


Decoupled Ray-marching of Heterogeneous Participating Media

Christopher Kulla
Sony Pictures Imageworks



(a) Brute force ray marching, 5h (b) Decoupled ray marching, 4min (c) Indirect lighting, 10min (d) With reflections, 2min

Figure 1: Ray traced heterogeneous volume densities (roughly 100 steps used in all images)

We describe the integration of heterogeneous participating media into our production ray tracer. We specifically focus on the problem of computing the lighting without incurring a quadratic number of shader evaluations. The main difficulty in rendering volumes comes from the inability to analytically importance sample the transmission term:

$$\tau(t) = e^{-\int_0^t (\sigma_s(x) + \sigma_a(x)) dx} \quad (1)$$

The classical ray-marching algorithm [Perlin and Hoffert 1989] has a long history of use in production rendering due to its simplicity. The volume equation is solved by marching down the ray in fixed or adaptive steps, evaluating the lighting at each sample. Unfortunately all shadow rays must themselves be ray marched, leading to a quadratic number of shader evaluations.

The typical solution to this problem has been to employ light caching techniques such as deep-shadow maps [Lokovic and Veach 2000]. Unfortunately these techniques require multiple passes, and are prone to many sources of artifacts. In contrast, our method reduces the quadratic complexity without giving up the benefits of accurate ray traced lighting.

Ray-marching Revisited We begin by observing that for homogeneous media, unbiased solutions are practical. This is due to the fact that equation 1 simplifies to: $e^{-t(\sigma_s + \sigma_a)}$ which can be both evaluated and sampled analytically, leading to simple Monte Carlo solutions. These benefits remain even if we confine the homogeneous media to a region of space, only the normalization constants change on the associated probability density function (PDF).

Our key insight is to observe that we can reformulate the task of ray-marching as transforming an unknown, spatially varying volume into a list of piecewise homogeneous segments. This gives us access to inexpensive analytical formulas for evaluating and sampling equation 1 at arbitrary points.

Our algorithm begins by marching down the ray in fixed (or adaptive) steps. We pick a random sample point within the segment of each step and run the volume shader. The volume properties σ_s and σ_a are recorded into an array from which we can build a piecewise linear PDF proportional to $\sigma_s(t)\tau(t)$. We can then perform light sampling outside the ray-marching loop (hence *decoupled*) by choosing positions along the ray using this PDF. For each light sample we trace shadow rays to the lights, which in turn runs the first half of our algorithm to estimate the transmission factor (traditional

ray marching). As the lighting calculations are decoupled, we can also evaluate indirect lighting recursively at a reduced cost.

Time Complexity If we assume that ray-marching requires $O(N)$ shader evaluations per ray, the traditional ray marching algorithm performs $O(N^2)$ shader evaluations. Our method only requires $O(N + LN)$ evaluations (L being the number of light samples) which is a substantial reduction when $L \ll N$. In fact we typically use $L = 1$ as we need to trace many primary rays for smooth anti-aliasing and motion blur.

Bias Naturally our method introduces some bias, however it is explicitly controlled by the step size parameter and can be made arbitrarily small. In fact, as the step size approaches the Nyquist rate of the volume, our algorithm converges on a ground truth result. In our implementation, the shader controls how many ray-marching samples are required within the bounds it occupies along the ray to allow for volumes of different frequency contents to be mixed in the same scene. We also point out that the homogeneous limit case is handled without bias as a single ray-marching sample captures the volume properties exactly.

In contrast to unbiased integration schemes [Yue et al. 2010] that use rejection sampling to invert equation 1, we do not require any knowledge of upper bounds on volume properties which is essential to deal with procedural and out-of-core volume representations. Moreover, since our PDF is proportional to $\sigma_s(t)\tau(t)$ instead of $\tau(t)$, it does not waste any samples on regions that do not scatter light.

References

- LOKOVIC, T., AND VEACH, E. 2000. Deep shadow maps. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, SIGGRAPH '00, 385–392.
- PERLIN, K., AND HOFFERT, E. M. 1989. Hypertexture. *SIGGRAPH Comput. Graph.* 23 (July), 253–262.
- YUE, Y., IWASAKI, K., CHEN, B.-Y., DOBASHI, Y., AND NISHITA, T. 2010. Unbiased, adaptive stochastic sampling for rendering inhomogeneous participating media. *ACM Trans. Graph.* 29 (December), 177:1–177:8.