Fast Furry Ray Gathering

Ivan Neulander
Rhythm & Hues Studios

We present several techniques for efficiently gathering diffuse and specular reflection rays originating from hair, recently put in production at Rhythm & Hues. We refine the hair BRDFs from [Neulander 2004] into a new *cone-shell* model that is more generally applicable and realistic, yet remains simple and fast to evaluate. Next, we address the chief shortcoming of [Neulander 2004]'s hair occlusion model by coupling it with rigorous occlusion-testing of skin geometry. Finally, we deploy a per-strand shading cache to improve performance, adopting a filtering technique in the spirit of trilinear mip-mapping.

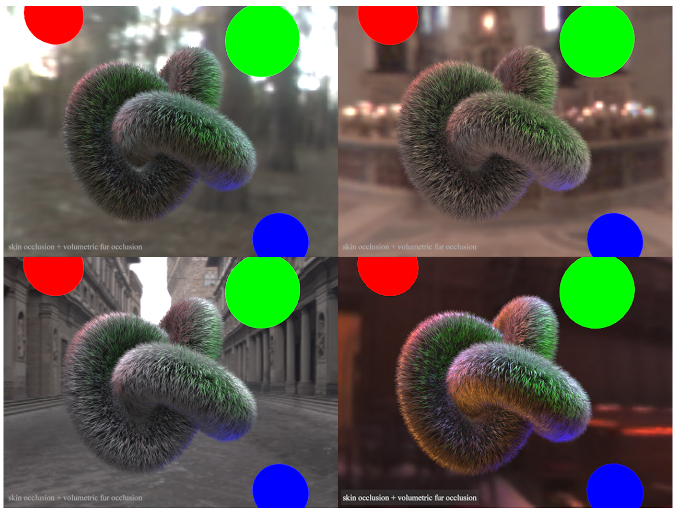

Figure 1: 640x480 images, each featuring 52k strands with 50 diffuse and 30 specular gather rays per shading point: rendered in 57 seconds per frame on an Athlon64 X2 2800MHz, single thread.

The Cone-Shell Model

Our cone-shell BRDF is based on the infinitesimal cylinder model of [Kajiya and Kay 1989] and is designed for efficient generation of scattered diffuse and specular reflection rays with nonzero weight. Unlike the IBL-centric BRDFs in [Neulander 2004], we sample rays within a nonzero solid angle: the volume between two concentric cones. In our model, $\theta_{min}, \theta_{max}$ are the angular bounds of the ray scatter cone-shell (measured from the z-axis), ranging from an infinitesimally thin surface when $\theta_{min} = \theta_{max}$ to the full sphere when $[\theta_{min}, \theta_{max}] = [0, \pi]$. Given a pair of random variables ξ_1 and ξ_2 in $U(0,1)$, we generate sample directions in a canonical coordinate system that aligns the z-axis with the hair tangent.

$$\begin{aligned} x &= \sqrt{1-z^2}\cos(2\pi\xi_2) \\ y &= \sqrt{1-z^2}\sin(2\pi\xi_2) \\ z &= \cos\theta_{max} + \xi_1\left(\cos\theta_{min} - \cos\theta_{max}\right) \end{aligned}$$

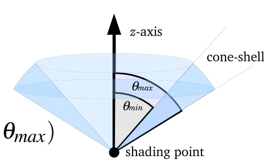

Our diffuse scattering over the full sphere matches the Kajiya model, though ours can be limited to a lesser solid angle if desired. Our specular scattering differs from Kajiya's in that our rays are restricted to the $[\theta_{min}, \theta_{max}]$ range with a linearly zero-tapered $\cos\phi$ weight[1] instead of being distributed over the full sphere with $\cos^n$ weight. This makes our specular sampling potentially more concise yet visually similar to Kajiya's for suitably chosen θ bounds.

[1] ϕ is the angle between a gather ray and the midpoint cone at $\frac{\theta_{min}+\theta_{max}}{2}$. For full-sphere scattering, the midpoint cone is the xy plane, so $\cos\phi$ equates to the sine of the angle between the ray and the z-axis, as per Kajiya's diffuse model.

SIGGRAPH 2010, Los Angeles, California, July 25 – 29, 2010.
ISBN 978-1-4503-0210-4/10/0007

For importance sampling, we can generate rays with density roughly proportional to our BRDF weight.[2] Assuming uniform incident radiance, full-sphere scattering would ideally draw ξ_1 from a Wigner Semicircle Distribution shifted to the interval [0,1]. Since incident radiance varies in practice, we need a flatter distribution than this. Based on empirical trials, we selected a *raised triangular distribution* with a pdf of $\frac{4}{9}(\frac{5}{2} - |x - \frac{1}{2}|)$. We remap ξ from $U(0,1)$ as follows:

$$\xi_1 = \begin{cases} \sqrt{\frac{1}{2}(8+9\xi)} - 2 & : \ \xi \le 1/2 \\ 3 - \sqrt{\frac{1}{2}(17-9\xi)} & : \ \xi > 1/2 \end{cases}$$

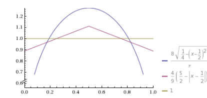

Volumetric Occlusion Approximation

Computing accurate hair occlusion tends to be slow due to the geometric complexity and shading cost involved. Although we provide this option with ray tracing, we have found that occlusion by relatively short, dense hair (e.g. most fur) can be modeled by a combination of accurate skin occlusion and approximate hair occlusion via the volumetric occlusion model in [Neulander 2004]. The latter works well at short range, is extremely fast, immune to noise, and highly art-directable. Its main drawback is its failure to account for nonlocal occlusion, as with cast shadows. By tracing rays to the skin and other hard-surface geometry, we address this with an acceptable balance of speed and fidelity.

We enhance [Neulander 2004]'s original volumetric model slightly by computing the occlusion height h_O at each control point as a projection of that point's position onto the occlusion normal $\mathbf{N_O}$. This noticeably improves accuracy for strands that run parallel to the skin, a common occurrence in practice.

Shading Cache

Traditional irradiance caching, as described by Ward, is unsuitable for hair due to abrupt surface variation perpendicular to each strand. However, irradiance and even reflected radiance vary more gradually along any given strand. Hence, we cache the results of our diffuse and specular gather operations at a discrete number of shading points along each strand, and interpolate in between.

We define a diffuse and specular cache density in terms of pixels per shading point, so that the total number of cached shading points varies with the screen-space length of each strand, subject to some limits. Our cache is populated on demand: only the cache points surrounding a requested shading sample will ever be populated.

For frame-to-frame coherence, each shading point uses the same set of canonical sample rays (driven by the same strand-based random number sequence), rotated to match the hair tangent at that point. Further, we constrain sample point counts to powers of two and evaluate the cache using both the full sample count and half the sample count, interpolating the final color between these intermediate results. This avoids popping as the cache density changes from frame to frame, similar to how trilinear texture filtering interpolates between two successive mipmap resolutions.

References

Kajiya, J. T., and Kay, T. L. 1989. Rendering fur with three dimensional textures. *SIGGRAPH Comput. Graph. 23*, 3, 271–280.

Neulander, I. 2004. Quick image-based lighting of hair. In *SIGGRAPH 2004 Sketches*, 43.

[2] Importance sampling is crucial for a specular lobe such as $\cos^n$ but less beneficial for ours, which spans a smaller nonzero range.