

Extinction Transmittance Maps

Pascal Gautron *

Cyril Delalandre *

Jean-Eudes Marvie *

Technicolor Research & Innovation



Figure 1: Production-quality clouds ($600 \times 400 \times 300$) rendered at 45 fps. Our solution gathers local extinction for a simple formulation of light transmission in heterogeneous media. We provide high quality results while remaining independent from the volume representation.

The interaction between light and participating media involves complex physical phenomena such as light absorption and scattering. The radiance transmitted through a medium then depends on the variations of scattering and extinction along potentially complex light paths, yielding soft light shafts and shadowing (Figure 1).

Computing light scattering in these media usually requires complex offline computations [Cerezo et al. 2005]. Some real-time applications are based on heavy precomputations [Zhou et al. 2008]. Some others introduce restrictions such as approximate diffusion schemes, or specific volume representations.

In particular, [Jansen and Bavoil 2010] extend the concept of deep shadow maps [Lokovic and Veach 2000] by representing the variations of opacity within the medium in Fourier space: each pixel of the *Fourier Opacity Map* stores a set of projection coefficients. Opacity variations are then unprojected from the coefficients to evaluate the overall opacity of the medium for each visible particle. This technique is highly effective, and the map is built by accumulating the contributions of each particle independently using alpha blending. However, this method only represents an opacity information and cannot account for actual light scattering.

This limitation is raised in [Delalandre et al. 2011]: instead of representing opacities, each pixel of the *Transmittance Function Map* (TFM) stores a Fourier transform of the medium transmittance along light rays. Single scattering is then estimated in a ray marching process, in which the light reduced intensity is deduced from the transmittance values. While providing accurate results and real-time performance, the evaluation of the medium transmittance at a point requires the knowledge of the overall extinction along the entire light path. The generation of the map then requires a ray marching through the participating medium. This technique is then usable only for voxelized media, overlooking other representations such as the dynamic particle clouds massively used in video games.

Extinction Transmittance Maps:

We introduce *Extinction Transmittance Maps* (ETM), a technique combining the advantages of both *Fourier Opacity Maps* and *Transmittance Function Maps* while avoiding their respective drawbacks.

As for [Jansen and Bavoil 2010] and [Delalandre et al. 2011], our method borrows from the principle of shadow mapping where a virtual camera is oriented towards the medium from the location of the light source. This camera is used to create the *Extinction Transmittance Map*, into which the contents of the medium are rendered to build a set of Fourier coefficients representing the local variations of the extinction parameters (Figure 2). To generate the final image we reformulate the transmittance function to evaluate the contribution of each visible volume sample directly from the ETM.

* {pascal.gautron, cyril.delalandre, jean-eudes.marvie}@technicolor.com

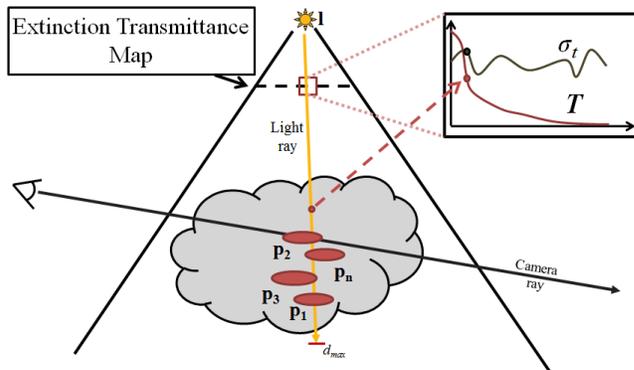


Figure 2: Rendering a participating medium involves evaluating the transmittance T from the light source to each sample point \mathbf{p}_i to estimate light transfer. The ETM is generated from the local extinction σ_t at volume samples such as particles. Analytical integration then provides the global transmittance T along light rays.

When Projecting Local Extinction...

In [Delalandre et al. 2011] each pixel of the TFM is generated by marching through the medium to evaluate its transmittance along light rays. In the case of particle-based media, marching through the particles is not tractable without prior voxelization. To overcome this problem, our approach is based on the projection of the local extinction coefficient of the medium as a set of Discrete Cosine Transform (DCT) coefficients c_j :

$$c_j = \frac{1}{d_{max}} \sum_i d_i \sigma_t(\mathbf{p}_i) b_j(x_i) \quad (1)$$

where d_{max} is the maximum traversal length of the light ray. The j -th DCT function is $b_j(x_i) = k_j \cos\left((2x_i + 1) \frac{i\pi}{2d_{max}}\right)$, k_j being its normalization factor. Each particle i is described by its location \mathbf{p}_i , distance to the light source x_i and traversal distance d_i .

A representation of the extinction function can then be obtained from any kind of volume representation, avoiding the need for a ray marching. In the next step we reformulate the transmittance directly in terms of the projection coefficients.

... Leads to Global Transmittance

The medium transmittance along a light ray follows an exponential rule according to the extinction σ_t , which can be expressed using the coefficients of the ETM:

$$T(\mathbf{p}_i) = \exp\left(-\int_1^{\mathbf{p}_i} \sigma_t(\mathbf{z}) d\mathbf{z}\right) = \exp\left(-\sum_j \int_1^{\mathbf{p}_i} c_j b_j(x_z) d\mathbf{z}\right) \quad (2)$$

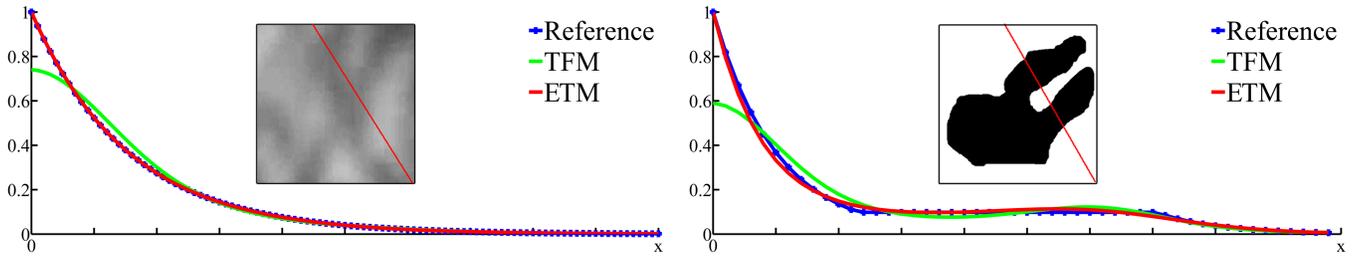


Figure 3: Transmittance function along a light ray through two media: smoke (left) and a test medium with sharp variations (right). Compared to the actual transmittance (blue), our method (red) features less oscillations than TFM (green) in both cases.

Each definite integral can then be rewritten in terms of the indefinite integral B_j of b_j , yielding a new formulation for T :

$$T(\mathbf{p}_i) = \exp\left(-\sum_j c_j (B_j(\mathbf{p}_i) - B_j(\mathbf{l}))\right) \quad (3)$$

where:

$$B_j(\mathbf{p}) = \frac{k_j d_{max}}{j\pi} \left[\sin\left(\left(1 + 2x_p\right) \frac{j\pi}{2d_{max}}\right) - \sin\left(\frac{j\pi}{2d_{max}}\right) \right]$$

This unprojection scheme emphasizes the duality of our representation: a set of potentially incoherent extinction samples is projected into the DCT basis regardless of the kind of representation used for the volume. Then, the transmittance expressing the effect of each particle on all the others along the same light ray is recovered analytically from the projection coefficients.

Our formulation also subtends a major efficiency difference between TFM and our approach. As shown in Equation 2 the transmittance function follows an exponential rule which gets encoded within the coefficients of the TFM. On the contrary, the exponential remains in its analytical form in Equation 3. Thus the coefficients only encode the variations of the extinction along a light ray, which tend to be oscillatory in many media. This property makes the DCT a method of choice for representing such variations, while Equation 3 naturally preserves the exponential shape of the transmittance (Figure 3).

In the spirit of [Jansen and Bavoil 2010], final images are obtained by rendering the particles from back to front, blending their contributions along camera rays. In voxelized media a simple ray marching is used. For each particle or sample point the medium transmittance is evaluated from the *Extinction Transmittance Map*. The amount of light scattered towards the camera is then deduced from the light reduced intensity and the local scattering properties.

Results

Our algorithm offers real-time performance on a 3.6GHz Xeon with an nVidia GeForce GTX480 GPU. The presented images were rendered at a resolution of 1280×720 using a 1024×1024 ETM containing 4 coefficients per pixel. We compare our approach to the TFM and a GPU-based brute force reference solution for single scattering. For comparison purposes, our tests are performed by ray marching through a voxelized medium. As shown in Figure 4 our method tends to reduce ringing artifacts, typically present in images of very dense media generated using TFM. Such artifacts appear as stripes perpendicular to the lighting direction, and are particularly visible in the difference images. In our approach the error tends to be spread over the entire medium, hence being less noticeable especially in animated media.

Regarding rendering speed our method performs similarly to TFM, while avoiding the restrictions on the representation of the participating media.

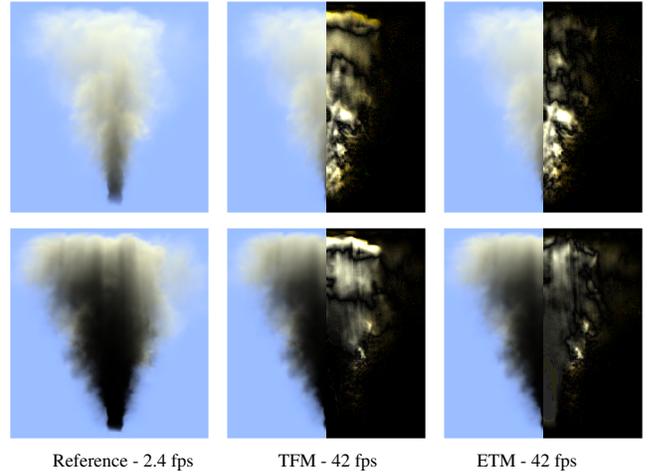


Figure 4: A plume of smoke rendered using *Extinction Transmittance Maps* and TFM, both with 4 coefficients per pixel. The right part of the images show the difference with a reference solution, with pixel intensities multiplied by 40 for legibility. The reconstruction errors plotted above typically appear as horizontal stripes when using TFM with dense media.

Production rendering also benefits from our method for higher rendering speed without compromising image quality. The genericity of our method allowed us to implement our approach as a shader for Pixar’s RenderMan using its built-in, representation-agnostic volume dicing. Our method is then implicitly compliant with other PrMan components such as depth-of-field and motion blur.

Extinction Transmittance Maps can then benefit to a wide range of domains, not only for real-time applications such as video games and interactive previsualization, but also for production rendering.

References

- CEREZO, E., PEREZ, F., PUEYO, X., SERON, F., AND SILLION, F. 2005. A survey on participating media rendering techniques. *The Visual Computer* 21, 5, 303–328.
- DELALANDRE, C., GAUTRON, P., MARVIE, J.-E., AND FRANÇOIS, G. 2011. Transmittance function mapping. In *Proceedings of I3D*, 31–38.
- JANSEN, J., AND BAVOIL, L. 2010. Fourier opacity mapping. In *Proceedings of I3D*, 165–172.
- LOKOVIC, T., AND VEACH, E. 2000. Deep shadow maps. In *Proceedings of SIGGRAPH*, 385–392.
- ZHOU, K., REN, Z., LIN, S., BAO, H., GUO, B., AND SHUM, H.-Y. 2008. Real-time smoke rendering using compensated ray marching. In *Proceedings of SIGGRAPH*, 1–12.