

Deriving basic Precomputed Radiance Transfer

by Bert Peers - Last change: 20/07/2009

This text brings a number of concepts related to PRT together, trying to turn them into a straight A-to-Z derivation of the basic technique (diffuse, no bounce, no scattering). "Straight" here does not mean simple; this is not really a tutorial. Depending on how much you already know about the subject, this is a math cheat sheet, a quick refresher, or a starting point to learn what's involved.

Let's start with the rendering equation, which tells us that radiance leaving a point in some direction is what that point is emitting, plus any incoming radiance reflected in that direction:

$$L(x \rightarrow \Theta) = L_e(x \rightarrow \Theta) + \int_{\Omega} f_r(x, \Psi \leftrightarrow \Theta) L(x \leftarrow \Psi) \cos \Psi d\omega_{\Psi} \quad (1)$$

We can simplify this if we assume that all the light is coming from some external environment (ie. a sky or environment map), meaning L_e is zero everywhere. Further, a diffuse-only scene simplifies the BRDF to just a (normalized) reflectivity:

$$L_e = 0 \quad (2)$$

$$f_r = \frac{\rho_d}{\pi} \quad (3)$$

$$\Rightarrow L(x \rightarrow \Theta^*) = \frac{\rho_d}{\pi} \int_{\Omega} L(x \leftarrow \Psi) \cos \Psi d\omega_{\Psi} \quad (4)$$

The star over theta indicates that the outgoing direction doesn't matter, since the reflection is diffuse; but at the same time it's important to keep in mind that the radiance *does* have a direction, since it's radiance and not radiosity.

Simplifying to a lightsource that's infinitely far away, and looking at direct light only for now (ie. no indirect bounces or other transport), means that the incoming radiance is independent of the location x . The indirect radiance at x simply becomes the radiance arriving from the direction Ψ , provided that x can see the infinitely far away light in that direction:

$$L(x \leftarrow \Psi) = L(\Psi)V(x, \Psi) \quad (5)$$

$$\Rightarrow L(x \rightarrow \Theta^*) = \frac{\rho_d}{\pi} \int_{\Omega} L(\Psi) V(x, \Psi) \cos \Psi d\omega_{\Psi} \quad (6)$$

At this point, the incoming radiance is approximated by spherical harmonics; the function L is rewritten as a weighted sum of these basis functions y_i :

$$L(\Psi) \approx \sum_0^{n-1} w_i y_i(\Psi) \quad (7)$$

where $i = l(l+1) + m$ and $n = \text{order}^2$. Substituting and shuffling a bit gives:

$$L(x \rightarrow \Theta^*) \approx \frac{\rho_d}{\pi} \sum w_i \int_{\Omega} y_i(\Psi) V(x, \Psi) \cos \Psi d\omega_{\Psi} \quad (8)$$

It follows that if we can compute that integral offline for every point x , then we can recover an approximate radiance leaving x by calculating the dot product with those weights, and scaling.

So how do we compute the integral? Via *Monte Carlo sampling*, which says that the *Expected Value* E of a function can be approximated by averaging a bunch of samples, taken according to the probability distribution function (pdf) used in evaluating E :

$$E(f(x)) = \int f(x)p(x)dx \approx \frac{1}{N} \sum_1^N f(x), x \text{ chosen via } p \quad (9)$$

It follows that if we estimate the expected value of f/p , instead of f , we end up estimating the integral of f , rather than its expected value:

$$E\left(\frac{f(x)}{p(x)}\right) = \int f(x)dx \approx \frac{1}{N} \sum_1^N \frac{f(x)}{p(x)} \quad (10)$$

This is known as importance sampling. The integral we want in equation 8 is the product of two parts: the spherical harmonic y , and the cosine weighted visibility term. We could importance sample the harmonic, but then we would have to evaluate the visibility term for every sample, for every harmonic. Since the visibility term is the expensive part (it's where the raytracing happens), the other way around is more efficient: importance sample the cosine part, and use the sample to update the estimate for every harmonic in parallel.

Sampling according to the cosine lobe requires a pdf that ideally follows these values (unnormalized for now):

$$\int_{\Omega} \cos \Psi d\omega_{\Psi} = \int_0^{2\pi} \int_0^{\pi/2} \cos \theta \sin \theta d\theta d\varphi \quad (11)$$

$$= \int_0^{2\pi} d\varphi \int_0^{\pi/2} \cos \theta \sin \theta d\theta \quad (12)$$

The sine term appears to compensate for the fact that an equal jump along φ , will cover less distance as θ gets closer to azimuth, resulting in a smaller solid angle.

To normalize this pdf properly, and also to efficiently generate samples accordingly, we need the cumulative distribution function (cdf) P :

$$P(\varphi', \theta') = \int_0^{\varphi'} d\varphi \int_0^{\theta'} \cos \theta \sin \theta d\theta \quad (13)$$

$$= \varphi' \frac{1 - \cos^2 \theta}{2} \Big|_0^{\theta'} \quad (14)$$

$$= \varphi' \frac{1}{2} (1 - \cos^2 \theta') \quad (15)$$

We substitute the full range to figure out the normalization factor for p :

$$P(2\pi, \pi/2) = 2\pi \frac{1}{2} (1 - 0) = \pi \quad (16)$$

Dividing p by this, and in turn dividing f by p , means that we simply multiply f by π , eliminating the divide-by-pi coming out of normalizing the diffuse BRDF.

To generate samples according to p/π , invert the separable P :

$$\varphi = 2\pi\zeta_1 \quad (17)$$

$$\theta = \cos^{-1} \sqrt{1 - \zeta_2} = \cos^{-1} \sqrt{\zeta_2} \quad (18)$$

where the ζ 's are uniform numbers chosen via stratified sampling, low discrepancy sampling, etc.

Now remember that the point of choosing samples via p was so that we can divide those terms away in the monte carlo estimator:

$$\int_{\Omega} y_i(\Psi) V(x, \Psi) \cos \Psi d\omega_{\Psi} \approx \frac{\pi}{N} \sum y_i(\varphi, \theta) V(x, \varphi, \theta) \quad (19)$$

So at runtime, the radiance becomes:

$$L(x \rightarrow \Theta^*) \approx \sum_{SH} \frac{\rho_d}{\pi} w_i \sum \frac{\pi}{N} y_i(\Psi) V(x, \Psi) \text{ where } \Psi = (\varphi, \theta) \quad (20)$$

$$\approx \sum_{SH} w_i t_{i,x} \quad (21)$$

$$\text{where } t_{i,x} = \frac{\rho_d}{N} \sum y_i(\Psi) V(x, \Psi) \quad (22)$$

Now that we have the transfer coefficients t , how do we figure out the weights w for a given lighting setup? Spherical harmonic functions are orthonormal, and real valued. The former means we can process them one by one, and not worry about normalization. The latter means they are their own complex conjugate. The result is we simply integrate the light with the harmonic:

$$w_i = \int L(\Psi) y_i(\Psi) d\omega_\Psi \quad (23)$$

Some lights have been defined as a sum of spherical harmonics straight out of the box, for example sky models.

References:

- Advanced Global Illumination; Philip Dutré, Kavita Bala, Philippe Bekaert
- Principles of Digital Image Synthesis; Andrew Glassner