Art-Directable Multiple Volumetric Scattering

Magnus Wrenninge*
Pixar Animation Studios

Figure 1: *Cloud rendered using single scatter lighting, using the contrast approximation, and using our multiple scattering technique.*

1 Introduction

Multiple scattering is a crucial part of photorealistic rendering of high-albedo media. In the production rendering context, current techniques include the wavefront tracking described in [Miller et al. 2012] as well as an approximation based on modified shadow calculations, described in [Wrenninge et al. 2013] (hereon referred to as the *contrast approximation*). In order to convincingly render optically thick media with multiple scattering effects, high scatter orders must be included, upwards of one hundred bounces. Also, anisotropic effects must be handled. Existing path tracing methods, even ones using importance sampling, become computationally impractical as each camera ray spawns multiple secondary rays for which importance functions must be constructed.

Also, while physically accurate multiple scattering is a worthwhile goal, it doesn't necessarily provide the most useful solution when it comes to art-directability. We have found that changing the amount of diffusion in a cloud greatly impacts the perceived scale of the scene, but that contrast can be adjusted substantially to suit the image without compromising scale. With this distinction in mind, we developed a method that computes efficient and physically accurate diffusion, but that allows independent control over contrast.

2 Method

We compute multi-bounce diffusion of light in volumes using a path tracing formulation. Initial scatter events along the camera ray are chosen according to the *density sampling* technique from [Kulla and Fajardo 2012]. Subsequent scatter events are chosen using an approximation to Woodcock Tracking: we make the approximation $d_{\max} \approx d$, i.e. the local density d is used instead of the global maximum in order to select a next step length $l = -\ln \xi_0 / d$. Figure 2 illustrates the process. While not mathematically rigorous, the local density provides a reasonable measure of the maximum extinction $d_{\max}$ in the region of the sample. In order to choose the new direction $\vec{v}'$, we invert the phase function p such that the new scatter direction $\vec{v}' = p^{-1}(\vec{v}, \xi_1, \xi_2)$. The combination of density sampling and our random walk method allows us to reach high-order scattering events with a minimal number of computations: one interpolated voxel value per bounce and one phase function inversion, rather than a full raymarch in order to build an importance function, as would be required by the density sampling method.

Our method uses the technique described in [Wrenninge et al. 2013] to provide overall contrast control of the volume. At each bounce depth i, direct light is calculated according to Equation 1. By using only a single bank of the contrast approximation at each bounce, the overall brightness is maintained while diffusion is added. While the diffusion effect is expensive due to the increased number of indirect rays to trace, the contrast approximation calculation is computationally inexpensive, and has proven intuitive to use.

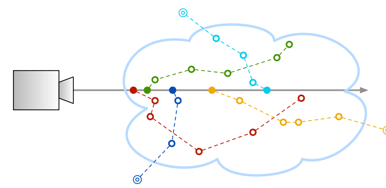

Figure 2: *Initial scatter locations are chosen along the camera ray using density sampling (solid). Subsequent scatter locations (dashed) make up a random walk based on local density and the phase function. Direct lighting is computed at each vertex.*

$$L_i = \sigma_s b^i L_{light}(\omega_i) p\left(\omega_i, \omega_o, c^i g\right) e^{-a^i \int_0^t \sigma_t(s) ds} \quad (1)$$

The three images in Figure 1 show a cloud scape rendered with various methods. The first uses single scattering only. The second shows the contrast approximation, and the third shows our method using 16 bounces and 256 rays/pixel. Render times were 5m 16s, 6m 4s and 63m 11s, respectively. Although computationally expensive, our method allows unbiased rendering of high order multiple scattering. It is also independent of the resolution of the underlying volumes, which is an advantage over methods that operate on the volume itself. Because our method works without pre-computation or other caching methods, it applies equally well to static and animated volumes. We are currently using the technique to render multiple scattering effects in both clouds and whitewater on the production of *The Good Dinosaur*.

References

KULLA, C., AND FAJARDO, M. 2012. Importance sampling techniques for path tracing in participating media. *Comp. Graph. Forum 31*, 4 (June), 1519–1528.

MILLER, B., MUSETH, K., PENNEY, D., AND BIN ZAFAR, N. 2012. Cloud modeling and rendering for Puss in Boots. In *ACM SIGGRAPH 2012 Talks*, ACM Press, New York, SIGGRAPH 2012.

WRENNINGE, M., KULLA, C., AND LUNDQVIST, V. 2013. Oz: The Great and Volumetric. In *ACM SIGGRAPH 2013 Talks*, ACM, New York, NY, USA, SIGGRAPH '13, 46:1–46:1.

*e-mail:magnus@pixar.com