Single Scattering in Heterogenous Participating Media

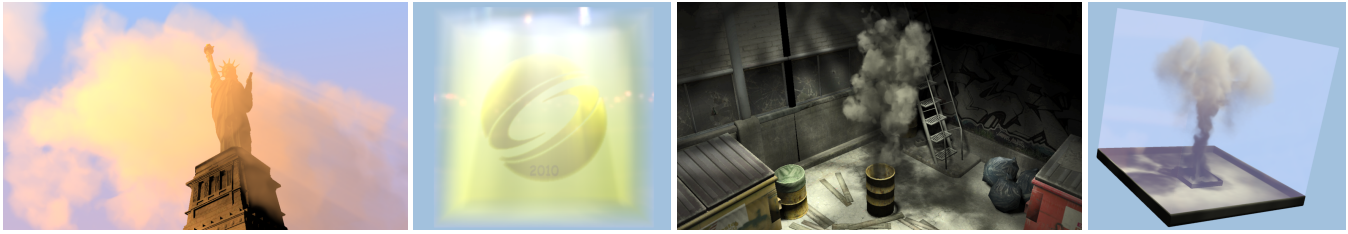Cyril Delalandre   Pascal Gautron   Jean-Eudes Marvie   Guillaume François

{cyril.delalandre, pascal.gautron, jean-eudes.marvie}@technicolor.com   guillaume-f@moving-picture.com

Technicolor Research & Innovation   The Moving Picture Company

(a) Liberty Cloud - 10 fps   (b) Siggraph Volume- 17 fps   (c) Smoke - 20 fps   (d) Smoke in a Box- 17 fps

Figure 1: Our algorithm introduces *attenuation function maps* for the computation of light scattering within heterogeneous media - 500 steps

The interaction between light and light-transmitting objects, known as *participating media*, involves complex physical phenomena such as light absorption and scattering. Media such as clouds, smoke and translucent materials often feature heterogeneous scattering properties. Hence the radiance transmitted by the medium potentially depends on such varying properties on the entire light paths, yielding soft light shafts and opacity variations (Figure 1).

Computing light scattering in these media usually requires complex offline computations [Cerezo et al. 2005]. Real-time applications are either based on heavy precomputations [Zhou et al. 2008], limited to homogeneous media [Gautron et al. 2009] or relying on simplistic rendering techniques such as billboards. We propose a generic method for fast estimation of single scattering within heterogeneous participating media. Introducing the concept of *attenuation function maps*, our method leverages graphics hardware to support dynamic light sources, viewpoints and participating media interactively. Our method also accounts for the shadows cast from solid objects, providing a full-featured solution for fast rendering of participating media which potentially embrace the entire world.

The Attenuation Function Map

The classical method for estimating single scattering in heterogeneous media relies on a dual ray marching. When entering the medium, the viewing ray is sampled into n sample points P_i, $i \in \{1..n\}$. For each sample point, the attenuation of the light incoming at P_i is computed by ray marching along the ray from P_i to the light source. However, in most cases, this computation-intensive method cannot be used for real-time or even interactive high quality rendering. Our method cuts the cost of reduced intensity calculations by generating an *attenuation function map* which is then queried to estimate single scattering in the medium.

The first step of our method is inspired from shadow mapping, a virtual camera being placed at the location of the light source. For each pixel of the image captured by this camera, we first determine the entry and exit points K_{in} and K_{out} of the ray in the bounding box of the medium (Figure 2). As we aim at representing the reduced intensity at any point K_i along $K_{in}K_{out}$, we sample the light attenuation on $K_{in}K_{out}$ at a number of points K_i on the light ray. To avoid the costs of explicit storage of the attenuation function values, we project this function using the well-suited Discrete Cosine Transform. The projection coefficients are computed on graphics hardware, and are then stored in the channels of multiple floating-point textures. Compared to the explicit storage of the attenuation values, this method both reduces the storage and provides continuous, smooth attenuation estimates. As we also focus on the influence of opaque objects on the lighting within the medium, we generate a classical shadow map to account for such shadowing.
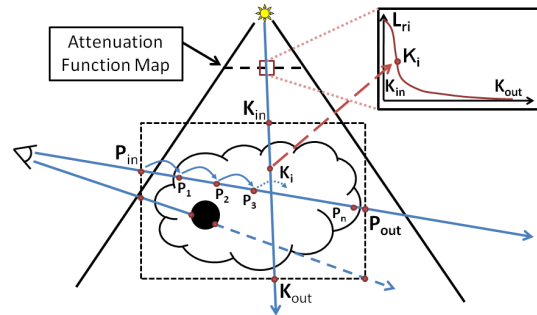

Figure 2: The *attenuation function map* is used to directly obtain the reduced intensity value at each point P_i.

In the second step, we render the scene from the actual viewpoint. We first compute the amount of light reflected from opaque objects. The incoming lighting at each surface point is obtained by first querying the classical shadow map and evaluating the attenuation function. A second pass performs the single scattering computations: For each pixel of the image, we compute the entry and exit points P_{in} and P_{out}, and then perform a classical ray marching. However, at each sample point P_i, the reduced light intensity is directly fetched from the shadow and attenuation function maps.

Results

Our algorithm offers interactive performance on a 3.6GHz Xeon with an nVidia GeForce GTX280 GPU. The presented images were rendered at a resolution of 1280×720 using a 1024×1024 attenuation function map with 16 coefficients per pixel. Note that the sole tuning of the number of coefficients and marching steps is sufficient to generate images ranging from real-time to production-quality, providing a unified and generic solution for fast high quality rendering of heterogeneous participating media. Future work will particularly consider multiple scattering and image-based lighting.

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.

GAUTRON, P., MARVIE, J.-E., AND FRANÇOIS, G. 2009. Volumetric shadow mapping. In *SIGGRAPH 2009 talks*.

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 2008*, 1–12.