

## **Pixar Project Description RIPS 2007**

When Computer Graphics (CG) films are rendered it is often desirable to include the effects of global illumination, specifically the illumination of geometry with indirect light that has been scattered from one or more surfaces on its way from the light source to the imaged object. While more efficient mechanisms exist for special cases, the most general approaches to solving for the indirect light amplitude is Monte-Carlo integration of the light transport over all of the possible scattering paths. This approach is computationally quite expensive due to the intrinsically slow convergence rate of such methods and the lack of good tools for importance sampling of rays departing the light sources.

Most modern stochastic renderers incorporate an approach known as photon mapping where the indirect illumination process is divided into a scatter pass where direct light is deposited onto a surface based illumination cache (photon map) followed by a gather pass where the visibility of the photon map surfaces is evaluated from the target imaging point. Variations of photon mapping are widely considered the most general way to render still images with indirect illumination. When rendering animated sequences with photon mapping there are a number of issues. Most of the computation time in a photon mapping renderer is incurred in the gather phase and is determined by the need to reduce sampling noise to values as small as possible. Sampling noise is particularly objectionable for animated sequences since single frame sample noise that might appear as non objectionable grain in a still image can appear as distracting background buzzing in an animation.

Recent work at Pixar (see Meyer and Anderson, Siggraph 2006, Statistical Acceleration for Animated Global Illumination, available at [graphics.pixar.com](http://graphics.pixar.com)) has focused on using space – time information to improve estimates of diffuse light transport when using stochastically calculated indirect light transport for rendering sequences of images. This approach has been shown to achieve significant speed improvements for the same image quality and can often determine the light transport for an entire sequence at roughly the same cost as a single frame of that sequence. A current research area associated with this work is the interaction between our approach and importance sampling for the renderer, which is the identification of regions of the image where additional samples should be computed due to local high time frequency variability in the light transport paths.

We would like the Pixar IPAM team to formulate and test some approaches to address this topic in a simple rendering environment. The project combines elements of digital signal processing, optimal filtering, and estimation theory.