

# Rendering Translucent Materials Using Photon Diffusion

Craig Donner

Henrik Wann Jensen<sup>†</sup>

University of California, San Diego

---

## Abstract

*We present a new algorithm for rendering translucent materials that combines photon tracing with diffusion. This combination makes it possible to efficiently render highly scattering translucent materials while accounting for internal blockers, complex geometry, translucent inter-scattering, and transmission and refraction of light at the boundary causing internal caustics. These effects cannot be accounted for with previous rendering approaches using the dipole or multipole diffusion approximations that only sample the incident illumination at the surface of the material. Instead of sampling lighting at the surface we trace photons into the material and store them volumetrically at their first scattering interaction with the material. We hierarchically integrate the diffusion of light from the photons to compute the radiant emittance at points on the surface of the material. For increased accuracy we use the incidence plane of the photon and the viewpoint on the surface to blend between three analytic diffusion approximations that best describe the geometric configuration between the photon and the shading point. For this purpose we introduce a new quadpole diffusion approximation that models diffusion at right angled edges, and an attenuation kernel to more accurately model multiple scattering near a light source. The photon diffusion approach is as efficient as previous Monte Carlo sampling approaches based on the dipole or multipole diffusion approximations, and our results demonstrate that it is more accurate and capable of capturing several illumination effects previously ignored when simulating the diffusion of light in translucent materials.*

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computing Methodologies]: Computer GraphicsThree-Dimensional Graphics and Realism

---

## 1 Introduction

Simulating light transport in natural materials such as milk, marble, leaves, skin, and wax is a challenging problem in computer graphics. Light penetrates these translucent materials, scatters inside, and may exit at a different location and direction than it entered. This scattering of light causes many complex phenomena, such as volumetric caustics, translucent inter-scattering between surfaces, and volumetric shadows. These effects are important for conveying realism when rendering images of translucent materials.

The most general methods in computer graphics capable of simulating these effects are Monte Carlo methods. Volumetric caustics and shadows have been rendered with Monte Carlo ray tracing [LW96, JLD99, PKK00] and photon map-

ping [JC98, DEJ\*99]. Unfortunately, when rendering highly scattering materials, these methods have high computational cost.

Many recent appearance models for translucent materials are based on the diffusion approximation [Sta95]. These methods are limited to homogeneous materials with high albedo, and cannot directly simulate full volumetric effects. Jensen et al. [JMLH01] use an analytic diffusion approximation to compute light transport in scattering materials, using Monte Carlo methods to sample the direct illumination. Donner and Jensen [DJ05] extend this diffusion model to thin and layered translucent materials, but use the same sampling technique. To in part capture global effects such as indirect illumination, Jensen and Buhler [JB02] sample the full irradiance of surfaces directly. Sampling the surface can be expensive for complex geometry, and may lead to a biased result without sufficient sampling density. Others have used shadow maps [DS03] or importance sampling [MKB\*03] to estimate

---

<sup>†</sup> {cdonner, henrik}@graphics.ucsd.edu

the incident illumination. All of these methods simulate only surface to surface transport, and cannot render effects such as inter-scattering and volumetric caustics and shadows.

Haber et al. [HMBR05] extend the multigrid method to account for internal visibility in translucent material, but cannot render caustics or translucent inter-scattering.

Other “hybrid” methods couple diffusion with a pre-computation step. Chen et al. [CTW\*04] generate shell texture functions using volumetric photon mapping [JC98]. Their method requires discretizing geometry, and a complex acquisition step to determine the shell volume’s structure. Photons are traced into the material from many directions over the sphere, and contribute to both single and multiple scattering. This requires a correction term during rendering to handle non-diffuse incident light. Photons that pass through the shell are not directly used, and inter-scattering is not simulated by this method.

The hybrid method by Li et al. [LPT05] uses path tracing until eye rays reach a central “core” of the material, at which point the diffusion dipole is used. This method, while general, does not account for multiple inter-scattering between translucent surfaces, or caustics, and is computationally expensive for highly scattering materials.


In this paper, we present a method that incorporates global illumination effects such as caustics, inter-scattering, and volumetric shadows, for translucent materials. Our technique uses a photon tracing step to distribute photons in the material, but uses diffusion to account for multiple scattering. We also extend previous work to handle oblique incident illumination and introduce a new technique for approximating the boundary conditions caused by complex geometry. This avoids the potentially complex surface sampling step associated with previous work, and makes no assumptions on the incident lighting distribution (e.g. normal, diffuse). Photon diffusion can simulate light transport in translucent materials not handled by existing techniques based on diffusion, and would be prohibitively expensive to compute using other methods. These effects include:

- refraction at the interface leading to internal caustics,
- anisotropic surface profiles due to oblique illumination,
- transmission of light through thin areas,
- volumetric effects such as shadows and caustics, and
- translucent inter-scattering from multiply-scattered light.

In addition, our source-based approach lends itself to efficient volumetric hierarchical integration by clustering sources together that are near each other volumetrically.

## 2 Diffusion from point sources

Light scattering in translucent materials is described by the bidirectional scattering surface reflectance distribution function (BSSRDF)  $S$  [NRH\*77]. The BSSRDF defines the general transport of light between two points and directions as the



**Figure 1:** (a) An incident beam is usually transformed into an isotropic source one mean free path below the incident position. (b) The actual source distribution is an infinite set of point sources with exponentially decaying intensity.

fraction of the flux,  $\Phi_i(\vec{x}_i, \vec{\omega}_i)$  incident at  $\vec{x}_i$  that contributes to the radiance  $L_o(\vec{x}_o, \vec{\omega}_o)$  exiting at position  $\vec{x}_o$  in direction  $\vec{\omega}_o$

$$dL_o(\vec{x}_o, \vec{\omega}_o) = S(\vec{x}_i, \vec{\omega}_i; \vec{x}_o, \vec{\omega}_o) d\Phi(\vec{x}_i, \vec{\omega}_i). \quad (1)$$

In highly scattering materials the diffusion of light can be approximated by a diffuse BSSRDF  $S_m$  [JMLH01, JB02]

$$S_m(\vec{x}_i, \vec{\omega}_i; \vec{x}_o, \vec{\omega}_o) = F(\vec{\omega}_i) R(|\vec{x}_i - \vec{x}_o|) F(\vec{\omega}_o) \quad (2)$$

where  $R(d)$  is a diffusion kernel, and  $F(\vec{\omega}_i)$  and  $F(\vec{\omega}_o)$  are the Fresnel terms for a smooth surface. Jensen et al. [JMLH01] use a dipole diffusion approximation for  $R_d$  that computes the diffusion of light due to a virtual light source placed one mean free path inside the material. In the following we generalize the placement of this virtual light source and show how this can be used to compute diffusion with higher accuracy.

Consider the case of a beam of light normally incident on an isotropically scattering semi-infinite homogeneous material. The distance light in this beam travels before interacting (absorbing or scattering) with the material is described by the PDF  $\sigma'_t e^{-\sigma'_t x}$ . Thus the average distance light travels is the expected value of the PDF

$$\int_0^\infty x \sigma'_t e^{-\sigma'_t x} dx = \frac{1}{\sigma'_t} \quad (3)$$


or one mean free path. Here  $\sigma'_t = \sigma_a + \sigma'_s$  is the reduced extinction coefficient, and  $\sigma_a$  and  $\sigma'_s$  are the absorption and reduced scattering coefficients. As it propagates through the material, the beam has a reduced intensity  $L_{ri}$  due to absorption and out-scattering of light [Ish78]

$$L_{ri}(z) = L_i e^{-\sigma'_t z} \quad (4)$$

where  $L_i$  is the irradiance at the surface. Since the material scatters isotropically, the beam is equivalent to an infinite set of point sources along its length with decaying intensity. The average intensity scattered out of the beam is

$$\sigma'_s \int_0^\infty L_i e^{-\sigma'_t x} dx = \alpha' L_i. \quad (5)$$

where  $\alpha' = \sigma'_s / \sigma'_t$  is the reduced albedo. Equations 3 and 5



**Figure 2:** Distributing diffusing sources along the incident beam closely approximates the asymmetric exitant radiance. The top row of images was rendered using the standard diffusion dipole (DD), the second row using Monte Carlo path tracing (MC), while the third row uses sources distributed along the incident beam (PD). Below the images are intensity plots of the green channel horizontal center scanline of the above images. The incident angle of the beam increases from  $0^\circ$  to  $80^\circ$  from left to right, and the beam has unit radius.

imply that a normally incident beam of light can be approximated by a source embedded in the material at one mean free path, with an intensity of  $\alpha' L_i$ . Using the diffusion approximation to compute the exitant surface radiance due to the point source leads to the diffusion dipole approximation [JMLH01]

$$R_d(r) = z_r(1 + \sigma_{tr}d_r) \frac{e^{-\sigma_{tr}d_r}}{4\pi d_r^3} + z_v(1 + \sigma_{tr}d_v) \frac{e^{-\sigma_{tr}d_v}}{4\pi d_v^3} \quad (6)$$

where the radiant emittance is found by summing the contribution of a real source at the average distance light travels, and a virtual negative source mirrored around an extrapolated boundary of height  $2AD$ .  $\sigma_{tr} = \sqrt{3\sigma_a\sigma_t}$  is the effective transport coefficient,  $z_r = 1/\sigma_t'$  and  $z_v = z_r + 4AD$  are the  $z$ -coordinates of the real and virtual sources relative to the surface at  $z = 0$ . Also,  $d_r = \sqrt{r^2 + z_r^2}$  and  $d_v = \sqrt{r^2 + z_v^2}$  are the distances to the sources from a given point on the surface of the object.  $D = \frac{1}{3\sigma_t'}$  is the diffusion constant, and  $A = (1 + F_{dr})/(1 - F_{dr})$  represents the change in fluence due to internal reflection at the surface. The diffuse Fresnel reflectance  $F_{dr}$  is approximated by the following polynomial expansions [EHR73]

$$F_{dr} \approx \begin{cases} -\frac{0.4399}{\eta} + \frac{0.7099}{\eta^2} - \frac{0.3319}{\eta^2} + \frac{0.0636}{\eta^3}, & \eta < 1 \\ -\frac{1.4399}{\eta^2} + \frac{0.7099}{\eta} + 0.6681 + 0.0636\eta, & \eta > 1 \end{cases} \quad (7)$$

or, in the case of a rough surface, is replaced by an average BRDF reflectance [DJ05]. Note that in the standard dipole approximation the source power is the average outscattered power  $\alpha' L_i$ .

## 2.1 Oblique illumination


To use the diffusion dipole, incident rays of light are converted to dipole sources [JMLH01]. Regardless of the incident angle of the illumination, the sources are always mirrored about the incident point (Figure 1a). Thus, the dipole embedded at one mean free path always predicts a radially symmetric radiant emittance (Figure 2, top).

Oblique angles of incidence cause a significant shift in the reflectance, and produce noticeable asymmetric surface effects (Figure 2, middle). This is true in many common materials, where the scattering and absorption parameters are spectral.

We approximate these effects from oblique illumination by placing sources along the incident beam, and using the diffusion approximation to compute the radiant emittance from each of these sources. This is effectively an integration of Equation 6 over the propagating beam, taking into account the exponential decay in Equation 4

$$R_{ri}(r) = \int_0^\infty \kappa(d_r) e^{-\sigma_t' x} \left[ z_r(1 + \sigma_{tr}d_r) \frac{e^{-\sigma_{tr}d_r}}{4\pi d_r^3} + z_v(1 + \sigma_{tr}d_v) \frac{e^{-\sigma_{tr}d_v}}{4\pi d_v^3} \right] dx \quad (8)$$

Note that we redefine the depth of the source as  $z_r = x \cos \theta_i$ , where  $x$  is the distance along the refracted beam and  $\theta_i$  is the angle the refracted beam makes with the internal normal of the surface. The other values  $z_v$ ,  $d_r$  and  $d_v$  are all defined in terms of  $z_r$  as before.  $\kappa(x)$  is an attenuation function discussed below. Unfortunately, this integral has no apparent closed form solution [FPW92]. Rather than attempting to



**Figure 3:** Modulating sources that are close to the surface closely approximates the actual multiple scattering radiant emittance.

numerically integrate it, we sample the integral by distributing diffusing sources along the refracted beam using the PDF  $\sigma'_r e^{-\sigma'_r z}$ . We treat each photon as a positive source of a dipole. The total reflectance is the sum of contributions from these real sources and their corresponding virtual sources. We discuss how to distribute sources in translucent materials in Section 4.1.

Because we use diffusion to compute the multiple scattered radiance from the distributed sources, those close to the surface, where single scattering dominates, will over-estimate the exitant power. Rather than clamp the dipole, we modulate it with a smooth kernel that estimates how much of the source's power contributes to non-single-scattered light

$$\kappa(x) = 1 - e^{-\sigma_r x} \quad (9)$$

Note that this is an approximation, we modulate by the approximate probability that light is not scattered before exiting the surface. We have tested this attenuated dipole with a range of input parameters, and found that it is more accurate than the dipole approximation (see Figure 3).

Figure 2 compares the multiple scattering reflectance of a large body of skim milk illuminated from several incident angles by a beam of light with unit radius. The standard diffusion dipole gives an axisymmetric reflectance, as each photon in the beam transforms into a single source below the incidence point. In this material, however, red light is more highly scattering than green and blue light, and emerges closer to the leading edge of the beam, to the right. As the angle of incidence increases, this effect becomes more prominent. Less scattering blue light emerges more at the trailing edge of the beam, to the left. Using a distributed set of sources captures this shift. It also models the shift of the beam to the right with increasing angle.

### 3 Diffusion in arbitrary geometry

The previous section described placing diffusing sources along refracted incident beams and their virtual sources above

the surface. To increase the accuracy of rendering using sources at arbitrary locations, we introduce a new method to approximate light transport in complex geometry.

The distribution of sources along incident beams of light satisfies an approximate boundary condition on the radiance at the surface in the same way as the original dipole. This approximate condition

$$\phi(r)(-2AD) = 0 \quad (10)$$

states that the diffuse flux  $\phi$  is zero at an extrapolated boundary  $2AD$  above the surface.

If the material is a finite slab of thickness  $d$ , light may scatter out through the bottom of the slab, reducing the observed reflectance. This is modeled by a similar boundary condition at the bottom of the slab

$$\phi(r)(d + 2AD) = 0. \quad (11)$$


Satisfying both conditions requires mirroring the dipole of the top face around the extrapolated boundary of the bottom face, and then repeating this mirroring about the top interface, *ad infinitum*, producing the multipole [DJ05]. The  $z$ -coordinates of the multipole sources are

$$\begin{aligned} z_{r,i} &= 2i(d + 4AD) + \ell \\ z_{v,i} &= 2i(d + 4AD) - \ell - 2z_b, \quad i = -n, \dots, n, \end{aligned} \quad (12)$$

where the number of dipoles is  $2n + 1$ , and  $\ell$  is the depth of the original real source. The total reflectance predicted by the multipole is the sum of the contributions of each of the dipole sources

$$R_m(r) = \sum_{i=-n}^n \frac{z_{r,i}(1 + \sigma_{tr}d_{r,i})e^{-\sigma_{tr}d_{r,i}}}{4\pi d_{r,i}^3} - \frac{z_{v,i}(1 + \sigma_{tr}d_{v,i})e^{-\sigma_{tr}d_{v,i}}}{4\pi d_{v,i}^3} \quad (13)$$

where  $d_{r,i} = \sqrt{r^2 + z_{r,i}^2}$  and  $d_{v,i} = \sqrt{r^2 + z_{v,i}^2}$  are the distances to each source from a point on the surface.



**Figure 4:** The four point sources of the quadpole satisfy the boundary conditions in the case of a right angle between adjacent faces.

### 3.1 A quadpole diffusion approximation

Both the dipole and multipole make the assumption that the surface of the material is flat and infinite in extent. Real objects, however, are often approximated by finite polygonal meshes, and may have sharp corners. For example, in a translucent cube, light that enters at the top face may exit a side face, rather than returning to the top surface as predicted by the dipole.

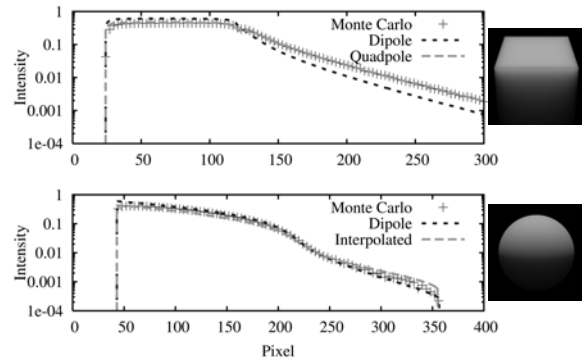
We account for light exiting a face adjoining the lit face at a right angle by imposing an additional boundary condition (see Figure 4). In addition to the extrapolated boundary above the lit face, there is another above the adjoining face. Mirroring the original dipole about this side boundary gives a symmetric set of four sources, a *quadpole*. Note that unlike the multipole, four sources are sufficient to satisfy both boundary conditions, due to symmetry. The reflectance from the quadpole on the side face is

$$R_q(r) = x_r(1 + \sigma_{tr}d_r) \frac{e^{-\sigma_{tr}d_r}}{d_r^3} + x_r(1 + \sigma_{tr}d_v) \frac{e^{-\sigma_{tr}d_v}}{4\pi d_v^3} + x_v(1 + \sigma_{tr}d_{rm}) \frac{e^{-\sigma_{tr}d_r}}{d_{rm}^3} + x_v(1 + \sigma_{tr}d_{vm}) \frac{e^{-\sigma_{tr}d_v}}{4\pi d_{vm}^3} \quad (14)$$

where  $d_{rm}$  and  $d_{vm}$  are the distances from the shading point to the second dipole, and  $x_r$  and  $x_v$  are the distances to the side face (see Figure 4). When calculating the reflectance at the top face,  $x_r$  and  $x_v$  are the distances to the top face, as in the original dipole.

The quadpole sources are the corners of a square centered at the intersection of the extrapolated boundaries (see Figure 4). The plane of the quadpole is perpendicular to the planes of the two faces [Kie05].

For a closed mesh, one could impose boundary conditions on each face and try to satisfy them simultaneously, but this would be costly for non-trivial meshes. Instead, we use only information about the face the source entered the material through, and the viewed face. This allows interpolation between the dipole, quadpole, and multipole to obtain an ap-



**Figure 5:** (top) The quadpole more accurately captures the change in illumination across a right angle than the dipole. (bottom) Interpolating between the dipole, quadpole, and multipole closely approximates the reflectance of smooth surfaces. Both plots are of the blue channel of the center vertical scanline of the image.

proximate reflectance. We discuss this in more detail in Section 4.4.

Figure 5 (top) compares the dipole and quadpole reflectance to a Monte Carlo path tracing simulation on two faces of an optically thick marble box with a corner at about pixel 125. Note the quadpole accurately captures the discontinuity at the corner, and closely approximates the reflectance along the adjoining face. The dipole underestimates the reflectance of the side face of the box, as it does not take the additional boundary condition into account. The bottom of the figure shows the same comparison in the case of a sphere, using an interpolation method described in Section 4.4.

## 4 Implementation


To render with photon diffusion we use a two-pass algorithm. In the first pass, photons are traced through the scene, and stored at their single scattering interactions with translucent materials. Photons may continue to scatter, but are only stored again if they exit and re-enter the material, e.g., in the case of reflection, transmission, or complex geometry (see Figure 6). We then build a hierarchical acceleration structure from the photons for fast evaluation.

In the second pass, during rendering, we traverse the acceleration structure to volumetrically integrate the contributions of the photons using the the dipole, quadpole, and multipole.

### 4.1 Source distribution

Photons are traced from light sources towards translucent objects using standard Monte Carlo methods (e.g., [Jen01]). When a photon hits a scattering material and refracts in, it propagates until it

- (a) hits another interface, in which case the photon is re-



**Figure 6:** (Left) The diffusion dipole is evaluated by sampling the surface of an object to determine the incident flux, and generating sources one mean free path below the incident location. (Right) Tracing photons gives a better distribution of sources, and captures important global illumination effects, such as (a) actual source depth and shifts in illumination due to refraction, (b) global illumination from reflections off other surfaces, (c) volumetric caustics, (d) diffuse caustics, (e) translucent inter-scattering, and (f) indirect illumination on diffuse surfaces from translucent effects.

fracted out of the material and continues on (internal reflection is handled by the diffusion models),

- (b) is absorbed probabilistically or hits some absorbing blocker, in which case the photon is terminated,
- (c) or scatters, in which case we store the photon's position, power, and its incidence plane. The photon then continues to scatter until (a) or (b) occur, but is only stored again if it exits and re-enters the material.

Photons that continue to scatter throughout the material are not stored, as they are in photon mapping, unless they exit and re-enter the material. Multiple scattering is handled by treating the photons as diffusion sources.

Tracing photons in translucent materials captures many important illumination effects, such as refraction at the interface (Figure 6a). If a photon intersects the back side of the geometry, it refracts out, creating caustics on diffuse surfaces (Figure 6d), or self-lighting through an optically thin areas (Figure 6c). Since photons may interact with other objects before the translucent material, there may also be lighting from indirect global illumination (Figure 6b). Photons that scatter out of the material may produce sources in other areas through translucent inter-scattering (Figure 6e), or contribute to indirect illumination on other surfaces (Figure 6f).

The total exitant radiance at a point on the surface is found by summing the contributions of all stored sources. Note that since we are distributing power from the lights into the material, there is no area associated with the photons. This photon tracing step also avoids the potentially expensive task of explicitly sampling the geometry [JB02]. Photons that exit the material and contribute to caustics, translucent inter-scattering, or indirect illumination discretely represent exitant

global illumination due to the multiple scattering of light, just as in photon mapping, thus energy is conserved.

## 4.2 Hierarchical approximation

Using photon tracing to seed diffusion can produce millions of sources. To efficiently calculate their contribution to a shading point, we hierarchically cluster sources together. As Jensen and Buhler [JB02] note, diffusion's exponential fall-off with distance facilitates this approximation. Our hierarchical method is similar to theirs, but we group sources together volumetrically and aggregate different information about them.

Specifically, we construct an octree, with a maximum of eight sources per leaf. For each source, we store its position, power, and the plane of the face at which it entered the material. Once the photon tracing stage is complete, we aggregate the photons in each voxel to create a representative source. Each node computes and stores the total power, an average position weighted by power, and an average plane. We average planes by taking the average normal of the sources, and an average depth below the plane. To ensure the stability of the diffusion calculations we check that the depth is positive, otherwise we move the plane along its normal based on the average depth.

## 4.3 Single scattering

Most approaches for rendering highly scattering materials use an analytical approximation for single scattering [JMLH01]. To capture global illumination effects such as volumetric caustics, we use the single scattered photons stored in the photon tracing step. We march rays through the volume and evaluate volume radiance estimates from a photon map [JC98]. This gives accurate refracted beams, and volumetric caustics.


## 4.4 Multiple scattering

The multiple scattering contribution uses the octree built after the photon tracing stage. We traverse from the root of the octree until we reach a voxel with solid angle (from the perspective of the shading point) greater than a user-set threshold, or until a leaf is reached. The representative source (or sources of a leaf) in the node are then used to calculate the radiant emittance as described in Section 2. We construct a dipole, quadpole, and multipole depending on the angle  $\gamma$  between the shading point normal  $\vec{n}_i$  and the source normal  $\vec{n}_o$

$$R(r) = \frac{2}{\pi} \begin{cases} (\frac{\pi}{2} - \gamma) R_d(r) + \gamma R_q(r), & 0 \leq \gamma \leq \frac{\pi}{2} \\ (\pi - \gamma) R_q(r) + (\gamma - \frac{\pi}{2}) R_m(r), & \frac{\pi}{2} \leq \gamma \leq \pi \end{cases} \quad (15)$$

where  $R_d(r)$  is the dipole,  $R_q(r)$  is the quadpole, and  $R_m(r)$  is the multipole.

In the case of arbitrary geometry, we compute the dipole as in [JMLH01]. We estimate the thickness of the material for the multipole as the sum of the depth of the source relative



**Figure 7:** Photon diffusion simulates indirect lighting in the complex geometry of an ear. Photons that scatter many times distribute multiple sources deep into the ear. This causes in particular the brightening of the frontal ear lobe, and the lighting in the inner ear. These effects are not captured by the dipole or multipole. Also, note that photon diffusion retains more high frequency geometry, as it accounts for the high angle of incident illumination.

to its incident plane, and relative to the view plane. Note we need only compute the multipole reflectance using the observed source depth, and do not need separate reflectance and transmittance multipole computations. To use the quadrupole, we assume the geometry between the source's incident location and the shading point is two planes at a right angle.

#### 4.5 Volumetric shadowing

We use the diffusion sources to simulate first order volumetric shadows. During the photon tracing step, photons may interact with any blocking geometry inside the translucent volume. If the material is absorbing, the photon is terminated. During rendering, we trace a shadow ray from the shade point to the volumetric source. If the ray intersects a blocking material, no contribution from the source is added. Although this method does not account for reflecting blockers, or light paths around blockers, we have found it gives a good first order shadowing approximation.


## 5 Results

We have implemented photon diffusion in a Monte Carlo ray tracer that supports photon tracing as described in [Jen01].

All images in this paper were rendered on an Intel Core 2 Duo 2.4GHz. All images took about 15 minutes to render, there was little difference in render time when using Monte Carlo surface sampling [JMLH01] or photon diffusion. Although there is no sampling noise with photon diffusion, a moderate number of sources is required for accurate sampling of the lighting.

It is unfortunately extremely difficult to render reference images using path tracing methods that capture the effects shown in Figures 7 and 8. The probability of traced rays scattering and hitting a light source is astronomically low. Photon mapping quickly exhausts memory for highly scattering materials. Thus, we directly compare photon diffusion to the diffusion dipole.

Figure 7 shows an example of a backlit ear rendered using photon diffusion and the diffusion dipole. Roughly 3 million photons were stored in the octree, and took about 5 minutes to trace to fully capture translucent inter-scattering. The brightening of the thin parts of the ear in the right image is because the dipole underestimates transmittance [DJ05], whereas with photon diffusion the multipole calculation correctly captures the high translucency. The inner ear, frontal ear lobe, and the side of the head are lit by sources deposited from multiply



**Figure 8:** Our method simulates global illumination effects in translucent materials. The left image was rendered using the diffusion dipole, while the middle and right images using photon diffusion. Note the volumetric caustic at the surface of the milk, and the volumetric shadow cast by the rim of the glass. In the far right image, the brightening at the top of the milk is from light penetrating into the volume, and from indirect illumination from the floor. Light refracting through the bottom of the glass under the milk forms the caustic. Some of this light reflects off the bottom of the glass, causing a subtle glow.

scattered photons. This translucent inter-scattering cannot be captured by previous diffusion techniques that sample only the surface. In addition, photon diffusion shows more high frequency surface geometry, such as the pores of the skin, as it accurately handles the high angles of illumination. The dipole, however, blurs the illumination and smooths the skin's appearance.

In Figure 8 three images of a glass of skim milk are shown. The left image was rendered using the diffusion dipole, while the middle and right images with photon diffusion. Roughly 6 million photons were stored in the middle and right images, taking about one minute to trace, with 3000 photons in the single scattering radiance estimate. Note that photons bouncing off the back side of the glass form a cardioid caustic in the milk. This caustic is *volumetric*, the photons have penetrated into the milk volume. In the right image, the back of the milk is lit by light penetrating through the milk and light bouncing off the floor plane, while the bottom glows from light bouncing off the bottom surface of the glass, and the floor plane. Also, note that the shadow from the rim of the glass is volumetric. The caustic on the ground plane is formed from light refracting through the glass below the milk. Simulating these effects with Monte Carlo path tracing or photon mapping would be prohibitively expensive.

Figure 9 shows an example of volumetric shadows cast from the bones in a backlit hand. The left image was rendered without bones in the hand, while the right image includes the bones. Roughly 0.4 million photons were stored, and took about 15 seconds to trace both with and without the bones geometry. Photons that intersected the bones were absorbed, giving a first order approximation of shadowing as described in Section 4.5. The shadows clearly show the bones in the fingers and in the hand.

## 6 Discussion and Conclusions

We have presented a method for rendering global illumination effects in translucent materials using photon diffusion. We trace photons into the material and store them at their first scattering interactions. Photons continue to scatter in the medium and may contribute to translucent inter-scattering between surfaces. This allows penetration of light into deeply shadowed areas, and simulates volumetric shadows and caustics in translucent materials. For efficiency we cluster photons together and integrate their contribution hierarchically. To improve the accuracy of diffusion calculations, we have introduced a quadpole diffusion approximation, and combined it with the dipole and multiple to simulate light transport in complex geometry. In the future we would like to improve the speed of rendering with photon diffusion, and extend it to simulate light transport in a wider range of materials.


## 7 Acknowledgments

This research was supported by CalIT<sup>2</sup>, the UCSD FWGrid Project (NSF EIA-0303622) and the National Science Foundation (NSF 0305399). The ear model was provided by XYZRGB, and the hand model is courtesy MPII through the AIM@SHAPE Model Repository. Thank you to Toshiya Hachisuka for the bone geometry.

## References

- [CTW\*04] CHEN Y., TONG X., WANG J., LIN S., GUO B., SHUM H.-Y.: Shell texture functions. *ACM Trans. Graphic.* 23 (2004), 343–353.
- [DEJ\*99] DORSEY J., EDELMAN A., JENSEN H. W., LEGAKIS J., PEDERSEN H. K.: Modeling and rendering





**Figure 9:** Standard sampling methods for diffusion do not account for the internal transport of light within the medium, only the transport between surface points. With photon diffusion, we obtain a first order approximation of visibility within the volume. This image of a backlit hand with and without bones show the significance of including internal geometry that block light. In the right image it is easy to see the bones in the hand due to the volumetric shadows.

- of weathered stone. In *Proceedings of ACM SIGGRAPH 1999* (1999), pp. 225–234.
- [DJ05] DONNER C., JENSEN H. W.: Light diffusion in multi-layered translucent materials. *ACM Trans. Graphic.* 24, 3 (2005), 1032–1039.
- [DS03] DACHSBACHER C., STAMMINGER M.: Translucent shadow maps. In *Rendering Techniques* (2003), pp. 197–201.
- [EHR73] EGAN W. G., HILGEMAN T. W., REICHMAN J.: Determination of absorption and scattering coefficients for nonhomogeneous media. 2: Experiment. *Appl. Opt.* 12 (1973), 1816–1823.
- [FPW92] FARRELL T. J., PATTERSON M. S., WILSON B.: A diffusion theory model of spatially resolved, steady-state diffuse reflections for the noninvasive determination of tissue optical properties *in vivo*. *Med. Phys.* 19, 4 (1992), 879–888.
- [HMBR05] HABER T., MERTENS T., BEKAERT P., REETH F. V.: A computational approach to simulate subsurface light diffusion in arbitrarily shaped objects. In *Proceedings of the 2005 conference on Graphics interface* (2005), pp. 79–86.
- [Ish78] ISHIMARU A.: *Wave Propagation and Scattering in Random Media*. Oxford University Press, 1978.
- [JB02] JENSEN H. W., BUHLER J.: A rapid hierarchical rendering technique for translucent materials. *ACM Trans. Graphic.* 21 (2002), 576–581.
- [JC98] JENSEN H. W., CHRISTENSEN P. H.: Efficient simulation of light transport in scenes with participating media using photon maps. In *Proceedings of ACM SIGGRAPH 1998* (1998), pp. 311–320.
- [Jen01] JENSEN H. W.: *Realistic Image Synthesis Using Photon Mapping*. AK Peters, 2001.
- [JLD99] JENSEN H. W., LEGAKIS J., DORSEY J.: Rendering of wet materials. In *Rendering Techniques* (1999), pp. 273–282.
- [JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. In *Proceedings of ACM SIGGRAPH 2001* (2001), pp. 511–518.
- [Kie05] KIENLE A.: Light diffusion through a turbid par- allelepiped. *J. Opt. Soc. Am.* 22, 9 (2005), 1883–1888.
- [LPT05] LI H., PELLACINI F., TORRANCE K.: A hybrid monte carlo method for accurate and efficient subsurface scattering. In *Rendering Techniques* (2005), pp. 283–290.
- [LW96] LAFORTUNE E. P., WILLEMS Y. D.: Rendering participating media with bidirectional path tracing. In *Rendering Techniques* (1996), pp. 91–100.
- [MKB\*03] MERTENS T., KAUTZ J., BEKAERT P., REETH F. V., SEIDEL H.-P.: Efficient rendering of local subsurface scattering. In *Proceedings of the 11th Pacific Conference on Computer Graphics and Applications* (2003), pp. 51–58.
- [NRH\*77] NICODEMUS F. E., RICHMOND J. C., HSIA J. J., GINSBERG I. W., LIMPERIS T.: *Geometrical Considerations and Nomenclature for Reflectance*. National Bureau of Standards, 1977.
- [PKK00] PAULY M., KOLLIG T., KELLER A.: Metropolis light transport for participating media. In *Rendering Techniques* (2000), pp. 11–22.
- [Sta95] STAM J.: Multiple scattering as a diffusion process. In *Rendering Techniques* (1995), pp. 41–50.