9

Conclusions

"This is not the end. It is not even the beginning of the end. But it is, perhaps, the end of the beginning."

—Winston Churchill, 1874–1965

IN this dissertation we sought to develop efficient algorithms for simulating light transport within arbitrary participating media. We provided missing derivations for the split-sphere model, which is central to the success of the irradiance caching algorithm. We also provided full derivations of the irradiance gradient computation, extended them to participating media, and have developed a reformulation of the volumetric photon mapping method. These theoretical examinations have provided insights which have directly led to several significant steps towards our goal.

We introduced two concrete approaches for efficient light transport within participating media: volumetric radiance caching and volumetric photon mapping using beam radiance estimation. These techniques share many similarities. Both approaches are based around a Monte Carlo ray tracing framework. This grants the algorithms a level of generality which is not attainable with finite element or analytic methods. Both methods also exploit the smooth nature of illumination within participating media by reusing illumination computations across the scene. This property is key to the efficiency of the methods. The techniques impose no restrictions on the properties and representation of the medium being rendered. They can efficiently compute single and multiple scattering in heterogeneous media with anisotropic phase functions.

We have also introduced improvements to the irradiance caching method to properly account for participating media. We derived gradients of the radiative transfer equation which correctly account for scatter, absorption, and emission, and which incorporate changes in occlusion. This allows irradiance caching to be used effectively even in scenes containing participating media.

Volumetric radiance caching and photon mapping also contain significant differences. Radiance caching is a view-dependent approach which only computes lighting within the visible portion of the scene. This makes the algorithm well suited for scenes with a large extent or scenes where only a very small portion of the scene is visible to the camera. In contrast, photon mapping simulates lighting within the whole scene. This can be a significant drawback if there are many light sources that do not contribute much to the chosen view. On the other hand, photon mapping is the only available method which can efficiently simulate volumetric caustics. These differences actually provide a great benefit by making the approaches complementary. We believe a more general and robust algorithm can be achieved by using volumetric radiance caching as a final gather pass for photon mapping. The photon mapping reformulation we presented also suggests interesting possibilities for other custom-designed radiance estimates. A final gather radiance estimate could be designed which collects photons along all hemispherical beams instead of just along the beam of a single ray.