Vol. [VOL], No. [ISS]: 1–8

Wrap Shading

Peter-Pike Sloan
Disney Interactive Studios

Derek Nowrouzezahrai
University of Toronto, Disney Research Zürich

Hong Yuan
UMass Amherst, Disney Interactive Studios

Abstract. Shading models that wrap around the hemisphere have been used to approximate subsurface scattering, area light sources, and softer reflectance profiles. We generalize a specific technique that has been used in games by parameterizing the amount of wraping while retaining important mathematical properties. We include details on how to incorporate theses models with spherical harmonic lighting.

1. Wrap Shading with Point and Directional Lights

While standard Lambertian shading is used quite often, a more flexible technique is the “wrap” shading model, where light is allowed to wrap around the edge of the horizon instead of sharply cutting off when the cosine between the surface normal and the light is zero. An example of this in Figure 1, where the subject is partially backlit and the effect of increasing the amount of wrap becomes clear. This model is also used as a cheap approximation to subsurface scattering [Green 04] or as a more expressive base shading model [Mitchell et al. 06]. One wrap shading formulation is

$$f(\theta, w) = (cos(\theta) + w)/(1 + w) \ , \tag{1}$$

1086-7651/06 $0.50 per page

Figure 1. Left to right: no wrap shading, wrap shading with $a = 0.5$, and $a = 1.0$. Note how wrap shading softens the appearance.

where w is a parameter between $[0, 1]$ that represents how far light is allowed to wrap around the sphere, $cos(\theta)$ is the dot product of the normal and light vector, and the value is clamped to zero. When $w = 0$ this is just conventional diffuse shading. The formulation used by Valve [Mitchell et al. 06] always wraps completely around the sphere, but the fall off takes on a different shape:

$$f(\theta) = 0.25\ (cos(\theta) + 1)^2\ . \tag{2}$$

One property of the above formulation is that when $\cos(\theta) = 1$ it matches the first two derivatives of the cosine function, while the formulation in Equation 1 only matches one derivative. This property is beneficial since it ensures f more closely matches a normal diffuse surface for a larger set of angles θ, as can be seen in Figure 3. This also retains the important diffuse highlight as θ, the angle between the light and the surface normal, approaches 0.

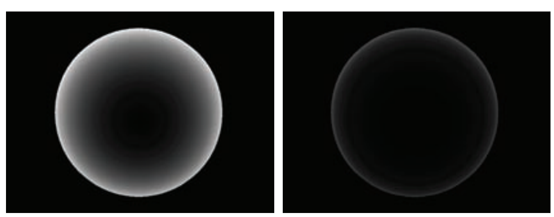

Figure 2. Wrap shading with a fairly bright light directly behind the sphere. The GPUGems approach (*left*) and the more subtle model used by Valve (*right*).

We generalize Mitchell's model [Mitchell et al. 06] while retaining its important derivative property. Our new parameterized model is

$$f(\theta, a) = \begin{cases} ((\cos\theta + a)/(1 + a))^{1+a} & \text{, if } \theta \le \theta_m \\ 0 & \text{, otherwise} \end{cases} \tag{3}$$

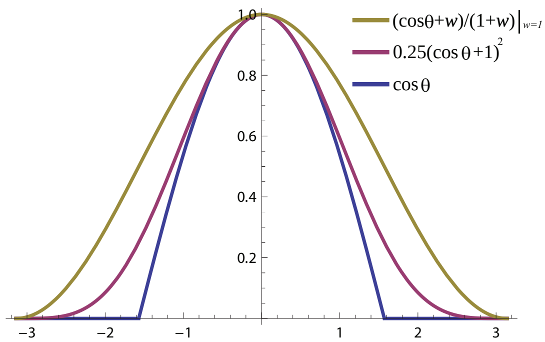

Figure 3. Comparing different shading models: the blue curve is $cos(\theta)$, purple is Equation 2, tan is Equation 1 with $w = 1$. Our model better fits the important diffuse highlights present at small values of θ, all while maintaining the ability to smoothly wrap around the bottom hemisphere.

where we clamp the domain of the function to the region where $\cos\theta + a \geq 0$, and so $\theta_m = \arccos(-a)$ is the boundary of the kernel.

Like Green [Green 04] this model is a clamped cosine at $a = 0$, reproduces Equation 2 at $a = 1$ and the derivative properties are maintained for $a \in [0, 1]$. In Figure 4 we see that when $a = 1$ the shading from the approach of [Green 04] is much flatter than our model's shading; in this case, the light is in the upper right corner.

Figure 4. Shading with Green's model (*top*) and our model (*bottom*) for different values of w and a: (*left* to *right*) w and $a = \left\{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\right\}$.

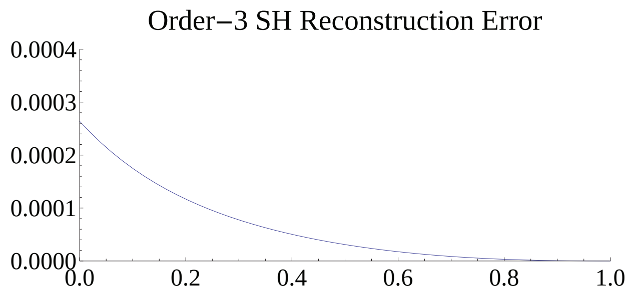

Figure 5. The L_2 reconstruction error (over the sphere) of the order-3 SH approximation of $f(\theta, a)$ is plotted as a function of a. At $a = 1$ there is no error.

2. Wrap Shading with Environmental Lights

While the models presented in Section 1 are suitable for shading from point and directional light sources, they can be adapted for computing shading from environmental light sources represented in the spherical harmonic (SH) basis.

We derive simple analytic representations for SH convolution coefficients [Sloan 08] of the shading model in [Green 04] as well as our proposed model. The order-6 convolution coefficients for the shading model in [Green 04] are

$$\begin{aligned}\mathbf{f} &= \sqrt{\frac{4\pi}{2l+1}} \int_0^{\theta_m} \int_0^{2\pi} \frac{1}{\pi}\; f(\theta, w)\; y_l^0(\theta, \phi)\; \sin(\theta)\, d\theta\, d\phi \\ &= \Bigg[w+1, -\frac{1}{3}(w-2)(w+1), \frac{1}{4}(w-1)^2(w+1), -\frac{1}{4}(w-1)^2 w(w+1), \\ &\quad \frac{1}{24}(w-1)^2(w+1)\left(7w^2-1\right), -\frac{1}{8}(w-1)^2 w(w+1)\left(3w^2-1\right) \Bigg] .\end{aligned}$$

The order-6 per-band scaling coefficients for our shading model, $f(\theta, a)$, are

$$\begin{aligned}\mathbf{f} &= \sqrt{\frac{4\pi}{2l+1}} \int_0^{\theta_m} \int_0^{2\pi} \frac{1}{\pi}\; f(\theta, a)\; y_l^0(\theta, \phi)\; \sin(\theta)\, d\theta\, d\phi \\ &= \Bigg[\frac{2(a+1)}{a+2}, \frac{4(a+1)}{(a+2)(a+3)}, \frac{2(a+1)\left(a^2-2a+3\right)}{(a+2)(a+3)(a+4)}, \\ &\quad -\frac{2(a-1)a(a+1)(5a-7)}{(a+2)(a+3)(a+4)(a+5)}, \frac{2(a-1)(a+1)\left(36a^3-46a^2-7a+15\right)}{(a+2)(a+3)(a+4)(a+5)(a+6)}, \\ &\quad -\frac{2(a-1)a(a+1)\left(329a^3-401a^2-151a+219\right)}{(a+2)(a+3)(a+4)(a+5)(a+6)(a+7)} \Bigg] .\end{aligned}$$

While this model works well for character lighting, one limitation is that it is not normalized; namely, energy is gained as you wrap more around the sphere.

The model in [Green 04] is also not normalized. The necessary normalization for our model is

$$\left(\int_0^{\theta_m}\int_0^{2\pi}\frac{1}{\pi}\,f(\theta,a)\,\sin(\theta)\,d\theta\,d\phi\right)^{-1}=\left(2\int_0^{\theta_m}f(\theta,a)\,\sin(\theta)\,d\theta\right)^{-1}=\frac{(2+a)}{2(1+a)}\;.$$

This normalization ensures that, as the model reflects light from a larger region about the shade point, it does so in a manner that *distributes* energy without *adding* any. We recommend using the normalization for physically-based rendering or environmental shading with SH, however in the case of point or directional lighting the normalization causes the diffuse highlight to no longer match; this property is only retained if the normalization for a cosine lobe is applied.

One very interesting property of our model is that, at $a = 1$, it can be perfectly represented using an **order-3** SH reconstruction, and so the higher-order terms need not be computed in this case. Furthermore, we found that an order-3 SH reconstruction provided a good trade-off between performance and accuracy; Figure 5 plots the order-3 SH reconstruction error as a function of a. As seen in Figure 5, the error is largest at $a = 0$ so, as noted in [Ramamoorthi and Hanrahan 01], an order-3 SH approximation is a reasonable choice.

If a higher-order shading model is necessary[1], it is interesting to note that, as $a \to 1$, the energy of the kernel shifts from the higher to lower-order bands, as illustrated in Figure 6 for an order-6 kernel. As noted earlier, at $a = 1$, **all** of the energy is contained within the first 3 bands.

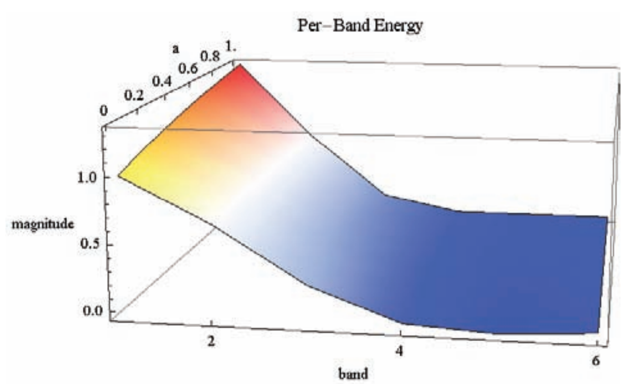

Figure 6. Visualizing how $f(\theta, a)$'s energy rapidly shifts from higher to lower-order bands with increasing a.

We include source code for the point/directional model, as well as the SH per-band scaling coefficients in Appendix A.

[1] For example, for integration into a triple product integration framework, where lighting and/or visibility may be represented with higher-order SH expansions.

3. Discussion

Wrap shading integrates seamlessly with techniques like normal mapping or glossy shading. If shadow mapping is applied, simple lighting can induce a high contrast edge (compared to conventional Lambertian shading) at shadow boundaries where the shading is abruptly clamped. This artifact can also occur to some extent whenever vertex normals are used. We expect that with environmental lighting, or larger number of light sources, these issues become unnoticeable.

4. Conclusion

We introduce a generalized wrap shading model that not only maintains all of the useful mathematical properties of existing models [Green 04, Mitchell et al. 06], but also introduces several additional important properties. Our model retains the derivative properties of Valve’s model, is parameterized like Green’s model, and reduces to standard Lambertian shading when necessary.

We extend the model from simple point and directional light sources to realistic environmental light sources represented in the SH basis, as well as deriving an equivalent extension for the parametric model in [Green 04].

Acknowledgments

We thank Morgan McGuire for identifying and discussing integration with shadow maps, and we thank Stephen Hill for spotting an error in our text and bringing it to our attention.

References

[Green 04] Simon Green. *GPU Gems*, Chapter Real-Time Approximations to Subsurface Scattering. Addison Wesley, 2004.

[Mitchell et al. 06] Jason Mitchell, Gary McTaggart, and Chris Green. “Shading in Valve’s Source Engine.” In *Advanced Real-Time Rendering in 3D Graphics and Games, SIGGRAPH ’06*, 2006.

[Ramamoorthi and Hanrahan 01] Ravi Ramamoorthi and Pat Hanrahan. “An efficient representation for irradiance environment maps.” In *SIGGRAPH 2001 Conference Proceedings, August 12–17, 2001, Los Angeles, CA*, edited by ACM, pp. 497–500. ACM Press, 2001.

[Sloan 08] Peter-Pike Sloan. “Stupid Spherical Harmonics (SH) Tricks.” In *GDC ’08*, 2008.

A. HLSL Code for Green’s Model and Our Generalization

Code listing 1 assumes that the w and a parameters of Green’s model, $f(\theta, w)$, and our model, $f(\theta, a)$, are passed to the shader via the CPU. Furthermore, we include the normalization term for our model only in the SH case, used for environmental lighting.

Listing 1. HLSL Code Listing.

```
float fW, fA; // Wrap shading parameters passed via the CPU.

// Green's model for simple point and directional light sources.
float GreenWrap_Simple(float fCosTheta) {
   return max(0, fCosTheta + fW)/(1.0f + fW);
}

// Our model for simple point and directional light sources.
float GeneralWrap(float fCosTheta) {
   return pow( max(0, fCosTheta + fA)
               /(1.0f + fA), 1.0f + fA);
}

// Function to generate order-3 SH per-band coefficients for Green's model.
float3 GreenWrapSH(void) {
    float3 t = float3( fW+1, fW-1, fW-2 );
    return float3( t.x, -t.z*t.x/3.f, 0.25f*t.y*t.y*t.x );
}

// Function to generate order-3 SH per-band coefficients for our model.
float3 GeneralWrapSH(void) {
    // Normalization factor for our model.
    float norm = 0.5*(2 + fA)/(1 + fA);
    float4 t = float4( 2*(fA+1), fA+2, fA+3, fA+4 );
    return norm*float3( t.x/t.y, 2*t.x/(t.y*t.z),
                        t.x*(fA*fA-t.x+5)/(t.y*t.z*t.w));
}
```

Web Information:

Peter-Pike Sloan, Disney Interactive Studios
(Peter-Pike.Sloan@disney.com)

Derek Nowrouzezahrai, University of Toronto, Disney Research Zürich
(derek@disneyresearch.com)

Hong Yuan, UMass Amherst, Disney Interactive Studios
(Hong.Yuan@disney.com)

Received [DATE]; accepted [DATE].