Improved Radiance Gradient Computation

Jaroslav Křivánek*
IRISA/INRIA Rennes
Czech Technical University
Univ. of Central Florida

Pascal Gautron†
IRISA/INRIA Rennes
Univ. of Central Florida

Kadi Bouatouch‡
IRISA/INRIA Rennes

Sumanta Pattanaik§
Univ. of Central Florida

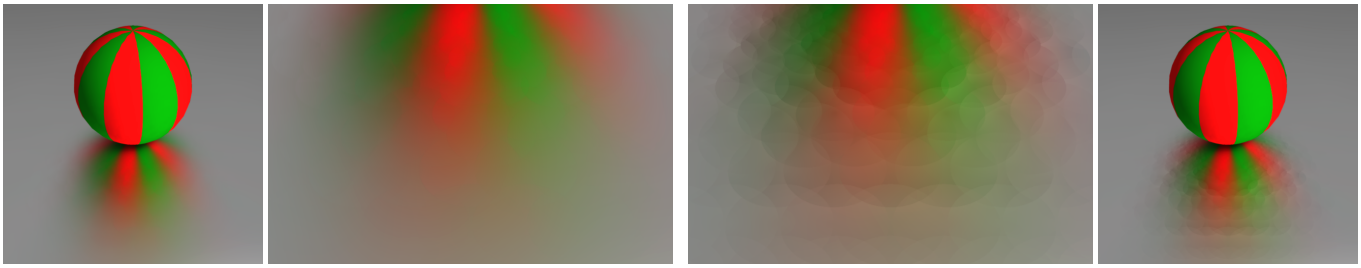

New gradients | Gradients by [Křivánek et al. 2005]

Figure 1: **Right:** The gradient computation proposed by [Křivánek et al. 2005] does not properly handle significant change of occlusion in the sampled environment and leaves visible interpolation artifacts. **Left:** The radiance gradient computation proposed in this paper handles occlusion changes and leads to a smoother indirect illumination interpolation on the glossy floor. The two images in the middle are cut out from the two images on the very left and very right.

Abstract

We describe a new and accurate algorithm for computing translational gradients of incoming radiance in the context of a ray tracing-based global illumination solution. The gradient characterizes how the incoming directional radiance function changes with displacement on a surface. We use the gradient for a smoother radiance interpolation over glossy surfaces in the framework of the radiance caching algorithm. The proposed algorithm generalizes the irradiance gradient computation by [Ward and Heckbert 1992] to allow its use for non-diffuse, glossy, surfaces. Compared to previous method for radiance gradient computation, the new algorithm yields better gradient estimates in the presence of significant occlusion changes in the sampled environment, allowing a smoother indirect illumination interpolation.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Shading, Shadowing

Keywords: radiance gradients, radiance caching, irradiance gradients, global illumination, ray tracing

*jkrivane@irisa.fr
†pgautron@irisa.fr
‡kadi@irisa.fr
§sumant@cs.ucf.edu

1 Introduction and Previous Work

Irradiance caching [Ward et al. 1988; Křivánek et al. 2005] is a ray tracing-based method for computing indirect diffuse illumination. It exploits the indirect illumination coherence by sparsely sampling, caching and interpolating the indirect diffuse illumination over surfaces, instead of computing it independently for each pixel. Indirect illumination at a caching point is computed by sampling the hemisphere of incoming directions through a number of secondary rays, and then stored as a new record in the cache. Later, when a ray hits a surface sufficiently close to the stored records, indirect illumination can be simply interpolated instead of computing it again by a costly hemisphere sampling.

Ward and Heckbert [1992] found that the interpolation quality on diffuse surfaces can be significantly improved by the use of translational and rotational irradiance gradients. The translational gradient characterizes the change of irradiance with a small displacement on a surface and the rotational gradient describes the change with surface rotation. The gradients are computed simultaneously with the hemisphere sampling and stored in the cache to be used later, during the interpolation, to effectively raise the interpolation order.

In our earlier work on *radiance caching* [Křivánek et al. 2005], we used the workings of irradiance caching to efficiently compute indirect illumination on glossy surfaces. We achieved it by caching the full directional incoming radiance function (represented by hemispherical or spherical harmonics [Gautron et al. 2004]) instead of just a scalar irradiance value. Similarly to Ward and Heckbert [1992], we also used translational gradients to improve the interpolation quality. However, some important differences between the ways hemisphere is sampled in irradiance and radiance caching prevented us from directly adapting Ward's and Heckbert's gradient computation algorithm. Instead, we proposed a more general translational gradient computation method that was applicable to the hemisphere sampling schemes used in both radiance and irradiance caching. This previous method works well for most situations, but it breaks down when there is a significant occlusion change with translation (Figure 1).

In this paper we propose a novel algorithm for translational gradient computation which works well in the presence of occlusion changes and is general enough to be applicable to radiance caching. The algorithm is based on the gradient computation proposed in [Ward and Heckbert 1992], but we reformulate the problem without assuming uniform projected area hemisphere sampling used in irradiance caching. Our formulation allows an arbitrary weighting function to be used for incoming radiance samples, allowing projection of the incoming radiance function onto an arbitrary hemispherical basis.

Apart from the papers already mentioned, the following work on translational illumination gradient computation exists. Arvo [1994] computes the irradiance Jacobian due to partially occluded polygonal emitters of constant radiosity, whereas Holzschuch and Sillion [1995] handle polygonal emitters with arbitrary radiosity. The gradient computation presented by Annen et al. [2004] is equivalent to our previous gradient computation method described above.

The rest of the paper is organized as follows. Section 2 gives a short review of radiance caching, Section 3 presents the main contribution — a new gradient computation method. Results are showcased in Section 4 and Section 5 concludes the work.

2 Radiance Caching Overview

Radiance caching is based on sparse sampling, caching and interpolating incoming radiance function over visible glossy surfaces. Whenever a ray hits a surface, the cache is queried for nearby incoming radiance records and, if some are found, the indirect illumination is evaluated by a gradient-based interpolation as described in [Křivánek et al. 2005].

In this paper, we deal with the situation when a ray hits a surface and the interpolation is not possible due to the lack of nearby records. In this case, hemispherical incoming radiance function and its translational gradient are computed by Monte Carlo quadrature and stored as a new record in the radiance cache. For efficient representation, the incoming radiance function L^i is projected onto the basis of spherical or hemispherical harmonics and represented by a vector of projection coefficients $\Lambda = \{\lambda_l^m\}$ as $L^i(\theta,\phi) \approx \sum_{l=0}^{n-1}\sum_{m=-l}^{l}\lambda_l^m H_l^m(\theta,\phi)$. In this formula, H_l^m are the basis functions and n is the representation order. The coefficients λ_l^m are computed by a stratified Monte Carlo quadrature with uniform hemisphere sampling:

$$\lambda_l^m = \frac{2\pi}{N\cdot M}\sum_{j=0}^{M-1}\sum_{k=0}^{N-1} L_{j,k}^i H_l^m(\theta_{j,k},\phi_{j,k}), \qquad (1)$$

where

$L_{j,k}^i$ is the incoming radiance from the sampled direction $(\theta_{j,k},\phi_{j,k}) = \left(\arccos(1-\frac{j+\zeta_j}{M}), 2\pi\frac{k+\zeta_k}{N}\right)$. It is computed by tracing a secondary ray from the hemisphere center in the direction $(\theta_{j,k},\phi_{j,k})$.

ζ_j, ζ_k are uniformly distributed random variables in $[0,1)$,

$N\cdot M$ is the total number of sampled directions and $N \approx 4M$.

See [Křivánek et al. 2005] for further details on radiance caching.

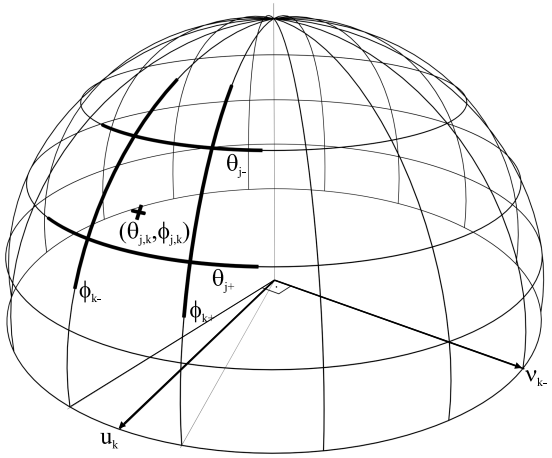

Figure 2: Uniform hemisphere subdivision for radiance sampling. The hemisphere is subdivided into rectangular cells of the same area. One random direction is selected within each cell for tracing secondary rays.

3 Novel Radiance Gradient Computation

This section presents the main contribution of this paper — a novel radiance gradient computation algorithm. The gradient computation problem consists in estimating the translational gradient $\vec{\nabla}\lambda_l^m$ for each coefficient λ_l^m from the hemisphere samples distributed according to Equation (1).

We keep the general idea of the irradiance gradient computation method of Ward and Heckbert [1992], which consists in dividing the hemisphere into cells and estimating each cell's contribution to the gradient separately. A cell's gradient is given by the marginal change of incoming radiance with the differential translation of the hemisphere center.

However, we cannot use the irradiance gradient calculation directly, for two reasons. First, we distribute the samples uniformly over the hemisphere, whereas irradiance gradients are based on uniform *projected* solid angle sampling. That is to say, the hemisphere is sampled more densely near the pole for irradiance gradients than for radiance gradients. Second, to compute the projection coefficients, we weight the incoming radiance samples by the basis functions H_l^m evaluated in the appropriate direction. No weighting is applied in the hemisphere sampling scheme used for irradiance gradients.

In the rest of this section we describe our approach to gradient computation, which relaxes the assumptions about hemisphere sampling made by Ward and Heckbert [1992] in the irradiance gradients algorithm. The stratified sampling employed in radiance caching divides the hemisphere into cells of equal solid angle, or equal area on the unit hemisphere, given by

$$A_{j,k} = (\cos\theta_{j_-} - \cos\theta_{j_+})(\phi_{k_+} - \phi_{k_-}) = \frac{1}{M}\frac{2\pi}{N}. \qquad (2)$$

The symbols relating to the geometry of the hemisphere division are illustrated in Figure 2 and defined summarized here:

(j,k) is the cell index.

$A_{j,k}$ is the cell area.

θ_{j_-} is the polar angle at the boundary between the current cell (j,k) and the previous cell $(j-1,k)$, $\theta_{j_-} = \arccos(1-\frac{j}{M})$.

θ_{j_+} is the polar angle at the boundary between the current cell (j,k) and the next cell $(j+1,k)$, $\theta_{j_+} = \arccos(1-\frac{j+1}{M})$.

ϕ_{k_-} is the azimuthal angle at the boundary between the current cell (j,k) and the previous cell $(j,k-1)$, $\phi_{k_-} = 2\pi\frac{k}{N}$.

ϕ_k is the azimuthal angle at the center of the current cell (j,k), $\phi_k = 2\pi\frac{k+0.5}{N}$.

ϕ_{k_+} is the azimuthal angle at the boundary between the current cell (j,k) and the next cell $(j,k+1)$, $\phi_{k_+} = 2\pi\frac{k+1}{N}$.

$\hat{u}_k$ is the unit vector in direction $(\pi/2,\ \phi_k)$.

$\hat{v}_{k_-}$ is the unit vector in direction $(\pi/2,\ \phi_{k_-}+\pi/2)$.

For each cell (j,k), we observe how its area changes with respect to the hemisphere center displacement along two near perpendicular vectors $\hat{u}_k$ and $\hat{v}_{k_-}$ defined above. The displacement along $\hat{u}_k$ causes a shift of the wall separating the considered cell (j,k) and its neighboring cell $(j-1,k)$. The induced change of the cell area is

$$\begin{aligned}\nabla_{\hat{u}_k} A_{j,k} &= \nabla_{\hat{u}_k}\theta_{j_-}\cdot\frac{\partial A_{j,k}}{\partial\theta_{j_-}}\\ &= \frac{-\cos\theta_{j_-}}{\min\{r_{j,k},r_{j-1,k}\}}\cdot\sin\theta_{j_-}(\phi_{k_+}-\phi_{k_-})\\ &= \frac{2\pi}{N}\frac{\cos\theta_{j_-}\sin\theta_{j_-}}{\min\{r_{j,k},r_{j-1,k}\}},\end{aligned}$$

where $\nabla_{\hat{u}_k}$ denotes the (scalar) directional derivative in the direction of $\hat{u}_k$, and $r_{j,k}$ is the distance from the hemisphere center to the closest surface in the sampled direction $(\theta_{j,k},\phi_{j,k})$. The derivative $\partial A_{j,k}/\partial\theta_{j_-} = \sin\theta_{j_-}(\phi_{k_+}-\phi_{k_-})$ follows directly from Equation (2).

Similarly, the displacement along $\hat{v}_{k_-}$ causes a shift of the wall separating the considered cell (j,k) and its neighboring cell $(j,k-1)$, and the induced change of the cell area is

$$\begin{aligned}\nabla_{\hat{v}_{k_-}} A_{j,k} &= \nabla_{\hat{v}_{k_-}}\phi_{k_-}\cdot\frac{\partial A_{j,k}}{\partial\phi_{k_-}}\\ &= \frac{-1}{\sin\theta_{j,k}\min\{r_{j,k},r_{j,k-1}\}}\cdot(\cos\theta_{j_+}-\cos\theta_{j_-}) \quad (3)\\ &= \frac{1}{M\sin\theta_{j,k}\min\{r_{j,k},r_{j,k-1}\}}.\end{aligned}$$

Here we used the equality $-(\cos\theta_{j_+}-\cos\theta_{j_-}) = 1/M$, which holds for uniform hemisphere sampling as follows from the way directions are generated in Equation (1).

The change of incoming radiance arriving at the hemisphere center through the cell is given by interpolating the radiance from two neighboring cells with the area used as the blending factor and is given by

$$\begin{aligned}\nabla_{\hat{u}_k} L^i_{j,k} &= \nabla_{\hat{u}_k} A_{j,k}(L^i_{j,k}-L^i_{j-1,k})\\ \nabla_{\hat{v}_{k_-}} L^i_{j,k} &= \nabla_{\hat{v}_{k_-}} A_{j,k}(L^i_{j,k}-L^i_{j,k-1}).\end{aligned}$$

The final gradient for a coefficient λ^m_l is given by summing the marginal radiance changes over all hemisphere cells weighted by the basis functions H^m_l:

$$\vec{\nabla}\lambda^m_l = \sum_{k=0}^{N-1}\left[\hat{u}_k\frac{2\pi}{N}\sum_{j=1}^{M-1}\frac{\cos\theta_{j_-}\sin\theta_{j_-}}{\min\{r_{j,k},r_{j-1,k}\}}(L^i_{j,k}-L^i_{j-1,k})H^m_l(\theta_{j,k},\phi_{j,k})+ \hat{v}_{k_-}\frac{1}{M}\sum_{j=0}^{M-1}\frac{1}{\sin\theta_{j,k}\min\{r_{j,k},r_{j,k-1}\}}(L^i_{j,k}-L^i_{j,k-1})H^m_l(\theta_{j,k},\phi_{j,k})\right],$$

The computed gradients lie in the tangent plane at the hemisphere center; in other words, we disregard the derivative with respect to the local Z axis. This is justified by the assumption of locally near-flat surfaces, in which case the displacements along Z are very small and therefore hardly influence the radiance change.

The change of occlusion is accounted for by differentiating the incoming radiance from the neighboring cells and using the $\min\{r_{j,k},r_{j,k-1}\}$ and $\min\{r_{j,k},r_{j-1,k}\}$ terms for estimating the relative wall movement between the current and the neighboring cells. The minimum of the two distances is important since "*it is always the distance to the* closer *surface that determines rate of change in occlusion*", as pointed out by Ward and Heckbert [1992]. On the other hand, the gradient computation by [Křivánek et al. 2005] treats each hemisphere cell completely independently from each other, based solely on the local geometrical information at the point where the cell's sampling ray hits another surface. Such an approach does not lend itself to estimating the occlusion changes. To sum up, the proposed formula is more accurate than the radiance gradient formula of [Křivánek et al. 2005] and more general than the irradiance gradient formula of [Ward and Heckbert 1992].

Unlike Ward and Heckbert, we do not compute rotational gradients. They are replaced by a full rotation of the incoming radiance function which is required anyway in radiance caching [Křivánek et al. 2005]. This rotation is applied *after* the application of translational gradients and therefore the translational gradients themselves do not have to be rotated — they are applied in the local coordinate frame of the radiance record before the rotation takes place.

Irradiance Gradient. The gradient formula above was derived for uniform hemisphere sampling with radiance samples weighted by the basis functions H^m_l. However, the same approach can be used to infer an *irradiance* gradient for the hemisphere sampling employed by Ward and Heckbert. Two slight changes have to be made. First, the weighting by $H^m_l(\theta_{j,k},\phi_{j,k})$ is replaced by $\cos\theta_{j,k}$. Second, as a consequence of the uniform *projected* solid angle sampling, the definition of $\theta_{j,k}$, θ_{j_-} and θ_{j_+} is different from ours and the equality $-(\cos\theta_{j_+}-\cos\theta_{j_-}) = 1/M$ does not hold anymore. Instead, the values $-(\cos\theta_{j_+}-\cos\theta_{j_-})$ have to be used directly in the final irradiance gradient formula:

$$\vec{\nabla}E = \sum_{k=0}^{N-1}\left[\hat{u}_k\frac{2\pi}{N}\sum_{j=1}^{M-1}\frac{\cos\theta_{j,k}\cos\theta_{j_-}\sin\theta_{j_-}}{\min\{r_{j,k},r_{j-1,k}\}}(L^i_{j,k}-L^i_{j-1,k})+ \hat{v}_{k_-}\sum_{j=0}^{M-1}\frac{\cos\theta_{j,k}(\cos\theta_{j_-}-\cos\theta_{j_+})}{\sin\theta_{j,k}\min\{r_{j,k},r_{j,k-1}\}}(L^i_{j,k}-L^i_{j,k-1})\right]$$

The θ related quantities are defined as $\theta_{j,k} = \sin^{-1}\sqrt{\frac{j+\zeta_j}{M}}$, $\theta_{j_-} = \sin^{-1}\sqrt{\frac{j}{M}}$ and $\theta_{j_+} = \sin^{-1}\sqrt{\frac{j+1}{M}}$ [Ward and Heckbert 1992]. This formula yields numerically very similar results to that of Ward and Heckbert and the images generated with the two formulas are indistinguishable from each other. The computational performance of both methods is equal since all terms except from the hit distances $r_{j,k}$ and the sampled incoming radiances $L_{j,k}$ can be precomputed.

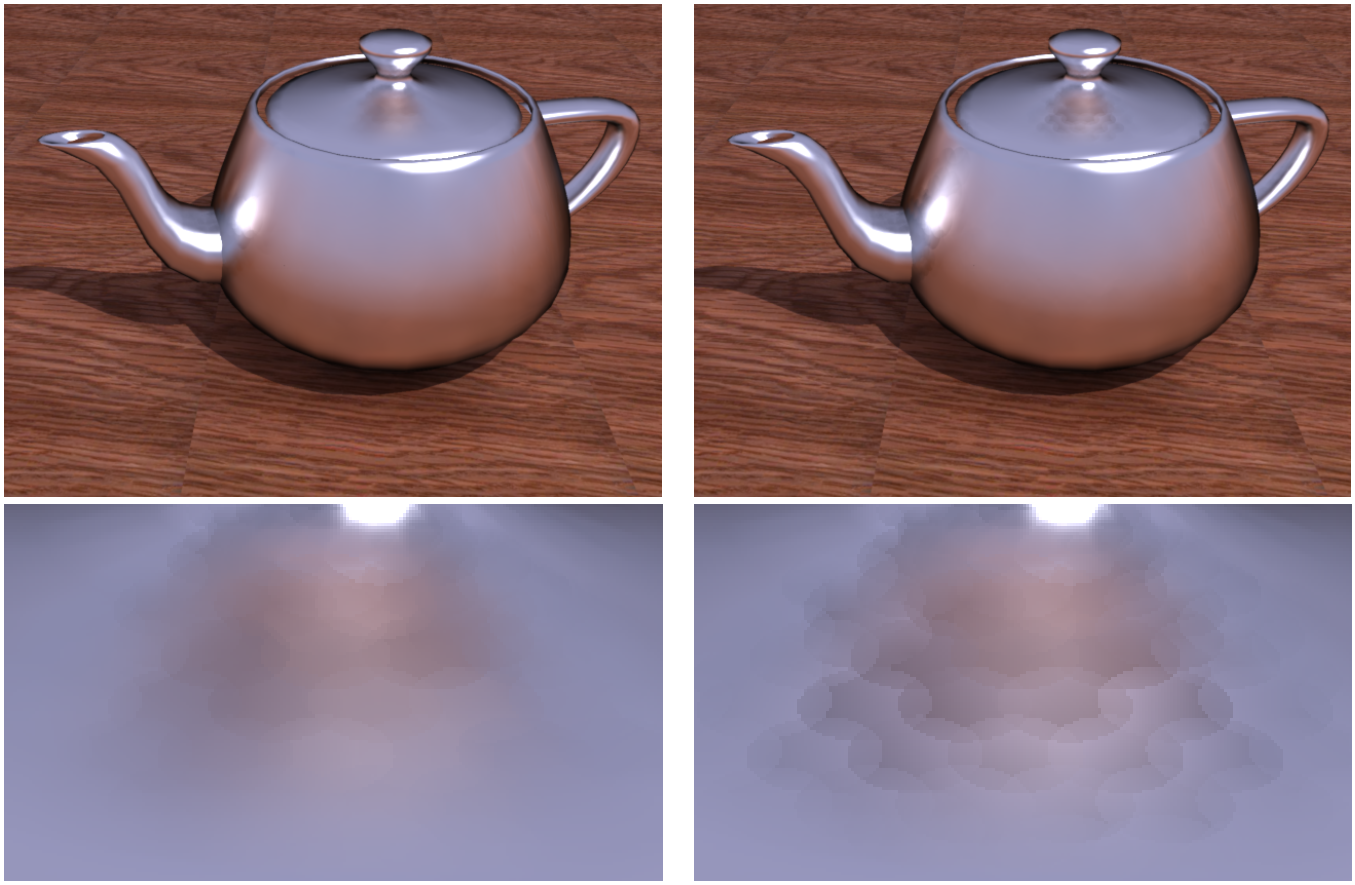

New gradients | Gradients by [Křivánek et al. 2005]

Figure 3: Indirect illumination on the teapot computed with radiance caching using gradient-based interpolation. The new gradients (on the left) provide smoother interpolation on the lid than the gradients of [Křivánek et al. 2005] (on the right).

4 Results

The scene in Figure 1 is a pathological case for the gradient computation of [Křivánek et al. 2005], but is handled correctly by the new gradient computation we have just presented. The scene consists of a glossy floor plane (isotropic Ward BRDF, $\rho_s = 0.9$, $\alpha = 1.5$, [Ward 1992]) with a diffuse sphere on it, which is the only source of indirect illumination for the glossy floor. Consequently, there is a large change of occlusion for the part of the floor that reflect the edge of the sphere. Since the gradient computation of [Křivánek et al. 2005] assumes continuity of the reflected environment, it does not handle such a situation correctly and the interpolation shows discontinuities (Figure 1 on the right). On the other hand, the new gradient computation does handle the occlusion change correctly and the sphere reflection on the floor is much smoother (Figure 1 on the right). Note that the interior of the sphere reflection is rendered near identically by both methods, since no occlusion changes prevent the method of [Křivánek et al. 2005] to perform well in that area. Both images took approximately same time to render.

Figures 3 and 4 illustrate again the same general observations. The method of [Křivánek et al. 2005] breaks down when a severe occlusion change is reflected on a glossy surface. In Figure 3 it is the part of the lid that reflects the lid handle (shown in the blow-up), in Figure 4 it is the floor reflecting the edges of the box and the pyramid (also shown in the blow-up). On the other hand, both methods perform similarly when no significant occlusion change is present.

The teapot in Figure 3 is assigned the isotropic Ward BRDF [Ward 1992] with $\rho_s = 0.9$, $\alpha = 2$ and the Cornell box floor in Figure 4 uses measured metallic BRDF fit with three Lafortune BRDF lobes [Westin 2000; Lafortune et al. 1997].

5 Conclusion

We presented a new algorithm for computing translational gradient of incoming radiance function projected onto an arbitrary hemispherical basis. On the test rendering we illustrated that the radiance interpolation based on the new gradient estimate leads to smoother results without visible discontinuity artifacts. Compared to the previous method [Křivánek et al. 2005], the new gradient computation does not involve any additional computational cost and is even easier to implement, since derivatives of the basis functions do not have to be evaluated.

In future work, we would like to devise a gradient formula for localizing bases such as spherical wavelets [Schröder and Sweldens 1995], for which the gradient computation proposed here is not directly applicable because of their hierarchical, adaptive nature. This should allow us to use smooth gradient-based interpolation on surfaces with higher frequency BRDFs.

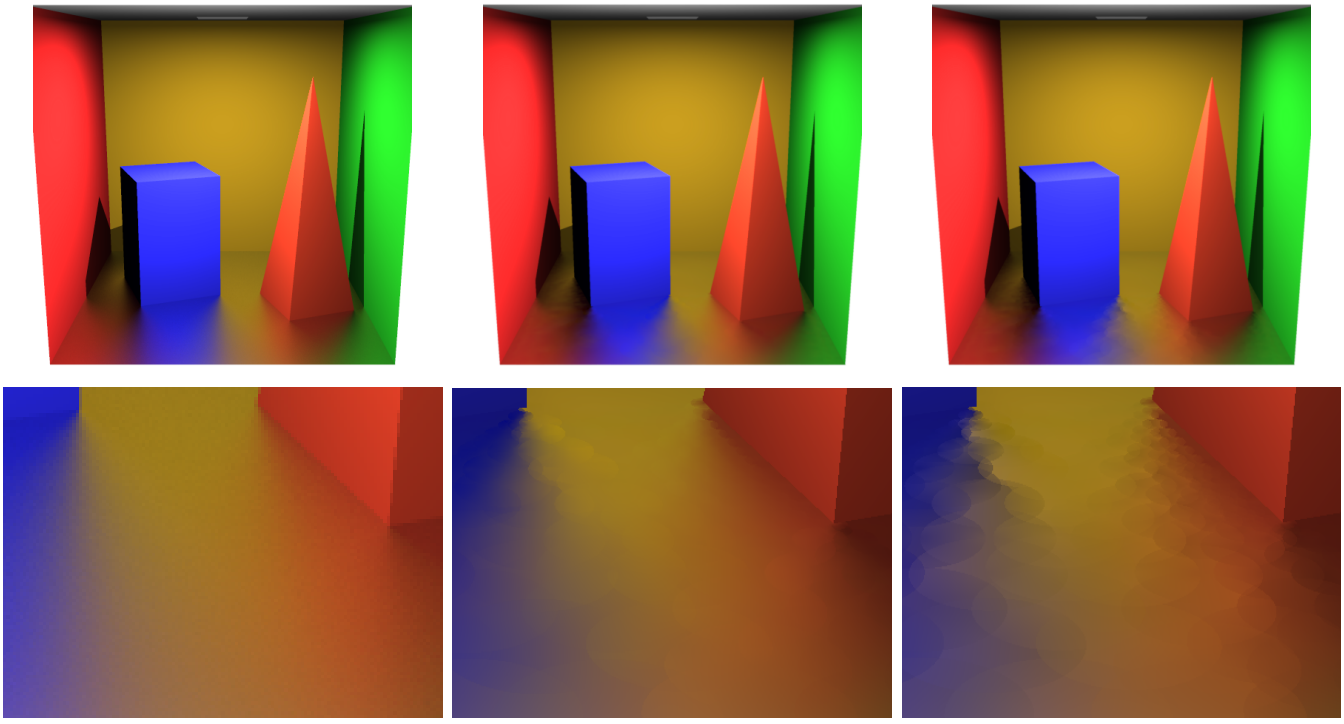

Reference solution | New gradients | Gradients by [Křivánek et al. 2005]

Figure 4: Indirect illumination on the glossy floor computed with path tracing (left) and with radiance caching using gradient-based interpolation (middle and right). Compared to the gradients of [Křivánek et al. 2005] (on the right), the new gradients (in the middle) provide smoother interpolation near the occlusion changes caused by the box and the pyramid, and provide an image visually much closer to the reference solution.

Acknowledgements

Thanks to Vlastimil Havran for letting us use his ray tracer. This work was partially supported by ATI Research and Office of Naval Research.

References

ANNEN, T., KAUTZ, J., DURAND, F., AND SEIDEL, H.-P. 2004. Spherical harmonic gradients for mid-range illumination. In *Proceedings of the Eurographics Symposium on Rendering 2004*.

ARVO, J. 1994. The irradiance jacobian for partially occluded polyhedral sources. In *Proceedings of SIGGRAPH '94*.

GAUTRON, P., KŘIVÁNEK, J., PATTANAIK, S. N., AND BOUATOUCH, K. 2004. A novel hemispherical basis for accurate and efficient rendering. In *Eurographics Symposium on Rendering*.

HOLZSCHUCH, N., AND SILLION, F. 1995. Accurate computation of the radiosity gradient with constant and linear emitters. In *Sixth Eurographics Workshop on Rendering*.

KŘIVÁNEK, J., GAUTRON, P., PATTANAIK, S., AND BOUATOUCH, K. 2005. Radiance caching for efficient global illumination computation. *Transactions on Visualization and Computer Graphics (accepted for publication)*. Also available as Technical Report #1623, IRISA, http://graphics.cs.ucf.edu/RCache/index.php.

LAFORTUNE, E. P. F., FOO, S.-C., TORRANCE, K. E., AND GREENBERG, D. P. 1997. Non-linear approximation of reflectance functions. In *Proceedings of SIGGRAPH '97*.

SCHRÖDER, P., AND SWELDENS, W. 1995. Spherical wavelets: efficiently representing functions on the sphere. In *Proceedings of SIGGRAPH*, ACM Press, 161–172.

WARD, G. J., AND HECKBERT, P. S. 1992. Irradiance gradients. In *Eurographics Workshop on Rendering*.

WARD, G. J., RUBINSTEIN, F. M., AND CLEAR, R. D. 1988. A ray tracing solution for diffuse interreflection. In *Proceedings of SIGGRAPH '88*, 85–92.

WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *Proceedings of SIGGRAPH '92*, ACM Press, 265–272.

WESTIN, S. H., 2000. Lafortune BRDF for RenderMan. http://www.graphics.cornell.edu/ westin/lafortune/lafortune.html.