

BSSRDF Importance Sampling

Ian King
Solid Angle

Christopher Kulla
Sony Pictures Imageworks

Iejandro Conty
Sony Pictures Imageworks

Marcos Fajardo
Solid Angle



figure 1: Simple BSSRDF without (a) and with (b) axis MIS (25 spp). Complex BSSRDF without (c) and with (d) weight MIS (16 spp).

Light propagation within translucent materials can be described by a BSSRDF [Jensen et al. 2001]. The main difficulty in integrating this effect lies in the generation of well-distributed samples on the surface within the support of the rapidly decaying BSSRDF profile. Jensen suggested that these points could be importance sampled but did not provide implementation details. More recently, Walter et al. [2012] and Christensen et al. [2012] proposed other sampling methods which can still suffer from excessive variance.




figure 2: Geometric setup for Equation 1.

Disk based sampling Our approach uses a disk distribution of samples which we project against the surface geometry using probe rays. We bound the radial term of the BSSRDF with a maximum distance R_m to define a bounding sphere around the shading point P . We cast probe rays against this sphere along an axis \mathbf{V} perpendicular to the disk and compute the incoming irradiance at all intersection points we find inside the sphere. The contribution of each point is modulated by:

$$\frac{R_d \|\mathbf{P}_{hit} - \mathbf{P}\|}{\text{pdf}_{\text{disk}}(r)} \frac{1}{|\mathbf{V} \cdot \mathbf{N}_{hit}|} \quad (1)$$

which accounts for the probability of generating the point, the change in differential area measure and the BSSRDF itself. We omit the view-dependent terms here for simplicity. We choose the pdf to be proportional to R_d , as for flat surfaces this results in perfect importance sampling. We have developed equations for both cubic and gaussian profiles. For a single planar gaussian defined as $R_d(r) = 2\pi v e^{-r^2/2v}$, we use the following normalized pdf and warping equation to get $r \in [0, R_m)$ from a random sample $\epsilon \in [0, 1)$:

$$\text{pdf}_{\text{disk}}(r) = R_d(r) / \left(1 - e^{-R_m^2/2v}\right)$$

$$r = \sqrt{2v \log \frac{1}{1 - \epsilon} - 1 - e^{-R_m^2/2v}}$$

The angle around the disk is chosen uniformly. We found $R_m = \sqrt{v}/12.46$ to be a reasonable compromise between accuracy and speed.

axis and Weight Multiple Importance Sampling Equation 1 will lead to high variance in regions of high curvature, or around sharp corners, because the dot product in the denominator can become arbitrarily small. In fact areas perpendicular to the disk may not be hit at all! To compensate for this, we randomly choose the axis \mathbf{V} among the three coordinate axes of the local shading frame: \mathbf{N} , \mathbf{T} , \mathbf{B} . When computing the contribution of each hit point, we account for the probability of having found the same point along the two other axes by multiple importance sampling. Since the three axes form an orthonormal basis, the denominator can never become small which ensures more uniform variance over the surface. The resulting improvement is shown in Figure 1(b).

For realistic material modelling, the BSSRDF is usually a weighted sum of multiple gaussians with unique variances and weights per color channel [D'Eon and Irving 2011]. In skin for example, red light scatters much further than in the other channels. We apply importance sampling to combine all gaussian distributions per color channel by their weight via Veach's one-sample MIS model, which significantly reduces the noise of complex profiles. This improvement is shown in Figure 1(d). As our method is dominated by the cost of tracing the probe rays, the increased complexity of the weight calculation does not affect the overall speed.

Discussion We implemented this technique in the Arnold renderer where it has replaced a much more complicated implementation based around point clouds. The new algorithm is much easier to control for artists, has no memory overhead, is trivially parallelizable, and handles recursive effects without the need for multiple passes. Another benefit is that tight diffusion kernels can be used without penalty, while point cloud based methods must ensure a minimum spacing close to the mean-free path or suffer loss of detail.

References

- CHRISTENSEN, P. H., HARKER, G., SHARDE, J., SCHUBERT, B., AND BERTLI, D. 2012. Multiresolution radiosity caching for global illumination in movies. *SIGGRAPH 2012 Talks*.
- D'EON, E., AND IRVING, G. 2011. Quantized-diffusion model for rendering translucent materials. *SIGGRAPH 2011*.
- JENSEN, H. W., MORSE, S. R., LEVOY, M., AND HANRATTY, P. 2001. A practical model for subsurface light transport. *SIGGRAPH 2001*.
- WALTER, B., KHUNGURN, P., AND BLOOM, K. 2012. Bidirectional lightcuts. *SIGGRAPH 2012*.