

# Noise-Free BSSRDF Rendering On The Cheap (sap\_0343)

Anders Langlands\*  
The Moving Picture Company

Tom Mertens†  
Hasselt University



(a) 615K samples

(b) 615K samples using our technique

(c) reference with 7.5M samples

**Figure 1:** (a) Detail of a character rendered using 615K samples. The low sampling rate introduces artifacts. (b) The same sampling rate, but rendered with our technique. Despite using less than a tenth of the samples, the image rendered using our technique is virtually indistinguishable from (c), the reference solution. The insets show a closeup of the right ear.


Convincingly rendering translucent materials is a precious feature for production-quality rendering. The most successful techniques are based on the hierarchical pointcloud approach of Jensen and Buehler [2002], which is based on the dipole BSSRDF [Jensen et al. 2001]. This method is efficient to evaluate and naturally incorporates all the lighting effects typically deployed in a production renderer such as soft shadows and global illumination. Our PRMan BSSRDF implementation at MPC bakes illumination samples from micropolygon vertices during a pre-render to use as input for the hierarchical evaluation of the BSSRDF surface integral. We notice that for a realistically scaled character in a typical film frame the required sample density to generate a smooth result is often an order of magnitude higher than that required to accurately represent the incident illumination. Such dense pointclouds are very expensive to generate (due to the often complex lighting involved) and store (due to their large size). Using lower sampling rates results in splotchy noise and flickering during animation, and is exacerbated by the non-uniform nature of PRMan’s micropolygon generation.

To render a translucent appearance, we need to compute a surface integral  $\int R_d(r)E(\mathbf{x})d\mathbf{x}$ , which basically convolves the BSSRDF  $R_d$  with the incident illumination  $E$  [Jensen and Buehler 2002]. The culprit for causing the splotchy noise is the steep falloff in  $R_d$ . We take the approach of Mertens and Langlands [2007], reducing the influence of nearby samples by weighting samples that lie within a disk of radius  $\alpha$  around the point of interest  $x$  using function  $\bar{w}$ . The nearby samples are shown in red in the figure on the right. To account for the local missing information, we integrate again over the disk using weighting  $w$ . By definition,  $w + \bar{w} = 1$ , and this enables us to split the BSSRDF integral into a local and global part, respectively:

$$\underbrace{\int w(r)R_d(r)E(\mathbf{x})d\mathbf{x}}_{2\pi E \int_0^\alpha w(r)R_d(r)rdr} + \underbrace{\int \bar{w}(r)R_d(r)E(\mathbf{x})d\mathbf{x}}_{\text{computed using pointcloud}}$$

By approximating the local incident illumination by a constant value  $E$ , the first (local) term can be reduced to a sim-

ple 1D integral after conversion to polar coordinates, and precomputed for any disk radius  $\alpha$ . The global term is calculated by a standard traversal of the octree [Jensen and Buehler 2002], with the addition that each point to be integrated is weighted by  $\bar{w}(r)$ .



We added the local disk integral to our existing BSSRDF shading plugin for PRMan. For each point to be shaded,  $\alpha$  is chosen to be proportional to the  $k$ th nearest neighbour in the pointcloud (we choose  $k=10$  empirically). The irradiance  $E$  is calculated by the surface shader at  $\mathbf{x}$  and passed through to the plugin. A lookup table for the local term is generated at frame start, taking only a couple of seconds.

	Low Res.	High Res.
#samples	615,902	7,518,091
Size in memory	86Mb	361Mb
Size on disk	22Mb	259Mb
Point Sampling	7 min	49.5 min
Rendering	8.5 min	18 min
Total	15.5 min	67.5 min

The disk integral can be added transparently to an existing pointcloud BSSRDF implementation for immediate savings in computational cost and storage requirements. The above table shows performance statistics for Fig. 1. While the current implementation uses the dipole diffusion approximation, the multipole model or indeed any other scattering kernel could be used instead.

## References

- JENSEN, H. W., AND BUEHLER, J. 2002. A rapid hierarchical rendering technique for translucent materials. In *SIGGRAPH 2002*.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *SIGGRAPH 2001*.
- MERTENS, T., AND LANGLANDS, A., 2007. Noise reduction for rendering with diffusion BSSRDFs. Submitted.

\*e-mail:anders-l@moving-picture.com

†e-mail:tom.mertens@uhasselt.be