N \cdot V = \sin\theta_N \sin\theta_V (\cos\phi_N \cos\phi_V + \sin\phi_N \sin\phi_V) + \cos\theta_N \cos\theta_V = 0$$

$$\cos(\phi_N - \phi_V) = -\cot\theta_N \cot\theta_V.$$

Let $\phi_R = |\phi_N - \phi_V| = \{\cos^{-1}\}(-\cot\theta_N \cot\theta_V)$, where $\{\cos^{-1}\}(x) = \cos^{-1}(x)$ if $-1 \le x \le 1$, and is 0 if $x > 1$, or π if $x < -1$. Then, interpreting ϕ_N mod 2π, the range R is $[\phi_V - \phi_R, \phi_V + \phi_R]$. So for one-sided polygons, the projected area A_R is

$$A_R = \int_R (N \cdot V)\frac{d\phi_N}{2\pi}$$

$$= \int_{\phi_V - \phi_R}^{\phi_V + \phi_R} \sin\theta_N \sin\theta_V \cos(\phi_N - \phi_V)\frac{d\phi_N}{2\pi} + \int_{\phi_V - \phi_R}^{\phi_V + \phi_R} \cos\theta_N \cos\theta_V \frac{d\phi_N}{2\pi}$$

$$= \frac{1}{\pi}(\sin\theta_N \sin\theta_V \sin\phi_R + \phi_R \cos\theta_N \cos\theta_V).$$

For two-sided polygons, we add

$$\int_{[0, 2\pi] - R} (-N \cdot V)\frac{d\phi_N}{2\pi}$$

$$= \frac{1}{\pi}(\sin\theta_N \sin\theta_V \sin\phi_R - (\pi - \phi_R)\cos\theta_N \cos\theta_V)$$

to get the total

$$\frac{1}{\pi}(2\sin\theta_N \sin\theta_V \sin\phi_R + (2\phi_R - \pi)\cos\theta_N \cos\theta_V).$$

To compute the average diffuse reflection intensity for one-sided opaque polygons, we must integrate over R the diffuse reflection $\{N \cdot L\}$, weighted by the projected area $N \cdot V$, and divide by the integral of the weights, which is the A_R calculated above. The region S where $N \cdot L \le 0$ is, by the above analysis, the interval $[\phi_L - \phi_S, \phi_L + \phi_S]$, where $\phi_S = \{\cos^{-1}\}(-\cot\theta_N \cot\theta_L)$. The product $(N \cdot V)(N \cdot L)$ is a polynomial in $\sin\phi_N$ and $\cos\phi_N$, with coefficients involving sines and cosines of θ_N, θ_V, and θ_L, and can be straightforwardly integrated over the intersection of the intervals R and S (which may be empty).

For two-sided translucent leaf polygons, the intersection of R and S and their respective complements divides the unit circle for ϕ_N up into ranges of four types: (1) front, front lit (*i. e.* the front surface of the leaf is visible to V, and also to L), (2) back, front lit, (3) front, back lit, and (4) back, back lit. The identification of these ranges and their types is found from a table, based on the topology of overlap of the intervals R and S on the unit circle. Over each range, $|(N \cdot V)(N \cdot L)|$ can be straightforwardly integrated.

In our current implementation, we assume that the front and back surfaces of the leaves have identical optical properties, so that ranges of types (1) and (2) contribute to the diffuse reflection, and ranges of types (3) and (4) contribute to the diffuse transmission. For each wavelength band, we allow separate coefficients for diffuse reflection and diffuse transmission, and in the images shown here, the diffuse transmission is yellower. The analysis above shows how the scattering for perfectly diffuse reflection and transmission can be integrated analytically over certain ϕ_N ranges, but any computed or table-based measured BRDF or BTDF could be integrated numerically over the same ranges.

The above computations for κ and diffuse r are integrated in θ_N over each angle bin band $S_i = [\theta_i, \theta_{i+1}]$ into which the surface normals are accumulated, using Simpson's rule, and then divided by the solid angle of the band. For the contribution of specular reflection, however, it is easier to consider the whole band S_i from the start, rather than first integrating over ϕ_N.

The Phong specular term $\{V \cdot R\}^c$ is constant on small circles on the unit sphere, centered at R. We show below that for a fixed viewing direction V, the contours of this specular term, as a function of the normal N, are, to first order, small ellipses centered at $H = (V+L)/|V+L|$. Blinn has suggested an alternate specular term [Blinn] of the form $\{N \cdot H\}^c$, consistent with a rotationally symmetric distribution of local microfacet normals H, centered on the global surface normal N. For a fixed viewing direction, the contours of this term are exactly circles, centered at H. We first show how to integrate the Blinn term over the band S_i, and then modify the method for the Phong term.

We flatten the unit sphere near H onto the (x, y) tangent plane T_H at H, by a map p_H, which takes great circle arclength distance from H into $r = \sqrt{x^2 + y^2}$, and precompute and tabulate

$$Q(t) = \int_0^t dy \int_{-\infty}^{\infty} \cos^c(r) \frac{\sin(r)}{r} dx .$$

The factor $\sin(r)/r$ converts differential area on the plane to differential area on the sphere. With appropriate additions and subtractions of $Q(t)$, we can find the integral over the band between any two parallel lines in the (x, y) plane, and thus, to first order, the integral for N in the band between any two large circles on the sphere, such as the band S_i. We then multiply by $V \cdot H$ to account for the forshortening of the polygon area, as viewed from V, assuming the normal N of any polygon showing a significant highlight is very close to H, and divide by the solid angle of the band S_i.

For the Phong term, consider the map f which takes a normal N to the corresponding reflection vector $R = 2(N \cdot L)N - L$. For fixed V and L, we need to integrate $\{V \cdot R\}^c = \{V \cdot f(N)\}^c$ as N varies over S_i. Since f is smooth, the differential df of f at H maps the tangent plane T_H at H to the tangent plane T_W at $W = f(H) = 2(H \cdot L)H - L$, and the composition of maps $p_W^{-1} \cdot df \cdot p_H$ is a first order approximation to f. Therefore, we can map S_i to T_W by $df \cdot p_H$, approximate the image as the region between two parallel lines, and integrate $\cos^c(r)$ there using the tabulated $Q(t)$. We must then divide by the Jacobian $J(f)$ of f to account for the change of the integration variables, and

multiply by $V \cdot H$ as above to account for forshortening. In spherical coordinates with the north pole at L, the formulas for f are straightforward: θ is doubled and ϕ is unchanged. The Jacobian of f turns out to be 4 $V \cdot H$. (See [Rens].) Since df is linear, df^{-1} maps circles to ellipses, as mentioned above.

References

[Allen] William Allen and Arthur Richardson, "Interaction of Light with a Plant Canopy," Journal of the Optical Society of America, Vol. 58, No. 8 (1968) pp. 1023 - 1028.

[Amb] V. A. Ambarzumanian, "Diffuse reflection of light in a foggy medium," Compt. Rend. (Doklady) Acad. Sci. USSR, Vol. 38 (1943) p. 229.

[Bhate] Neeta Bhate, "Application of Rapid Hierarchical Radiosity to Participating Media," in "ATARV-93: Advanced Techniques in Animation, Rendering, and Visualization" (B. Özgüç and V. Akman, eds.), Bilkent University (July 1993) pp. 43 - 53.

[Biu] Phong Bui-Tuong, "Illumination for Computer Generated Pictures," Communications of the ACM, Vol. 18. No. 6 (1975) pp. 311 - 317.

[Blinn] James Blinn, "Models of light reflection for computer synthesized pictures," Computer Graphics Vol. 11 No. 2 (1987, Siggraph '87 Conference Proceedings) pp. 192 - 198.

[Cabr] Brian Cabral, Nelson Max, and Rebecca Springmeyer, "Bidirectional Reflection Functions from Surface Bump Maps," Computer Graphics Vol. 21 No. 4 (July 1987, Siggraph '87 Conference Proceedings) pp. 273 - 281.

[Carp] Loren Carpenter, "The A-Buffer, An Antialiased Hidden Surface Method," Computer Graphics Vol. 18 No. 3 (1984; Siggraph '84 Proceedings) pp. 103 - 108.

[Chan] Subrahmanyan Chandrasekhar "Radiative Transfer", Oxford University Press, 1950.

[Cohen] Michael Cohen and John Wallace, "Radiosity and Realistic Image Synthesis," Academic Press Professional, Boston, 1993.

[Edd] Arthur Eddington "The Internal Combustion of Stars," Dover, New York, 1959.

[Gerstl] Seigfried Gerstl and Christoph Borel, "Principles of the Radiosity Method versus Radiative Transfer for Canopy Reflectance Modeling," IEEE Transactions Geoscience and Remote Sensing, Vol. 30 No. 2 (1992) pp. 271 - 275.

[Gregg] W. W. Gregg and K. L. Carder, "A simple spectral solar irradiance model for cloudless maritime atmospheres," Limnology and Oceanography Vol. 35 No. 8 (1990) pp. 1657 - 1675.

[Hanr] Pat Hanrahan and Wolfgang Krueger, "Reflection from Layered Surfaces due to Subsurface Scattering," ACM Computer Graphics Proceedings, Annual Conference Series (1994) pp. 165 - 174.

[Harr] A. W. Harrison and C. A. Coombes, "An opaque cloud cover model of sky short wavelength radiance," Solar Energy, Vol. 41 No. 4 (1988) pp. 387-392.

[Imm] David Immel, Michael Cohen, and Donald Greenberg, "A Radiosity Method for Non-Diffuse Environments," Computer Graphics Vol. 20 No. 4 (August 1986) pp. 133 - 142.

[Irons] James Irons, Gaylon Campbell, John Norman, David Grahm, and William Kovalick, "Prediction and Measurement of Soil Bidirectional Reflectance," IEEE Transactions Geoscience and Remote Sensing, Vol. 30 No. 2 (1992) pp. 249 - 256.

[Kaji] James Kajiya "The Rendering Equation," Computer Graphics Vol. 20, No. 4 (August 1986) pp. 165 - 174.

[Kast] Fritz Kasten and Gerhard Czeplak, "Solar and terrestrial radiation dependent on the amount and type of cloud," Solar Energy Vol. 24 (1980) pp. 177-189.

[Kub] Paul Kubelka and F. Munk, "Ein beitrag zur optik der farbenstriche," Zurich Tech. Physik, Vol 12 (1931) p. 543.

[Li] Xiaowen Li and Alan Strahler, "Geometric-Optical Bidirectional Reflectance Modeling of the Discrete Crown Vegetation Canopy: Effect of Crown Shape and Mutual Shadowing," IEEE Transactions Geoscience and Remote Sensing, Vol. 30 No. 2 (1992) pp. 276 - 291.

[Ma] Qinglan Ma, Akira Ishimaru, Phillip Phu, and Yasuo Kuga, "Transmission, Reflection, and Depolarization of an Optical Wave for a Single Leaf," IEEE Transactions Geoscience and Remote Sensing, Vol. 28 No. 5 (1990) pp. 865 - 872.

[Max95a] Nelson Max, "Efficient Light Propagation for Multiple Anisotropic Volume Scattering," in "Photorealistic Rendering Techniques" (G. Sakas, P. Shirley, and S. Müller, eds.) Springer Verlag, Heidelberg (1995) pp. 87 - 104.

[Max95b] Nelson Max, "Optical Models for Direct Volume Rendering," IEEE Transactions on Visualization and Computer Graphics, Vol. 1 No. 2 (1995) pp. 99 - 108.

[Max96] Nelson Max, "Hierarchical Rendering of Trees from Precomputed Multi-Layer Z-Buffers'", in "Rendering Techniques '96", X Pueyo and P. Schröder (eds.), Springer, NewYork (1996).

[Mill] Gene Miller and Robert Hoffman, "Illumination and reflection maps: Simulated objects in simulated and real environments," Advanced Image Synthesis Course Notes, Siggraph 1984 Conference.

[Milne] F. A. Milne, "Thermodynamics of the Stars," in Handbuch der Astrophysik, Vol. 3, Springer Verlag, Berlin (1930) pp. 65 - 255.

[Mobl94] Curtis Mobley, Light and Water: Radiative Transfer in Natural Waters, Academic Press, San Diego (1984).

[Mobl95] Curtis Mobley, "Hydrolight 3.0 Users' Guide," SRI International, 333 Ravenswood Avenue, Menlo Park, CA 94025 (1996).

[Myn] Ranga Myneni, Juhan Ross, and Ghassem Asrar, "A review on the theory of photon transport in leaf canopies," Agricultural and Forest Meteorology, Vol. 45 (1989) pp. 1 - 153.

[Nils] T. Nilson and U. Peterson, "A Forest Canopy Reflectance Model and a Test Case," Remote Sensing Environments, Vol. 37 (1991) pp. 131 - 142.

[Press] William Press, Brian Flannery, Saul Teukolsky, and William Vetterling, "Numerical Recipes: The Art of Scientific Computing," Cambridge University Press, Cambridge (1988).

[Rush] Holly Rushmeier and Ken Torrance, "The Zonal Method for Calculating Light Intensities in the Presence of a Participating Medium," Computer Graphics Vol. 21 No. 4 (July 1987) pp. 293 - 302.

[Rens] W. A. Rense, "Polarization Studies of Light Diffusely Reflected from Ground and Etched Glass Surfaces," Journal of the Optical Society of America Vol 40 (1950) pp. 55 - 59.

[Sieg]Robert Siegel and John Howell, "Thermal Radiation Heat Transfer, third edition," Hemisphere Publishing, Washington, 1992.

[Sill95a] Francois Sillion, "A Unified Hierarchical Algorithm for Global Illumination with Scattering Volumes and Object Clusters," IEEE Transactions on Visualization and Computer Graphics, Vol. 1 No. 3 (1995) pp. 240 - 254.

[Sill95b] Francois Sillion and George Dretakis, "Feature-based Control of Visibility Error: A Multi-resolution Clustering Algorithm for Global Illumination," ACM Computer Graphics Proceedings, Annual Conference Series (1995) pp. 145 - 152.

[Ward] Gregory Ward, "The RADIANCE Lighting and Rendering System," ACM Computer Graphics Proceedings, Annual Conference Series (1994) pp. 459 - 472.

[Will] Lance Williams, "Casting Curved Shadows on Curved Surfaces," Computer Graphics Vol. 12 No. 3 (August 1978, Siggraph '78 Conference Proceedings) pp. 270 - 274.

[Wool] Joseph Woolley, "Reflectance and Transmittance of Light by Leaves," Plant Physiology Vol. 47 (1971) pp. 656 - 662.

Technical Information Department • **Lawrence Livermore National Laboratory**
University of California • Livermore, California 94551