Efficient Physically-Based Perceptual Rendering of Participating Media (sketches_0310)

Diego Gutierrez,[*] Oscar Anson, Francisco J. Seron
Light Simulation Lab
University of Zaragoza

Veronica Sundstedt,[†] Alan Chalmers
University of Bristol

1 Introduction

To significantly reduce physically-based rendering times for participating media, we propose a novel perceptual strategy based on the combination of a saliency map with our new physically-based extinction map (X-map), which stores in image-space the exponential decay of light in the medium. This combination is then used to guide a selective rendering of the scene, with more accurate estimates in the most perceptually important parts of the scene, without visible degradation. The novelties of this work can then be summarized as a) the introduction of the X-map concept and b) its combination with a saliency map to guide a perception-based renderer for inhomogeneous participating media. Results show a decrease of almost 60% in rendering times for our test scene.

2 The X-map and the saliency map

In the general case of inhomogeneous participating media, the intensity of light as it transverses the medium is reduced by $e^{-\tau(x,x')}$, where $\tau(x,x') = \int_x^{x'} \sigma_t(s)ds$ is the *optical length* (with σ_t being the extinction coefficient and s being the distance through the medium). This attenuation can be evaluated by ray marching. Therefore, independently on how much irradiance reaches a given object, only a fraction of it will finally reach the eye. The first idea of this work is to pre-compute that fraction beforehand, storing in image-space attenuation values in an extinction map which we have named the X-map.

To obtain the X-map, we shoot rays from the eye into the scene, ray marching through the medium until the first intersection, and save the result of the exponential decay in the X-map, representing the percentage of light (both direct and diffuse) which will reach the eye for each pixel (or sub-pixel if the resolution of the map is increased). For homogeneous media, the attenuation is just $e^{-\sigma_t s}$, and faster ray tracing can be used instead of ray marching. In either case, we also save the distance to the intersection in a Z-buffer, thus allowing for almost instantaneous recomputations of the X-map if the description of the medium in the scene changes, since intersections do not need to be recalculated at each step.

To complete our setup for selective rendering of participating media, we further propose a novel combination of this X-map with a saliency map, which indicates the areas that automatically attract our attention in the environment [Itti et al. 1998]. This combination is given by $XS(\omega_x, \omega_s, op)$, where XS represents the combined X- and saliency maps, and ω_x and ω_s are coefficients applied to the values in the X-map and the saliency map respectively, which allow for different weighting to define relative importance of the maps. The coefficients are combined through the operator op (in this case an addition operator). The combined map therefore represents saliency modulated by the decay of light in the media, and will guide the selective rendering process by spending computational resources in areas of higher perceptual importance.

Additionally, by building a Cumulative Distribution Function (CDF) of the resulting combined map, the image can be importance-sampled during the rendering process, launching more rays in the steep areas of the CDF, which correspond to perceptually important parts of the scene. The user can then specify the maximum rendering time allowed for the image, or the maximum number of rays to be shot, and the load will be distributed according to the perceptual importance of the image.

3 Results and conclusions

We have introduced a novel concept, the X-map, to precompute light attenuation in inhomogeneous participating media (with a straightforward simplification for the homogeneous case). Combined with the saliency map, we can perform selective rendering of participating media scenes, with high quality antialiasing settings for the salient features foreseen to be less attenuated by the medium. We have tested our implementation on a scene rendered at 400x400 pixel resolution using volume photon mapping. We first rendered it without applying the maps presented here, shooting nine rays per pixel. We then applied our method, selectively rendering at 1,2...9 rays per pixel according to the values in the combined X- and saliency map (with ω_x and ω_s having equal values of 0.5). This second image was rendered in just 41% the time it took the first one, without any degradation in perceptual quality as can be seen in the figures included. The generation of the X-map, the saliency map and the combined map added an overhead of less than five seconds.

In this case, the acceleration using photon mapping occurs during the more expensive radiance estimate phase, whereas the relatively less costly photon tracing phase remains unchanged. The approach, however, is general: any other algorithm can be used to render the final image, once the maps have been generated. As immediate future work, a case-based study on the optimal values of ω_x and ω_s is to be performed, along with more sophisticated perception-based operators op.

References

ITTI, L., KOCH, C., AND NIEBUR, E. 1998. A model of Saliency-Based Visual Attention for Rapid Scene Analysis. In *Pattern Analysis and Machine Intelligence*, vol. 20, 1254–1259.

---

[*] e-mail: diegog@unizar.es
[†] e-mail: veronica@cs.bris.ac.uk