

# Volume-Aware Extinction Mapping

Pascal Gautron\*

Cyril Delalandre\*

Jean-Eudes Marvie\*

Pascal Lecocq\*

Technicolor Research & Innovation



**Figure 1:** Our technique renders this animated scene at 1.5fps, featuring volcano smoke ( $1024^3$ ) as well as clouds and plumes ( $75 \times 512^3$ ).

The simulation of the lighting effects produced by the interaction of participating media with the light contributes to the production of visually compelling effects such as translucence and volumetric shadowing. However, the complex inner structure of participating media requires vast amounts of memory for storage and costly computations for rendering.

The cost of offline lighting estimation within volumes is usually reduced by caching strategies such as Deep Shadow Maps [Lokovic and Veach 2000], in which lists of volume samples are built to represent light attenuation. These lists are then ordered, filtered and compressed prior to rendering. Real-time rendering methods such as [Jansen and Bavoil 2010; Delalandre et al. 2011] avoid the need for list management by projecting light attenuation into a Fourier basis. While effective, the empty sections between participating media are encoded as well in the Fourier transform, leading to ringing artifacts when used in complex environments.

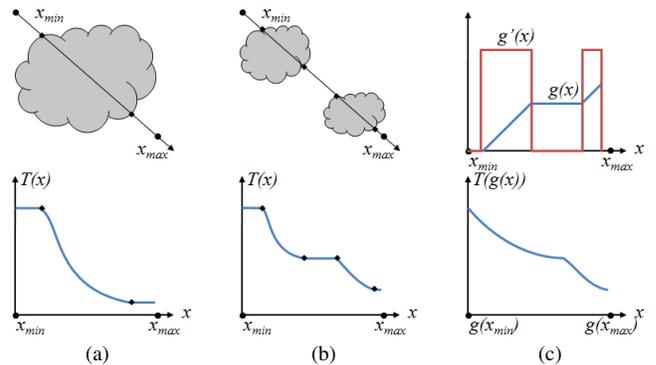
*Volume-Aware Extinction Maps* project and integrate spatially unordered sets of volume samples within a volume-aware functional space for high quality interactive rendering of production scenes.

## Volume-Aware Pseudometric

Our technique borrows a principle of [Lokovic and Veach 2000]. Empty sections of the scene can be cancelled out without any impact on the representation quality (Figure 2c). Intuitively this is equivalent to replacing the euclidean distance by a pseudometric only accounting for the sections traversing the participating media.

Instead of managing an explicit list, we aim at generating a *pseudometric map*, for which each texel contains a representation of the pseudometric using a small set of Fourier coefficients. This map is obtained by rendering the bounding boxes of the participating media, and projecting the traversal length of each ray through the boxes, yielding  $g'(x)$  (Figure 2). This function is then integrated through analytical integration of the Fourier basis functions, yielding a representation of  $g(x)$ .

We use a similar principle to represent the transmittance function along a ray. While [Delalandre et al. 2011] require a successive traversal of the media to evaluate transmittance, we first represent each available sample of local extinction in terms of the pseudometric, that is  $\sigma_t(x) = \sigma_t^g(g(x))$ . Then, we project each sample of this unsorted set into a Fourier basis, yielding a representation of the extinction along light rays. Then transmittance  $T$  at a point is then obtained by analytically integrating the projected extinction.



**Figure 2:** While the representation bounds can tightly fit a single medium (a), the transmittance function  $T$  features constant sections between multiple media (b). Using an appropriate pseudometric  $g$ , we represent transmittance without gaps (c).

## Results

We implemented Volume-Aware Extinction Maps to render  $1280 \times 720$  images on a Intel Xeon X5680 3.36GHz processor running a Nvidia GeForce GTX580 GPU. The skies of Figure 1 comprise a total of 10.5GB of volumetric data. We manage this complexity by introducing an easy-to-use out of core management of volumetric data for adaptive rendering. We also leverage the duality of our representation for interactive image-based lighting and multiple scattering. The Volume-Aware Extinction Maps used in these images have a resolution of  $1024 \times 1024$  with a total of 12 coefficients per pixel (4 for  $g$ , 8 for  $\sigma_t$ ).

Our technique finds a particular use for real-time editing and navigation in environments featuring massive volumetric datasets. As we do not introduce any precomputation, the components of the scene can be modified in real-time without impacting rendering speed. This ability makes our method a highly valuable tool for the previsualization and tuning of complex volumetric environments for production rendering.

## References

- DELALANDRE, C., GAUTRON, P., MARVIE, J.-E., AND FRANÇOIS, G. 2011. Transmittance function mapping. In *Proceedings of the 13D Symposium*, 31–38.
- JANSEN, J., AND BAVOIL, L. 2010. Fourier opacity mapping. In *Proceedings of the 13D Symposium*, 165–172.
- LOKOVIC, T., AND VEACH, E. 2000. Deep shadow maps. In *Proceedings of SIGGRAPH*, 385–392.

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