

Efficient Multiple Scattering in Hair Using Spherical Harmonics

Jonathan T. Moon

Bruce Walter

Steve Marschner

Cornell University

Abstract

Previous research has shown that a global multiple scattering simulation is needed to achieve physically realistic renderings of hair, particularly light-colored hair with low absorption. However, previous methods have either sacrificed accuracy or have been too computationally expensive for practical use. In this paper we describe a physically based, volumetric rendering method that computes multiple scattering solutions, including directional effects, much faster than previous accurate methods. Our two-pass method first traces light paths through a volumetric representation of the hair, contributing power to a 3D grid of spherical harmonic coefficients that store the directional distribution of scattered radiance everywhere in the hair volume. Then, in a ray tracing pass, multiple scattering is computed by integrating the stored radiance against the scattering functions of visible fibers using an efficient matrix multiplication. Single scattering is computed using conventional direct illumination methods. In our comparisons the new method produces quality similar to that of the best previous methods, but computes multiple scattering more than 10 times faster.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism

Keywords: Hair, multiple scattering, spherical harmonics

1 Introduction

Physically based simulation of light scattering in hair is needed to reproduce the appearance of real hair accurately. Previous work has shown that hair fibers transmit a large fraction of the light that falls on them [Marschner et al. 2003] and that, because of this, multiple scattering is important to the appearance of hair—in fact, multiply scattered light predominates in blond and other light colored hair [Moon and Marschner 2006]. An accurate hair renderer must therefore allow for interreflections among the fibers, accounting for all paths by which light can be scattered from fiber to fiber before reaching the eye. However, simulating global illumination in a model consisting of tens of thousands of densely packed fibers is a daunting problem. Previous hair rendering systems have sidestepped multiple scattering by introducing an unrealistically large diffuse component in the fibers’ scattering function, have used coarse approximations to multiple scattering that discard directional variation that is relevant to appearance, or have involved expensive computations that take hours of CPU time per frame.

Photon mapping methods have proven successful for hair rendering [Moon and Marschner 2006; Zinke 2008], but they are slow and require a large amount of storage for the photons. This paper proposes a new method for computing multiple scattering in hair. Like photon mapping, it uses a light tracing and a ray tracing pass, but it stores the position- and direction-dependent scattered radiance distribution in a 3D grid of coefficients for spherical harmonic basis functions. Compared to a photon map, this representation for radiance is more compact and better organized in memory, and it allows fast integration by sparse matrix multiplication during rendering. At the same time, we avoid the high cost of tracing paths from fiber to fiber by replacing the hair geometry with a voxel grid that stores density and orientation statistics, through which paths are traced by volumetric methods.

The end result of adopting smooth representations for the essentially smooth phenomena of volume scattering in hair is that our new method outperforms an implementation of photon mapping based on the same ray tracing infrastructure [Moon and Marschner 2006] by a factor of 20 while achieving results of equivalent quality.

2 Background

2.1 Prior work

This work is focused on the accurate calculation of multiple scattering in volumes of light colored hair fibers. Existing techniques handle multiple scattering in a number of different ways. The most common is to approximate its contribution as a constant diffuse term in single scattering calculations [Kajiya and Kay 1989; Marschner et al. 2003; Zinke and Weber 2007], ignoring directional effects altogether. All widely used hair rendering methods use this approach; for a survey see Ward et al. [2007]. The most accurate alternative would be to perform brute-force Monte Carlo path tracing [Kajiya 1986], but the rate of convergence makes it prohibitively slow. A more practical class of methods are generalized photon mapping techniques [Moon and Marschner 2006; Zinke 2008] which can smoothly approximate multiple scattering in volumes by estimating densities of traced particles. However, those approaches are memory intensive and still require several hours of computation to generate high quality still images. Our new method builds upon their ideas but avoids both of these limitations.

One of the core concepts in our work is to consider assemblies of hair as continuous volumetric media that vary both spatially and directionally. Multiple scattering within such media has been studied extensively in computer graphics [Cerezo et al. 2005], including ray tracing approaches [Kajiya and von Herzen 1984], two pass methods such as photon mapping [Jensen and Christensen 1998], and diffusion-based approximations [Stam 1995]. By posing the calculation of multiple scattering in hair as a continuous volumetric problem we can draw insights from these existing ideas, although the oriented nature of hair volumes prevents these methods from being applied directly. Other research has focused on the relationship between volumes of scattering geometry and continuous media [Moon et al. 2007] but assumed randomly oriented geometry.

Many methods for rendering multiple scattering in clouds use a

voxelized representation of a volume, and we take the same approach here. In those methods, a regular grid stores volume parameters and is also directly used to determine the flow of light through a volume, as in finite element methods [Bhate and Tokuta 1992; Rushmeier 1988] and energy shooting approaches [Max 1994; Nishita et al. 1996]. Our grid also stores hair volume parameters, but unlike these previous methods we trace light paths stochastically and simply store the resulting distribution in the grid.

We use spherical harmonic functions to represent the directional radiance distributions at positions within a volume of hair. Spherical harmonics are a popular method for representing spherical functions in applications from global illumination [Westin et al. 1992; Sillion et al. 1991] to efficient real-time rendering of soft lighting effects [Ramamoorthi and Hanrahan 2004; Sloan et al. 2002]. Our use of spherical harmonics on a grid is related to the idea of irradiance volumes [Greger et al. 1998], though we store radiance rather than irradiance and our application is different.

2.2 Spherical harmonics

The spherical harmonics are a basis for functions on the sphere that is used widely in graphics. More information on spherical harmonics, their definition, and their efficient evaluation can be found in the references [Green 2003; Ramamoorthi and Hanrahan 2004].

We use the real form of the spherical harmonics, which are a set of real-valued, orthonormal functions of direction. The spherical harmonics $Y_l^m(\omega)$ are indexed by two integers l and m , called the *degree* and *order* respectively, with $l \geq 0$ and $|m| \leq l$. In this paper, for notational simplicity we index the spherical harmonics by a single index $k > 0$, so that

$$Y_1 = Y_0^0; Y_2 = Y_1^{-1}; Y_3 = Y_1^0; Y_4 = Y_1^1; Y_5 = Y_2^{-2}; \dots$$

There are $(D+1)^2$ spherical harmonics of degree at most D , so a function f represented by spherical harmonics up to degree D is

$$f(\omega) = \sum_{k=1}^{(D+1)^2} c_k Y_k(\omega)$$

where the vector \mathbf{c} is called the spherical harmonic coefficients of the function f . We say that $f \in \mathcal{Y}_D = \text{span}\{Y_1, \dots, Y_{(D+1)^2}\}$.

The properties of the Y_k s that are relevant in this paper are:

- The spherical harmonics are orthonormal: $\langle Y_k, Y_k \rangle = 1$, and $\langle Y_k, Y_{k'} \rangle = 0$ for $k \neq k'$.¹
- The spherical harmonics up to degree D are closed under rotations. Starting with a function $f \in \mathcal{Y}_D$, a rotation of that function $f^R(\omega) = f(R\omega)$ is still in \mathcal{Y}_D . If the coefficients of f are \mathbf{c} , then the coefficients of f^R are $\mathbf{R}\mathbf{c}$ for a square matrix \mathbf{R} that depends only on R . A simple and efficient algorithm for multiplying by \mathbf{R} is given by Kautz et al. [2002].

3 Method

The input to our renderer is a list of 3D curves describing the geometry of the hair fibers, which are treated as elliptical tubes for ray intersection. Their scattering properties are defined by a scattering function $f_s(\mathbf{x}, \omega_o, \omega_i)$ that describes how much light, in total, is scattered from ω_i to ω_o by the fiber around \mathbf{x} , without regard for its distribution across the fiber (see Marschner et al. [2003] for a complete definition).

¹ $\langle f, g \rangle$ is the inner product with respect to solid angle over the sphere: $\int_{S^2} f(\omega)g(\omega)d\omega$

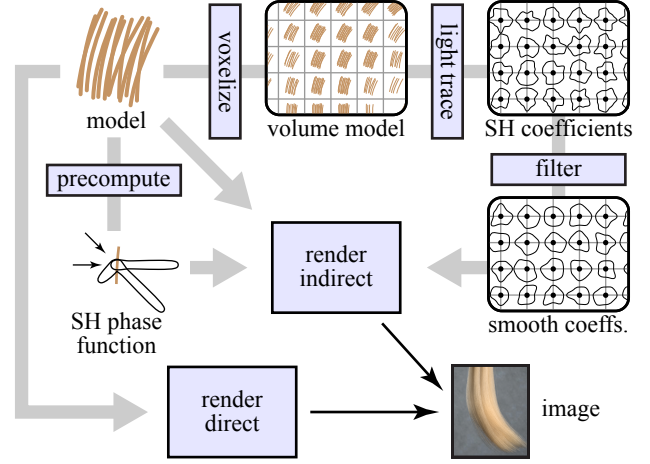


Figure 1: Block diagram of the phases of our rendering method, progressing from the model to the volume of spherical harmonic coefficients used by the rendering phase.

Our method proceeds as illustrated in Figure 1. Before rendering, the hair geometry is converted into a volume, producing a regular grid of density and orientation statistics throughout the hair. The scattering function of the fibers is also processed to convert it into a matrix suitable for transforming incident to scattered light in the spherical harmonic basis.

The rendering process centers around the scattered radiance distribution $L_s(\mathbf{x}, \omega)$, which is the average radiance found near the point \mathbf{x} and the direction ω , excluding light coming directly from sources. In the first pass, L_s is estimated everywhere in the hair by using the stored volume statistics to trace random paths from the light sources, updating spherical harmonic coefficients describing scattered radiance in each cell of the grid. The resulting coefficients are filtered to suppress ringing and noise. In the second pass, pixels are rendered by tracing rays from the eye, and direct illumination is calculated by sampling light sources by standard methods. Indirect illumination is computed using the stored L_s coefficients, which must be rotated into the fiber's local frame, and the precomputed scattering matrix.

We chose to use spherical harmonics for two reasons: because rotations are convenient and efficient, and because their axial structure makes them well-suited to representing fiber scattering functions, which naturally separate into longitudinal and azimuthal components [Marschner et al. 2003]. These special properties of the SH basis are only relevant in the rendering phase; the light path tracing phase would be the same for any orthogonal directional basis.

The spherical harmonics have some disadvantages in our application, including their global support, which means all the coefficients must be computed to evaluate a function for a single direction, and their propensity for ringing when approximating functions with sharp features. The measures we take to combat noise and ringing are discussed in Section 3.4.

In the following sections the different parts of the algorithm are described in more detail, beginning with the final rendering pass to establish notation.

3.1 Rendering

To render the output image, a standard ray tracing procedure is followed starting from the eye. When a ray hits a visible fiber at \mathbf{x}




Figure 2: During the rendering phase, the grid of coefficients is queried to find the scattered radiance incident on the visible fiber. The coefficients are interpolated from the grid, rotated into the fiber's coordinate frame, and multiplied by the precomputed scattering-function matrix to find the distribution of radiance scattered from the fiber, which is then evaluated in the viewing direction.

with fiber tangent u , we must evaluate the scattering integral:

$$L_o(\mathbf{x}, \omega_o) = \int_{S^2} f_s(\mathbf{x}, \omega_o, \omega_i) L_i(\mathbf{x}, \omega_i) \sin(\omega_i, u) d\omega_i \quad (1)$$

The scattered radiance L_o in the direction ω_o toward the eye is the integral over the sphere S^2 of the scattering function f_s against the incident radiance distribution L_i with respect to projected solid angle. L_i can be expressed as a sum of direct radiance L_d and scattered (indirect) radiance L_s . Then $L_o = L_{o,d} + L_{o,s}$, where

$$L_{o,d}(\mathbf{x}, \omega_o) = \int_{S^2} f_s(\mathbf{x}, \omega_o, \omega_i) L_d(\mathbf{x}, \omega_i) \sin(\omega_i, u) d\omega_i \quad (2)$$

$$L_{o,s}(\mathbf{x}, \omega_o) = \int_{S^2} f_s(\mathbf{x}, \omega_o, \omega_i) L_s(\mathbf{x}, \omega_i) \sin(\omega_i, u) d\omega_i \quad (3)$$

The direct illumination component $L_{o,d}$ is computed by sampling luminaires using standard ray tracing techniques. The multiple scattering (indirect) component $L_{o,s}$ is evaluated based on the stored representation of $L_s(\mathbf{x}, \omega_i)$ as a set of coefficients for spherical harmonic functions:

$$L_s(\mathbf{x}, \omega_i) \approx \sum_{k=1}^{(D+1)^2} c_k(\mathbf{x}) Y_k(\omega_i) \quad (4)$$

The coefficients $c_k(\mathbf{x})$ (collectively the vector \mathbf{c}) are found by tri-linearly interpolating coefficients (for spherical harmonics in the world frame) from the grid, then using the method of Kautz et al. [2002] to rotate \mathbf{c} to $\mathbf{c}' = \mathbf{R}(\mathbf{x}) \mathbf{c}$ so that they are defined relative to the local frame of the fiber at \mathbf{x} .

Once these coefficients are known, the integral in (3) can be computed. Because this integral acts as a linear operator mapping L_s to $L_{o,s}$, and this operator is fixed when directions are measured in the fiber's local frame, it can be expressed as a matrix transformation on the coefficients, so that the coefficients of the exiting radiance are $\mathbf{d}' = \mathbf{F} \mathbf{c}'$. The precomputation of the scattering matrix \mathbf{F} , which depends only on the properties of the fibers, is explained in Section 3.5. Figure 2 illustrates this computation.

If we have opted not to precompute \mathbf{F} , the exitant scattered radiance from a fiber in direction ω_o can be estimated by generating 20 to 50 directional “stabs” by sampling $f_s(\mathbf{x}, \omega_o, \omega_i) \sin(\omega_i, u)$. The filtered grid values can then be used to evaluate the incident radiance for each stab direction which, when weighted by their respective attenuation from scattering, average together to the exitant radiance. In our implementation, the matrix multiplication approach runs around 3 times faster than doing 20 stabs.




Figure 3: The attenuation coefficient for a group of hair fibers varies with respect to direction. Attenuation is minimized in the average fiber direction $\bar{\omega}$, seen here when $\theta = 0$ and π , and is maximized in directions ω perpendicular to $\bar{\omega}$, when $\theta = \pi/2$. As the directional deviation ν from $\bar{\omega}$ for a group of fibers increases, the attenuation coefficient becomes less directional.

These shading operations are still too expensive to do per eye ray, so we used a fiber radiance cache, as described by Moon and Marschner [2006]. Multiple scattering computations are cached along hair fibers after they are calculated, and subsequent computations are interpolated from nearby cached values when possible.

3.2 Voxelizing the hair

The complexity of hair models makes it expensive to carry out the first pass by light tracing the geometry. Because the details of the hair arrangement are not important for multiple scattering, a different path-tracing procedure can be substituted as long as it generates the same statistical distribution of paths. We adopt a volumetric approach that determines the attenuation and average phase function for regions within the volume, both of which vary with direction and depend on the underlying arrangement of hair fibers.

We use the voxel grid, which is already needed to store coefficients of radiance, to also store local measures of hair density and the distribution of fiber directions. We assume that within each region, hair fibers are somewhat aligned and clustered around a single average direction. We then consider three parameters: the average fiber direction $\bar{\omega}$, the standard deviation of fiber directions ν , and the perpendicular attenuation coefficient σ_t^\perp . This third quantity is the attenuation $\sigma_t(\omega)$ that would result if all fibers in a voxel were parallel and well separated and ω was perpendicular to $\bar{\omega}$.

In the parallel fiber case, the attenuation function $\sigma_t(\omega)$ is simply

$$\sigma_t(\omega) = \sigma_t^\perp \sin \theta \quad (5)$$

where θ is the angle between ω and $\bar{\omega}$. For non-parallel but still well separated fibers, the attenuation is (5) averaged over the distribution of fiber directions. This average can be precomputed for a range of values of θ and ν and interpolated; the precomputation takes a few seconds and is stored in a 2D table. Figure 3 shows plots of $\sigma_t(\theta)/\sigma_t^\perp$ for several values of ν . This estimate is approximate when fibers are not well separated, but we have found it to be sufficiently accurate in practical cases.

Similarly, the phase function of a voxel is simply f_s averaged over the distribution of fibers in that voxel. We can trace light paths without calculating it explicitly, as discussed in Section 3.3.

To estimate $\bar{\omega}$, ν , and σ_t^\perp on our voxel grid, we use a process we call *voxelization*. As illustrated in Figure 4, we rasterize all the hair segments into the voxel grid, for each segment visiting the voxels




Figure 4: A hair fiber is voxelized. For each sample point near the fiber, contributions are made to the density and the fiber direction distribution at that point, weighted based on the minimum distance to the fiber. The result at each sample point is the volume density of fibers and the mean and variance of the fiber direction distribution.

within a distance d from the segment. (Voxels that are closer to one of the adjacent segments are skipped to minimize double counting of hairs.) After this process, each voxel has a list of the nearby segment directions, with magnitudes scaled by a weight that drops linearly to zero at the distance d . The normalized sum of these directions is taken to be $\bar{\omega}$, ν is the standard deviation of the dot products of the (normalized) segment directions with $\bar{\omega}$, and

$$\sigma_t^\perp = \frac{2rN}{\pi d^2} \quad (6)$$

where $2r$ is the fiber diameter, N is the number of nearby fibers for this voxel, and d is the voxelization search distance.

Voxels that receive no contribution are marked as empty. Cells that are far enough from the nearest non-empty cell that they cannot participate in any lookup (even indirectly via the filtering pass described in Section 3.4) are marked inactive, and no coefficients will be allocated or computed for them.

3.3 Light tracing

The purpose of the light tracing phase is to estimate the scattered radiance in all directions everywhere in the hair volume, in terms of a 3D grid with a vector of spherical harmonic coefficients stored at each grid point. In this context the radiance is measured as a density of power-weighted path length per unit volume per unit solid angle:

$$\frac{(\text{power})(\text{path length})}{(\text{volume})(\text{solid angle})} = \frac{W \cdot m}{m^3 \cdot \text{sr}}.$$

We handle the density over volume by dividing space into disjoint rectilinear voxels, one per grid point, and computing the density for each point using the intersections of the paths with that cell. The density over directions is easy to estimate because of the orthogonality of our basis: the coefficient of Y_k at the j th grid point is

$$c_{jk} = \sum_p \frac{\ell_{pj} Y_k(\omega_p) \Delta \Phi_p}{\Delta V} \quad (7)$$

where the sum is over the intersections of all path segments with the cell j , which has volume ΔV . The contribution is proportional to the length ℓ_{pj} of the intersection between segment p and cell j , and to the power $\Delta \Phi_p$ associated with the light particle. This process is illustrated in Figure 5.

Light tracing proceeds by tracing paths from light sources in a distribution proportional to emitted power, in the same way as for photon mapping, keeping track of the power associated with each path.




Figure 5: A light path being traced through the voxel grid. For each cell a path segment intersects, a contribution to the directional radiance coefficients at that cell's sample point is made, proportional to the path length within the cell. The result at each sample point is a set of spherical harmonic coefficients representing the directional distribution of scattered radiance averaged over the cell.

After the first geometric scattering event, we walk through the voxel grid along each path segment, making a contribution according to (7) for each intersected grid cell. Scattering occurs according to $\sigma_t(\omega)$ for each voxel, which is derived from that voxel's stored parameters as discussed in Section 3.2. Upon scattering, we choose a fiber direction u from the distribution described by $\bar{\omega}$ and ν and then sample scattered direction ω_i from $f_s(\mathbf{x}, \omega_o, \omega_i) \sin(\omega_i, u)$.

As in earlier work [Moon and Marschner 2006], we use a ray tracing procedure to choose ω_i according to $f_s(\mathbf{x}, \omega_o, \omega_i) \sin(\omega_i, u)$. Since the model of Marschner et al. [2003] is difficult to sample, we instead compute the refracted direction for a random entry point on a rough elliptical fiber.² This produces an accurate phase function but is too expensive to evaluate relative to the cost of traversing the volume. To eliminate this cost we use a table-driven scattering procedure: instead of running the computation anew for each scattering event, we compute several thousand random directions for each possible inclination, then randomly choose one of these stored directions for each scattering event during the simulation. Incident rays are generated with random azimuthal angle and randomly across the width of a fiber, resulting in a 2D distribution that, though discrete, gives an excellent approximation to the correct behavior once it has been smoothed through multiple scattering and projection onto the spherical harmonics basis.

3.4 Filtering

Because the path density contains sharp peaks in some areas, directly using the accumulated coefficients will lead to ringing in the spherical harmonics. Unfortunately, the Gibbs phenomenon states that increasing the degree of the spherical harmonics will not reduce this ringing appreciably until the peak is fully resolved, which would require far more coefficients than is practical. Instead, we apply a gentle lowpass filter, attenuating the coefficients by a factor that drops smoothly to zero just beyond the maximum degree.

In the spatial domain, there is random noise in the coefficients that depends on the ratio of the number of cells to the number of paths. To reduce this noise to imperceptible levels, we can increase the number of paths, decrease the grid resolution, or use a smoothing filter to average nearby cells. For a given budget of paths, we normally use a dense enough grid that the noise would be evident in

²We simulate roughness by simply randomizing surface normals at fiber interfaces. This introduces a slight non-physical bias, but has little effect on the method's overall accuracy.

images, because it can be removed more controllably by spatial filtering than by decreasing grid resolution. Specifically, we convolve the grid with a box filter that has spherical support between 1.5 and 3 grid cells in radius. Less aggressive filtering is required for still frames than for animations, because the spatial noise is at fairly low frequency and is hard to notice until temporal incoherence turns it into perceptible flicker.

3.5 Scattering matrix estimation

Before the rendering of a particular model begins, we precompute the representation of the hair’s scattering function in terms of the spherical harmonic basis. As with any linear operator applied to functions in an orthonormal basis, the effect of applying the scattering integral to the radiance expressed in spherical harmonics amounts to a linear transformation of the coefficients. Substituting (4) into (3) and computing the coefficients $d'_j = \langle Y_j, L_{o,s} \rangle$:

$$\begin{aligned} d'_j &= \left\langle Y_j, \int_{S^2} f_s(\omega'_o, \omega'_i) \sum_{k=1}^{(D+1)^2} c'_k Y_k(\omega'_i) \sin(\omega'_i, u') d\omega'_i \right\rangle \\ &= \sum_{k=1}^{(D+1)^2} F_{jk} c'_k \end{aligned} \quad (8)$$

where ω'_i and ω'_o are directions measured relative to the local coordinate frame of the fiber, and u' is the fiber direction in the local frame (the z axis). The matrix entries are then:

$$F_{jk} = \iint f_s(\omega'_o, \omega'_i) Y_j(\omega'_o) Y_k(\omega'_i) \sin(\omega'_i, u') d\omega'_i d\omega'_o \quad (9)$$

The matrix \mathbf{F} depends only on the fibers’ intrinsic scattering properties. In a pre-process, which only needs to be run when the properties of the fibers are changed, we compute these integrals using straightforward Monte Carlo integration, then threshold the matrix to obtain a sparse approximation with about 2% nonzero entries. When the hair parameters need to be changed frequently, we resort to the somewhat slower “stabbing” alternative for the integration in the final pass (see Section 3.1).

In our models we have no fiber-to-fiber variation in properties; all texture arises from illumination effects. However, when such variation is required, multiple precomputed scattering matrices may be interpolated to obtain scattering matrices for a range of parameters.

4 Results

We implemented our method in Java, using a hierarchical grid to accelerate ray-hair intersection. Timings are reported for a single-threaded process running on a 3.0 Ghz Intel Core 2 Duo workstation. Spherical harmonics up to degree 15 (256 coefficients) were used for all the results.

Our method involves four different precomputations: we voxelize the hair geometry (Section 3.2), we tabulate the fibers’ volume attenuation function (Section 3.2), we tabulate the scattering function of an individual fiber (Section 3.3), and we (optionally) precompute a sparse scattering matrix \mathbf{F} (Section 3.5). All of these computations except for the generation of the scattering matrix are performed at run-time, and are thus included in the reported timings for the Tracing phase of our method. Typically, voxelization takes 10 seconds, and tabulating the attenuation function and scattering function each take 3 seconds.

The off-line precomputation of the scattering matrix for a particular hair type (color and eccentricity) requires 500,000 importance-sampled direction pair evaluations to produce a matrix that gives

visually indistinguishable results from those of the stabbing rendering method. This takes 700 seconds on a single processing core, but since this matrix depends only on the fiber properties this work can be amortized over many frames of an animation.

In Figure 6, we show the comparison between our method and the photon mapping method of Moon and Marschner [2006] for computing directionally varying multiple scattering in hair. We rendered the front-lit and back-lit ponytail scenes from the previous paper with both methods. With both models tuned to produce a level of noise acceptable in a still image, the previous method consumed around 1.5 hrs of CPU time and 1.5 GB of RAM, whereas our new method produces multiple scattering results of similar quality in 5 minutes using 270 MB of RAM. Note that direct illumination is calculated separately in both cases, and took an additional 10 minutes to compute.

For three models, *ponytail*, *curls*, and *braids*, we produced animations of the hair revolving in studio-like soft lighting (a key light 50 degrees to the right of the viewer, and rim light 150 degrees to the left for *ponytail* and *braids*) and of a small light source orbiting around them. These videos demonstrate the ability of our method to capture a variety of directional effects across a range of poses and lighting conditions.

Figure 7 shows a single frame from each of these animations. Some points to notice in these images and in the animations include:

- In the ponytail studio scene, note the strong directional effects near the bottom of the hair. When the hair is curving toward the camera, the camera is near the specular cone of the key light and the hair looks light; when it is facing away, although it is still illuminated, the hair appears dark because the light is scattered away from the camera.
- In the braids studio scene, note that our grid successfully resolves the intricate changes in multiple scattering produced by the braids. The multiple scattering component plays an important role in the appearance of the highlights on the braids. Also, note how light bleeds translucently into the shadows cast by the braids, taking on a more saturated color.

Rendering multiple scattering in these animations takes on average 13 minutes per frame on a single processing core, depending on the orientation of the geometry with respect to the light sources and camera. This is longer than a single still image because more paths, and somewhat increased smoothing, are required to completely eliminate flicker. Rendering single scattering takes around 1 hour for studio scenes and 10 minutes for orbit scenes – considerably longer than the full multiple scattering computation. Table 1 lists these timings in further detail, as well as other parameters of our method and of the scenes we rendered.

5 Conclusions

In this paper we have demonstrated that physically based multiple scattering simulation for hair does not have to be nearly as expensive as previously thought. The methods presented in this paper bring multiple scattering down in cost to where it can be used in production to achieve the subtle coloration and radiant “glow” that are characteristic of hair.

One major reason for our method’s speed is that it takes full advantage of the smoothness in the multiple scattering distribution, without giving up the ability to render important effects caused by directionally varying multiple scattering. Another result of this smoothness assumption, however, is that there is a limit to how directional radiance distributions can be and still be accurately represented by our spherical harmonic basis. Similarly, the spatial resolution of



Figure 6: Comparison against the method of Moon and Marschner [2006], using their model and both their (left) and our (right) methods. We obtain results of equivalent visual quality, but more than an order of magnitude faster.




Figure 7: Renderings of three hairstyles using our new method. (a-c) Renderings under studio-like lighting, producing realistic results under relatively soft lighting. (d-f) Renderings under strong back lighting, illustrating our method's ability to handle stronger directional effects. See the accompanying video for animation sequences corresponding to these six images.

Scene	Hairs	Segments	Grid size	Active cells	Paths traced	Smoothing r	Trace time	Render time	Total time
Ponytail	27K	1M	33x109x49	46074	160M	2.2 cells	11 min	2 min	13 min
Braids	48K	5M	74x107x32	160686	100M	2.0 cells	10 min	3 min	13 min
Curls	8K	1.5M	46x133x87	132165	100M	2.0 cells	8 min	4 min	12 min

Table 1: The parameters of our model, details of our scenes, and performance results. Timings refer to the computation of multiple scattering only; single scattering typically took about 1 hour per frame for studio scenes and 10 minutes for orbiting scenes. During animations, timings and memory usage vary based on the geometry’s orientation to the lights and camera. Images were rendered at 600 x 600 resolution, 64 rays per pixel. All tests were performed on a single core of a 3.0 Ghz Intel Core 2 Duo workstation restricted to 1 GB of RAM.

our method is limited by the spacing of our rectilinear grid. This results in an amount of spatial and angular blurring comparable to photon mapping methods, but with considerably less memory and time required.

An important factor limiting the quality of the results is the quality of the models. Our models were created in a state-of-the-art hair modeling system, but we still do not know whether they are arranged at all similarly to the fibers in real hair assemblies. Learning about the statistics of fibers in natural hair and creating tools that can produce physically plausible fiber arrangements are important areas of future work that will help physically based rendering of hair achieve its full potential.

Acknowledgments

This research was supported by funding from Unilever Corporation, NSF CAREER award CCF-0347303, NSF grant CCF-0541105, and an Alfred P. Sloan Research Fellowship. Computing equipment was donated by Intel Corporation.

References

- BHATE, N., AND TOKUTA, A. 1992. Photorealistic volume rendering of media with directional scattering. In *Eurographics Rendering Workshop 1992*, 227–245.
- CEREZO, E., PÉREZ, F., PUEYO, X., SERÓN, F. J., AND SILLION, F. X. 2005. A survey on participating media rendering techniques. *The Visual Computer* 21, 5, 303–328.
- GREEN, R., 2003. Spherical harmonic lighting: The gritty details. Game Developers Conference.
- GREGER, G., SHIRLEY, P. S., HUBBARD, P. M., AND GREENBERG, D. P. 1998. The irradiance volume. *IEEE Computer Graphics & Applications* 18, 2 (Mar./Apr.), 32–43.
- JENSEN, H. W., AND CHRISTENSEN, P. H. 1998. Efficient simulation of light transport in scenes with participating media using photon maps. In *Proceedings of ACM SIGGRAPH 98*, Computer Graphics Proceedings, 311–320.
- KAJIYA, J. T., AND KAY, T. L. 1989. Rendering fur with 3D textures. In *Computer Graphics (Proceedings of ACM SIGGRAPH 89)*, 271–280.
- KAJIYA, J. T., AND VON HERZEN, B. P. 1984. Ray tracing volume densities. In *Computer Graphics (Proceedings of ACM SIGGRAPH 84)*, 165–174.
- KAJIYA, J. T. 1986. The rendering equation. In *Computer Graphics (Proceedings of ACM SIGGRAPH 86)*, 143–150.
- KAUTZ, J., SLOAN, P.-P., AND SNYDER, J. 2002. Fast, arbitrary brdf shading for low-frequency lighting using spherical harmonics. In *Eurographics Rendering Workshop 2002*, 291–296.
- MARSCHNER, S. R., JENSEN, H. W., CAMMARANO, M., WORLEY, S., AND HANRAHAN, P. 2003. Light scattering from human hair fibers. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2003)* 22, 3, 780–791.
- MAX, N. L. 1994. Efficient Light Propagation for Multiple Anisotropic Volume Scattering. In *Eurographics Rendering Workshop 1994*, 87–104.
- MOON, J. T., AND MARSCHNER, S. R. 2006. Simulating multiple scattering in hair using a photon mapping approach. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2006)* 25, 3, 1067–1074.
- MOON, J. T., WALTER, B., AND MARSCHNER, S. R. 2007. Rendering discrete random media using precomputed scattering solutions. In *Eurographics Symposium on Rendering 2007*, 231–242.
- NISHITA, T., DOBASHI, Y., AND NAKAMAE, E. 1996. Display of clouds taking into account multiple anisotropic scattering and sky light. In *Computer Graphics (Proceedings of ACM SIGGRAPH 96)*, vol. 30, 379–386.
- RAMAMOORTHY, R., AND HANRAHAN, P. 2004. A signal-processing framework for reflection. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2004)* 23, 4, 1004–1042.
- RUSHMEIER, H. 1988. *Realistic Image Synthesis for Scenes with Radiatively Participating Media*. PhD thesis, Cornell University.
- SILLION, F. X., ARVO, J. R., WESTIN, S. H., AND GREENBERG, D. P. 1991. A global illumination solution for general reflectance distributions. In *Computer Graphics (Proceedings of ACM SIGGRAPH 91)*, 187–196.
- SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2002)* 21, 3, 527–536.
- STAM, J. 1995. Multiple scattering as a diffusion process. In *Eurographics Rendering Workshop 1995*, 41–50.
- WARD, K., BERTAILS, F., KIM, T.-Y., MARSCHNER, S. R., CANI, M.-P., AND LIN, M. 2007. A survey on hair modeling: Styling, simulation, and rendering. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 13, 2, 213–34.
- WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *Computer Graphics (Proceedings of ACM SIGGRAPH 92)*, 255–264.
- ZINKE, A., AND WEBER, A. 2007. Light scattering from filaments. *IEEE Transactions on Visualization and Computer Graphics* 13, 2, 342–356.
- ZINKE, A. 2008. *Photo-Realistic Rendering of Fiber Assemblies*. PhD thesis, University of Bonn.