

Real-time Shading with Filtered Importance Sampling

Jaroslav Krivánek
Czech Technical University in Prague

Mark Colbert
University of Central Florida

Abstract

We propose an analysis of numerical integration based on sampling theory, whereby the integration error caused by aliasing is suppressed by pre-filtering. We derive a pre-filter for evaluating the illumination integral yielding filtered importance sampling, a simple GPU-based rendering algorithm for image-based lighting. Furthermore, we extend the algorithm with real-time visibility computation. Free from any pre-computation, the algorithm supports fully dynamic scenes and, above all, is simple to implement.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism - Color, shading, shadowing, and texture

1. Introduction

Error of numerical integration is usually analyzed in terms of the integrand derivatives [QSS00]. Our main contribution is a novel approach based on sampling theory, whereby we view the error as a consequence of aliasing [Edw07]. We adopt pre-filtering to suppress the alias and derive the pre-filter shape for the illumination integral quadrature based on BRDF-proportional importance sampling (see Fig. 1).

Our analysis yields *filtered importance sampling*, a practical GPU-based algorithm for image-based lighting of glossy objects, described in detail in [CK07a]. BRDF importance sampling is implemented in a GPU fragment shader and pre-filtering is performed with a MIP-mapped lookup in an environment map texture. We extend the algorithm with a simple approximate approach to compute the visibility function for each surface point by sampling the visibility proportionally to the environment map luminance. We show that the combination of BRDF-proportional sampling for shading and environment map-proportional sampling for visibility yields a good compromise of rendering speed and simplicity vs. image quality.

Our main target application, interactive material and lighting design [CPK06], implies the required features our algorithm—support for a wide variety of BRDFs with spatially changing parameters and image-based lighting with shadows [FDA03]. BRDFs, lighting, and geometry can all be changed in real-time since no pre-computation is involved (the MIP-maps are generated on the fly during rendering).

Although similar algorithms have been proposed, e.g. [KVHS00, GKD07], we believe that they do not share

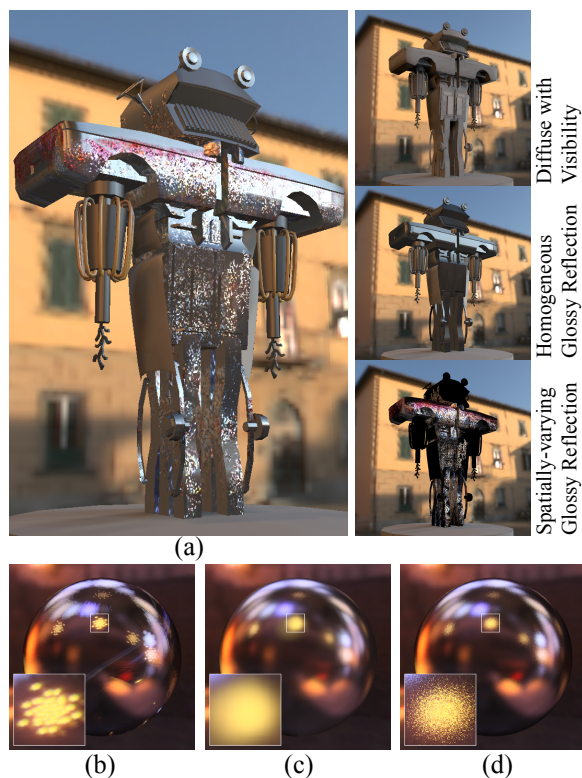


Figure 1: Using our technique, glossy as well as diffuse reflections can be rendered in real-time without any pre-computation. Here, the rusty robot (a) contains 7,722 triangles and renders at 75.8 FPS for a $1k \times 1k$ image on an NVIDIA 8800 GTX. The key idea is to suppress aliasing due to deterministic sampling (b) by pre-filtering (c) as opposed to randomization (d).

the simplicity, efficiency, and flexibility afforded by our approach. The price to pay is a limited accuracy of the generated shadows, the requirement of an efficient sampling formula for the supported BRDFs, and occasional blurring of highlights. However, in lighting or material design applications, the image quality delivered by our method is largely sufficient for an artist to understand the appearance of the scene [Col08].

2. Related Work

Aliasing and Numerical Integration. Edwards [Edw07] analyzed aliasing in numerical integration in the context of mechanical systems. Related work in computer graphics includes anti-aliasing by Gaussian quadrature [LJR96, GT96], and the pre-filtering for texture maps, notably the re-sampling framework by Heckbert [Hec89].

Real-Time Shading for Image-Based Lighting. Early techniques based on pre-filtered environment maps cannot change BRDFs and lighting in real-time because of the pre-filtering step [KVHS00]. Frequency space approaches using spherical harmonics [RH01, KSS02, RH02, SKS02] or other basis functions [NRH04, WTL04, LSSS04, TS06] partially lift this limitation but the BRDFs stay static. Global effects, such as shadows, can be included using pre-computed radiance transfer (PRT) but support for dynamic geometry remains difficult. Dynamic BRDF changes can be enabled at the expense of increased pre-computation time and storage cost [BAOR06, SZC*07]. However, our goal is to avoid pre-computation completely. McAllister et al. [MLH02] render spatially-varying Lafortune BRDFs by pre-convolving the environment map with cosine lobes of varying widths. Hensley and Scheuermann [HS06] approximate glossy reflections with a lookup into a summed-area table of the environment map. In a closely related work [GKD07], Green et al. approximate BRDFs by a sum of Gaussian kernels (fitted in an off-line process), which reduces rendering to a lookup in a pre-filtered environment map. In our algorithm, the off-line fitting is replaced by real-time BRDF importance sampling. We directly build upon the work of Colbert and Křivánek [CK07a, CK07b] from which we borrow the idea of BRDF importance sampling and environment map filtering. We extend their work by providing a theoretical analysis of the filtering that clearly identifies the approximations made, introducing support for real-time shadowing, and demonstrating a more detailed validation.

Sampling for Image-based Lighting. Methods that sample from environment lighting [ARBJ03, KK03, ODJ04, LRR05] may have problems with specular BRDFs. To address this issue, Havran et al. [HSK*05] increase the number of samples around the specular lobe; others sample from the product of the lighting and the BRDF [CJAMJ05, BGH05, TCE05, CETC06, CAM08]. Drawing a sample with these methods is involved and only pays off if a costly visibility test

is done for each. For highly specular BRDFs, rendering is more efficient using only the BRDF for sampling, either using analytic formulas [PH04] or with data-driven approaches [CPB03, LRR04]. We use BRDF importance sampling even for low-gloss BRDF and use environment map filtering to suppress aliasing.

Shadowing for Image-based Lighting. Approaches based on *visibility pre-computation* (e.g. PRT) use different ways to store the visibility information, such as various spherical basis functions [SKS02, NRH04, WTL04, LSSS04, ZHL*05, TS06, GKD07], discrete directions [MSW04], visibility cuts [AUW07] or a coherent shadow map [RGKM07]. However, the pre-computation step disables their use for scenes with deforming geometry. Techniques for real-time ambient occlusion in dynamic scenes exist [Bun05, SA07], but there is not directional information. As an alternative, Kautz et al. [KLA04] compute self-shadowing by a fast hemispherical rasterization at each vertex. Ren et al. [RWS*06] replace dynamic occluders by a sphere hierarchy and accumulates visibility in spherical harmonics. Annen et al. [ADM*08] replace an environment map by a number of area lights and uses a fast soft-shadow algorithm to generate the shadows. Compared to these approaches, we strive for simplicity, flexibility, and no pre-computation, though possibly at the cost of reduced accuracy.

3. A Sampling Theory-based Analysis of Numerical Integration

3.1. Uniform Sampling

Consider a function $g : \mathbb{R} \mapsto \mathbb{R}^+$. We want to numerically integrate $I = \int_{\alpha}^{\beta} g(u) du$. Suppose we take N regularly spaced samples and evaluate the integral as:

$$I \approx \frac{\beta - \alpha}{N} \sum_{i=1}^N g(u_i). \quad (1)$$

Since our goal is to express the integration in terms of filtering of g , we apply sampling theory treatment to the problem [OSB99]. Assuming g periodically repeats outside the integration domain, we interpret (1) as follows:

1. Sample the continuous signal $g(u)$ with the period $T = (\beta - \alpha)/N$ to get the sampled signal $g_s(u)$. In the frequency domain, this creates shifted replicas of the spectrum of g with spacing inversely proportional to T .
2. Reconstruct the continuous signal by convolving g_s with the ideal reconstruction kernel: $\hat{g}(u) = (g_s \otimes \text{sinc}_T)(u)$. In the frequency domain, this cuts off the spectrum replicas.
3. Retrieve the DC term (i.e. Fourier transform at $\omega = 0$) of the reconstructed signal. This is, indeed, an approximation of the estimated integral:

$$\hat{G}(\omega) \Big|_{\omega=0} = \int \hat{g}(u) e^{i\omega u} du \Big|_{\omega=0} = \int \hat{g}(u) du.$$

Aliasing. If the spectrum replicas due to sampling overlap, the original signal cannot be exactly reconstructed. This is aliasing. *Aliasing affects the integral estimate (1) if the spectrum replicas overlap in the DC term, i.e. at frequency $\omega = 0$.*

Pre-filtering. We adopt the customary means of suppressing aliasing by *pre-filtering* with a low-pass filter: $g'(u) = g \otimes h(u)$. Since we only need the DC term, *any normalized pre-filter that cuts off frequencies above the sampling rate ensures an alias-free integral estimate.* Such a pre-filter could also cut off frequencies below the sampling rate, blurring the signal, and, in the limit, pre-filtering could result in a constant function (equal to $I/(\beta - \alpha)$). However, we prefer pre-filters that preserve frequencies below the sampling rate, blurring the least amount possible, so that we can approximate the integral on any sub-domain.

Quadrature with filtered samples. We combine pre-filtering with sampling by taking *filtered samples*:

$$S_{US}(u_i) = \int g(u')h(u_i - u')du'.$$

Using filtered samples in (1) suppresses aliasing in the integral estimate. Filtered samples can be computed efficiently if a pre-filtered version of g , such as a MIP-map of an environment map, is available. Our analysis generalizes to multiple dimensions so we continue in 2D.

3.2. Importance Sampling

Importance sampling is an effective means for variance reduction in Monte Carlo integration. Samples are generated proportionally to a probability density function (PDF), p , so that many are taken where the integrand has large values. The *inversion method* generates a sample, x , by transforming a uniformly distributed random variable, u , by the inverse of the cumulative distribution function (CDF), $P(x) = \int_{-\infty}^x p(t)dt$, as follows: $x = P^{-1}(u)$. In 2D, the transformation is given by the inverse of the marginal and conditional CDF [Fis96]. For brevity, we use the same notation as in 1D: $\mathbf{x} = P^{-1}(\mathbf{u})$.

Although designed for Monte Carlo integration, importance sampling gives good results even for deterministic sampling, where the same set of sample directions are evaluated for each pixel. Nevertheless, the integral estimate may still suffer from aliasing. To apply our analysis (which only holds for uniform sampling), we first warp the integrand [CPL85] according to the inversion method, $\mathbf{x} = P^{-1}(\mathbf{u})$. Since the Jacobi determinant of P is equal to the PDF, p , we have:

$$I = \underbrace{\int g(\mathbf{x})d\mathbf{x}}_{\text{importance sampling}} = \underbrace{\int \frac{g(P^{-1}(\mathbf{u}))}{p(P^{-1}(\mathbf{u}))}d\mathbf{u}}_{\text{uniform sampling}}.$$

Estimating the first integral by importance sampling is equivalent to estimating the second one by uniform sampling. Having uniform sampling, we convolve with a pre-filter, yielding a filtered sample:

$$S_{IS}(\mathbf{u}_i) = \int \frac{g(P^{-1}(\mathbf{u}'))}{p(P^{-1}(\mathbf{u}'))}h(\mathbf{u}_i - \mathbf{u}')d\mathbf{u}'.$$

Warping back, we finally get pre-filtering in terms of the original, unwrapped function, $g(\mathbf{x})$:

$$S_{IS}(\mathbf{u}_i) = \int \frac{g(\mathbf{x}')}{p(\mathbf{x}')} \underbrace{[h(\mathbf{u}_i - P(\mathbf{x}'))p(\mathbf{x}')]d\mathbf{x}'}_{\substack{\text{warped pre-filter} \\ \text{(spatially variant)}}} \quad (2)$$

Quadrature based on importance sampling with filtered samples (2) yields an alias-reduced integral estimate. Unlike for uniform sampling, the pre-filter is now *spatially variant*, with size inversely proportional to the PDF—smaller filter is used in areas of high sample density.

4. Illumination Integral Evaluation

We now use the analysis from the previous section in evaluation of the *illumination integral* that gives the outgoing radiance in the direction of the camera, $\ell_o(\mathbf{o})$, as a hemispherical integral over all incoming lighting directions \mathbf{x} of the incoming radiance, $\ell_i(\mathbf{x})$, multiplied by the BRDF, $f_r(\mathbf{x}, \mathbf{o})$, and the cosine term, $\cos\theta_{\mathbf{x}}$:

$$\ell_o(\mathbf{o}) = \int \ell_i(\mathbf{x})f_r(\mathbf{x}, \mathbf{o})\cos\theta_{\mathbf{x}}d\mathbf{x}.$$

We simplify the notation by dropping \mathbf{o} , since it is fixed inside the integral, and by folding $f_r(\mathbf{x}, \mathbf{o})\cos\theta_{\mathbf{x}}$ into a single function $f(\mathbf{x})$ to get: $\ell_o = \int \ell_i(\mathbf{x})f(\mathbf{x})d\mathbf{x}$. We use *BRDF-proportional importance sampling combined with filtering of $\ell_i(\mathbf{x})$ to evaluate this integral.* The following facts motivate this choice. (1) The incoming radiance, ℓ_i , represented by an environment map, can be efficiently filtered using MIP-mapping. (2) The BRDF, given by an analytical formula, can be efficiently sampled, but not filtered.

4.1. Filtered Samples for the Illumination Integral

We use three interrelated spaces in the derivation of filtered samples for an alias-reduced estimation of the illumination integral:

- $\mathcal{M}_{\mathbf{x}}$ Unit sphere. Illumination integral is defined here.
- $\mathcal{M}_{\mathbf{u}}$ Unit square, $[0, 1]^2$. Related to the sphere, $\mathcal{M}_{\mathbf{x}}$, by the inverse CDF: $\mathbf{x} = P^{-1}(\mathbf{u})$. Therefore, importance sampling in $\mathcal{M}_{\mathbf{x}}$ corresponds to uniform sampling in $\mathcal{M}_{\mathbf{u}}$.
- $\mathcal{M}_{\mathbf{s}}$ Environment map representation, e.g. a cube map or a dual paraboloid map, related to the sphere, $\mathcal{M}_{\mathbf{x}}$, through $\mathbf{s} = R^{-1}(\mathbf{x})$. The the Jacobi determinant of R is denoted r .

First, we warp the illumination integral from $\mathcal{M}_{\mathbf{x}}$ to $\mathcal{M}_{\mathbf{u}}$. A filtered sample is then given by the convolution with the

pre-filter, h , in the domain of $\mathcal{M}_{\mathbf{u}}$. Then, we warp the convolution from $\mathcal{M}_{\mathbf{u}}$ to $\mathcal{M}_{\mathbf{s}}$ to get the filtered sample in terms of filtering the environment map with a spatially variant filter:

$$S_{\Pi}(\mathbf{u}_i) = \int \ell_i(\mathbf{x}') \frac{f(\mathbf{x}')}{p(\mathbf{x}')} [h(\mathbf{u}_i - P(\mathbf{x}')) p(\mathbf{x}') r(\mathbf{s}')] ds' \\ \approx \frac{f(P^{-1}(\mathbf{u}_i))}{p(P^{-1}(\mathbf{u}_i))} \int \ell_i(\mathbf{x}') \underbrace{[h(\mathbf{u}_i - P(\mathbf{x}')) p(\mathbf{x}') r(\mathbf{s}')] }_{\text{spatially variant pre-filter}} ds',$$

where $\mathbf{x}' = R(\mathbf{s}')$. On the last line, we moved f/p out from the integral, so that the filter only acts upon the environment map, ℓ_i . Effectively, we assume f/p is constant over the support of h . This is a fairly good approximation, since the support is small and the PDF, p , is proportional to the BRDF.

The above derivation can be extended to support shadowing by introducing the visibility function $v(\mathbf{x})$ in the illumination integral. In the derivation of the filtered sample, S_{Π} , the visibility would be taken out of the integral, assuming any occlusions are constant over the pre-filter support [GKD07].

4.2. Approximation of the Pre-filter

To make filtering efficient, we opt for a rather coarse approximation of the pre-filter—we replace the mapping $R \circ P$ locally around the i -th sample by isotropic scaling, i.e.

$$P(R(\mathbf{s})) \approx \mathbf{u}_i + (\mathbf{s} - \mathbf{s}_i) [p(\mathbf{x}_i) r(\mathbf{s}_i)],$$

where $\mathbf{x}_i = R(\mathbf{s}_i)$. Recall that p and r are the Jacobi determinants of P and R , respectively. With this approximation, shape of the pre-filter is not altered by the mapping; only its size is scaled by $1/p(\mathbf{x}_i)r(\mathbf{s}_i)$. Such a filter can be implemented by a MIP-mapped lookup in the environment map. A disadvantage of our choice is the loss of anisotropy in filtering, which may occasionally lead to blurring in the rendered images (see Fig. 4).

We experimented with replacing the $R \circ P$ mapping by its local affine approximation [Hec89] similarly to path differentials [SW01]. We use finite differences to estimate the partial derivatives and HLSL's `tex2Dgrad()` to perform the lookup. Surprisingly, the image quality is almost always worse than with isotropic filtering, except for highly anisotropic BRDFs. Moreover, there is nearly a tenfold drop in performance due to the partial derivatives computation and the anisotropic texture lookup.

5. Implementation

5.1. Filtered Importance Sampling

A detailed description of the implementation (bar the visibility, explained in the next subsection) is given in [CK07a]. However, we review it here for completeness. We use spherical harmonics for the diffuse component [RH01] of the material, and BRDF-proportional importance sampling with environment map filtering according to Section 4 for the glossy

component. Uniformly distributed numbers are computed on the CPU and uploaded into the constant registers. BRDF importance sampling code that transforms these into directions is implemented in a pixel shader, or partially pre-computed on the CPU, if using homogeneous materials. For each direction, we perform a MIP-mapped texture lookup of the environment map with the level, l , computed according to Section 4.2:

$$l = \max \left\{ \log_4 \left(K \frac{w \cdot h}{N} \frac{1}{p(\mathbf{x}_i)} \frac{1}{r(\mathbf{s}_i)} \right), 0 \right\}$$

Here, w and h are the dimensions of the environment map; N is the number of samples; and \mathbf{x}_i is the i -th sample direction. Constant K is the extent of the un-warped pre-filter h , derived empirically. For MIP-maps constructed with the box filter, we found the best image quality is obtained with $K = 4$. Finally, we attenuate the incoming radiance by the visibility evaluated for the sample direction. Pseudo-code of the shader in the Cg language is as follows:

```
float4 FIS(float3 viewing:TEXCOORD1
          uniform sampler2D env) : COLOR
{
    float4 c = 0;

    for (int i=0; i<N; i++) { // sample loop
        float2 ui = uniform_seq(i);
        float3 xi = sample_brdf(ui, viewing);
        float pdf = p(xi, viewing);
        float lod = compute_lod(xi, pdf);
        float2 si = dual_paraboloid(xi);
        float3 L = tex2Dlod(env, float3(si, lod));
        float vis = eval_visibility(xi);
        c += L * vis * f(xi, viewing) / pdf;
    }
    return c/N;
}
```

Environment Representation. We use dual paraboloid mapping [HS98] scaled by a factor of $b = 1.2$ so that the texture for each hemisphere contains more information from the opposite hemisphere, thus reducing artifacts due to independent MIP-map filtering for each hemisphere. The Jacobi determinant of this mapping is $r(z) = 4b^2(|z| + 1)^2$, where z is the z -component of the sample direction. We do not use cube maps since independent filtering of each face results in image artifacts (discontinuities).

5.2. Visibility Approximation

Representation. To support shadowing, we follow the observation of Kozłowski and Kautz [KK07] and Green et al. [GKD07] that a precise visibility representation on glossy surfaces is often unnecessary. Similar to Green et al., we represent the visibility at a point with 4th order spherical harmonics (SH) and assume constant visibility over the support of the pre-filter. The `eval_visibility()` function in the FIS shader uses SH coefficients stored in a visibility texture to compute the visibility factor for a particular direction. Representing the visibility multiplied by the cosine




Figure 2: (a) The execution time for rendering a glossy sphere at a resolution of 512×512 with (red) and without (blue) visibility. (b) Time for generating the spherical harmonics coefficients for visibility and rendering a diffuse surface with different polygon counts. Both times vary close to linearly with respect to the number of samples, where their sum represents the total rendering time for a frame in a dynamic scene. The non-linearity in (a) at 45 shadowed filtered samples results from a device-specific optimization shift on the GPU due to the increased number of instructions.

term reduces the illumination integral on diffuse materials to a scalar product and, additionally, reduces ringing artifacts.

Computation. Green et al. [GKD07] computes the visibility at each vertex in a pre-processing step with a ray tracer. To allow dynamic scenes, we instead propose to compute the visibility directly on the GPU. We perform importance sampling of the environment map using a standard approach based on marginal and conditional probabilities [PH04]. For each direction chosen, we sample the visibility function by (1) rendering the scene into an orthographic shadow map in one pass and (2) updating the SH coefficients for each texel in the visibility texture in another pass. We reconstruct the SH coefficients from non-uniform samples by a least-squares fit using the Moore-Penrose pseudoinverse on the normal equations [Bis06].

Using environment map *luminance* as the importance for sampling *visibility* may seem counterproductive and, indeed, may not be optimal for sampling the visibility function itself. However, doing so is beneficial since the visibility is captured for directions that produce the most pronounced shadows.

The angular discretization due to the visibility sampling is hidden by the smoothness of spherical harmonics. Nonetheless, spatial discretization caused by the use of shadow maps may lead to artifacts and requires further treatment. Although a rough approximation, filtering the SH coefficients spatially in texture space alleviates the problem.

6. Results and Discussion

Performance. Our results were gathered on a PC running Windows XP with an Intel 2.66GHz Dual Core2 CPU and an NVIDIA GeForce 8800 GTX GPU. We obtain real-time performance when rendering dynamic scenes with arbitrary BRDFs given by an analytic formula (we have tested Phong,

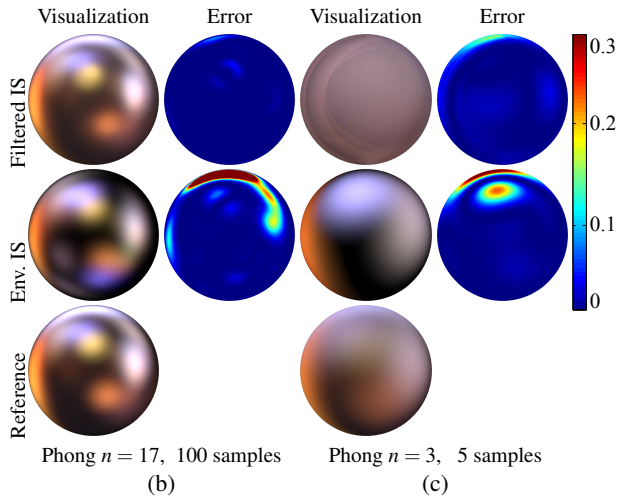



Figure 3: Comparison of our filtered BRDF importance sampling (IS) to environment IS. (a) shows the RMS error when rendering a sphere lit by the Grace light probe. While both methods converge at a rate of $N^{-1/2}$, filtered BRDF IS has a lower RMS overall, especially on more glossy materials. Visually, the BRDF IS better captures many of the subtle features of the environment (b). Although the RMS figures suggest better performance for BRDF IS, even for low-frequency glossy reflections with few samples, environment IS provides more visually pleasing images in such cases (c).

Lafortune [LFTG97], Ward [War92], and the Halfway vector disk [EBJ*06] models). The proposed filtering affords using only a small number of samples (40-50 in practice) for visually and numerically acceptable results, making the implementation amenable to the GPU. As seen in Figure 2, the run-time performance is directly proportional to the number of samples for evaluating both filtered importance samples and visibility samples (shadow maps).

Error. When comparing against the common technique of environment map proportional sampling, where the environment is represented by a set of directional lights, we obtain numerically and visually superior results for a wide variety of glossy BRDFs with the small additional cost of a tri-linear filtered texture lookup per sample (see Fig. 3). However, when using only a few samples with low-frequency BRDFs,

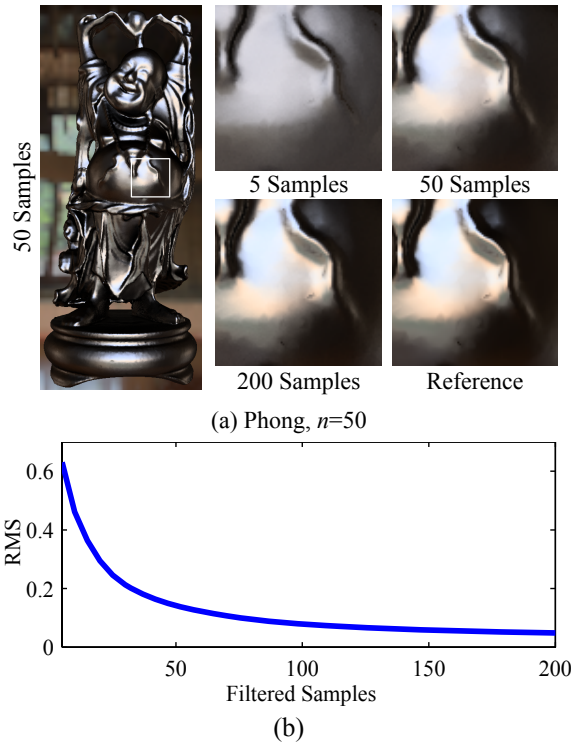


Figure 4: On complex geometry, undersampling appears as blurring and dulling of the highlights. However, visually acceptable results are found around 50 samples per pixel (a). The blurring on complex geometry affects the RMS (b) such that more samples are required to obtain an error similar to the spheres in Figure 3. However, the high-frequency contours on the geometry mask the error seen on smooth surfaces.

environment map importance sampling provides more pleasing images. Surprisingly enough, the RMS error for filtered importance sampling is still consistently lower.

The actual rendering error depends upon the complexity of the geometry. For instance, smoother meshes, such as the sphere in Figure 3, require fewer samples than more complex ones, Figure 4, to reach the same RMS error. However, the visual appearance of artifacts on the complex mesh is less noticeable due to the high frequency contours of the surface masking the errors.

Visibility. When evaluating the visibility, we use a rather rough approximation, often only 16 samples (shadow maps), but obtain fairly convincing results as demonstrated in Figure 5. A comparison with a reference spherical harmonic-based visibility function computed per pixel using 10,000 samples reveals that after 16 visibility samples, most of the error originates from the low-pass filter used to remove the spatial aliasing from the undersampled shadow maps.

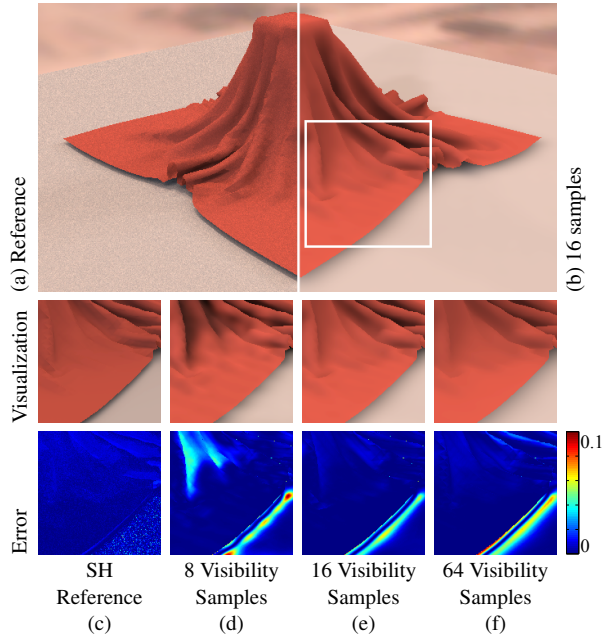


Figure 5: Comparison between a Monte Carlo estimate (a) computed with 30,000 shadow rays versus our solution (b) rendered with 16 shadow maps. To reveal the consequences of sparse visibility sampling, we compare our results to a reference SH fit computed per pixel using 10,000 stratified samples (c). This provides the best possible solution attainable by a SH-based visibility representation. The error incurred when using sparse sampling (d-f) is mostly gone by 16 samples with the error at 64 samples occurring mostly from the low-pass filter used to remove the shadow maps' spatial aliasing.

Dynamic scenes. The combination of our filtered importance sampling and visibility approximation provides a method for rendering spatially-varying materials in dynamic scenes (Figure 1a). Moreover, since we use deterministic sampling, the rendering does not suffer from any temporal flickering, noise, or aliasing. In our accompanying video, we demonstrate the ability to change the material properties in real time and view animated scenes without pre-computation.

7. Conclusion and Future Work

We have shown that analyzing numerical quadrature using sampling theory provides a new way of thinking about the problem and can be useful in designing efficient rendering algorithms. Applying our analysis to the evaluation of the illumination integral results in a simple and efficient GPU-based algorithm for real-time rendering of glossy objects under image-based illumination.

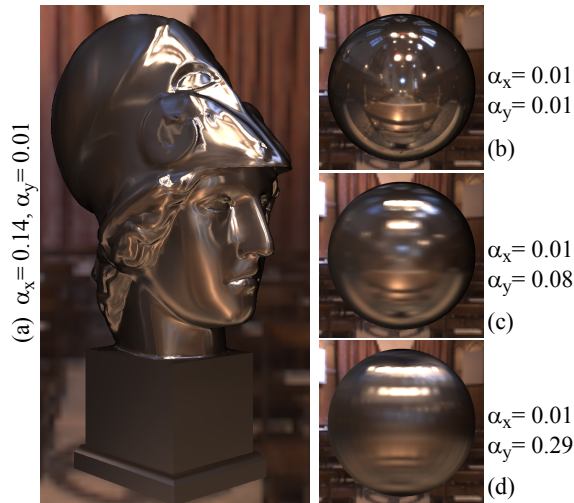


Figure 6: The Ward anisotropic BRDF rendered using filtered importance sampling. The method works well on complex surfaces (a) as well as high frequency materials (b) with slightly anisotropic reflections (c). However, the isotropic filter approximation breaks down when rendering more anisotropic materials (d).

Our current implementation can handle only moderately anisotropic BRDFs (Figure 6). Lifting this limitation is left for future work. Additionally, we would like to find a heuristic to automatically choose the number of samples based on the properties of the BRDF, environment, and geometry. Most importantly, further investigations should relate our theoretical analysis to the work of Ramamoorthi and Hanrahan [RH04] and Durand et al. [DHS*05]. We believe that our theory could eventually link the frequency-space and Monte Carlo approaches to image synthesis.

Acknowledgements

This work has been supported by the Ministry of Education, Youth and Sports of the Czech Republic under the research program LC-06008 (Center for Computer Graphics), the Aktion Kontakt grant no. 48p11 and by the US National Science Foundation grants SES0527675 and DRL0638977. Many thanks to Dan Sýkora for the discussions and Petr Olšák for checking the math.

References

- [ADM*08] ANNEN T., DONG Z., MERTENS T., BEKAERT P., SEIDEL H.-P., KAUTZ J.: Real-time, all-frequency shadows in dynamic scenes. *ACM Trans. Graph.* 27, 3 (2008).
- [ARBJ03] AGARWAL S., RAMAMOORTHI R., BELONGIE S., JENSEN H. W.: Structured importance sampling of environment maps. *ACM Trans. Graph.* 22, 3 (2003), 605–612.

- [AUW07] AKERLUND O., UNGER M., WANG R.: Precomputed visibility cuts for interactive relighting with dynamic BRDFs. In *Proc. of Pacific Graphics* (2007).
- [BAOR06] BEN-ARTZI A., OVERBECK R., RAMAMOORTHI R.: Real-time BRDF editing in complex lighting. *ACM Trans. Graph.* 25, 3 (2006).
- [BGH05] BURKE D., GHOSH A., HEIDRICH W.: Bidirectional importance sampling for direct illumination. In *Proc. of Eurographics Symposium on Rendering* (2005), pp. 147–156.
- [Bis06] BISHOP C. M.: *Pattern Recognition and Machine Learning*. Springer, 2006.
- [Bun05] BUNNELL M.: Dynamic ambient occlusion and indirect lighting. In *GPU Gems 2* (2005), Addison-Wesley Professional.
- [CAM08] CLARBERG P., AKENINE-MÖLLER T.: Practical product importance sampling for direct illumination. *Comp. Graph. Forum (Proc. of Eurographics – to appear)* 27, 2 (2008).
- [CETC06] CLINE D., EGBERT P. K., TALBOT J. F., CARDON D. L.: Two stage importance sampling for direct lighting. In *Proc. of Eurographics Symposium on Rendering* (2006).
- [CJAMJ05] CLARBERG P., JAROSZ W., AKENINE-MÖLLER T., JENSEN H. W.: Wavelet importance sampling: efficiently evaluating products of complex functions. *ACM Trans. Graph.* 24, 3 (2005), 1166–1175.
- [CK07a] COLBERT M., KRIVÁNEK J.: GPU-based importance sampling. In *GPU Gems 3* (2007), Addison-Wesley Professional.
- [CK07b] COLBERT M., KRIVÁNEK J.: Real-time shading with filtered importance sampling. In *SIGGRAPH Sketches* (2007).
- [Col08] COLBERT M.: *Appearance-Driven Material Design*. PhD thesis, University of Central Florida, 2008.
- [CPB03] CLAUSTRES L., PAULIN M., BOUCHER Y.: BRDF measurement modelling using wavelets for efficient path tracing. *Comp. Graph. Forum* 22, 4 (2003), 701–716.
- [CPK06] COLBERT M., PATTANAİK S., KRIVÁNEK J.: BRDF-Shop: Creating physically correct bidirectional reflectance distribution functions. *IEEE Comp. Graph. Appl.* 26, 1 (2006).
- [CPL85] CLARK J. J., PALEMER M. R., LAWRENCE P. D.: A transformation method for the reconstruction of functions from nonuniformly spaced samples. *IEEE Trans. on Acoustic, Speech, and Signal Processing ASSP-33*, 4 (1985).
- [DHS*05] DURAND F., HOLZSCHUCH N., SOLER C., CHAN E., SILLION F. X.: A frequency analysis of light transport. *ACM Trans. Graph.* 24, 3 (2005).
- [EBJ*06] EDWARDS D., BOULOS S., JOHNSON J., SHIRLEY P., ASHIKHMIN M., STARK M., WYMAN C.: The halfway vector disk for BRDF modeling. *ACM Trans. Graph.* 25, 1 (2006).
- [Edw07] EDWARDS T. S.: Effects of aliasing on numerical integration. *Mechanical Systems and Signal Processing* 21, 1 (January 2007), 165–176.
- [FDA03] FLEMING R. W., DROR R. O., ADELSON E. H.: Real-world illumination and the perception of surface reflectance properties. *Journal of Vision* 3 (July 2003), 347–368.
- [Fis96] FISHMAN G. S.: *Monte Carlo: Concepts, Algorithms, and Applications*. Springer-Verlag, 1996.

- [GKD07] GREEN P., KAUTZ J., DURAND F.: Efficient reflectance and visibility approximations for environment map rendering. *Comp. Graph. Forum* 26, 2 (2007), 495–502.
- [GT96] GUENTER B., TUMBLIN J.: Quadrature prefiltering for high quality antialiasing. *ACM Trans. Graph.* 15, 4 (1996).
- [Hec89] HECKBERT P. S.: *Fundamentals of Texture Mapping and Image Warping*. Master's thesis, University of California, Berkeley, 1989.
- [HS98] HEIDRICH W., SEIDEL H.-P.: View-independent environment maps. In *Proc. of Graphics Hardware* (1998).
- [HS06] HENSLEY J., SCHEUERMANN T.: Dynamic glossy environment reflections using summed area table. In *ShaderX⁴* (2006), Charles River Media.
- [HSK*05] HAVRAN V., SMYK M., KRAWCZYK G., MYSZKOWSKI K., SEIDEL H.-P.: Interactive system for dynamic scene lighting using captured video environment maps. In *Proc. of Eurographics Symposium on Rendering* (2005).
- [KK03] KOLLIG T., KELLER A.: Efficient illumination by high dynamic range images. In *Proc. of Eurographics Symposium on Rendering* (2003), pp. 45–51.
- [KK07] KOZŁOWSKI O., KAUTZ J.: Is accurate occlusion of glossy reflections necessary? In *Proc. of Symposium on Applied Perception in Graphics and Visualization* (2007).
- [KLA04] KAUTZ J., LEHTINEN J., AILA T.: Hemispherical rasterization for selfshadowing of dynamic objects. In *Proc. of Eurographics Symposium on Rendering* (2004).
- [KSS02] KAUTZ J., SLOAN P.-P., SNYDER J.: Fast, arbitrary BRDF shading for low-frequency lighting using spherical harmonics. In *Proc. of Eurographics Workshop on Rendering* (2002), pp. 301–308.
- [KVHS00] KAUTZ J., VÁZQUEZ P.-P., HEIDRICH W., SEIDEL H.-P.: A unified approach to prefiltered environment maps. In *Proc. of Eurographics Workshop on Rendering* (2000).
- [LFTG97] LAFORTUNE E. P. F., FOO S.-C., TORRANCE K. E., GREENBERG D. P.: Non-linear approximation of reflectance functions. In *Proc. of SIGGRAPH* (1997), pp. 117–126.
- [LJR96] LIU N., JIN H., ROCKWOOD A. P.: Antialiasing by gaussian integration. *IEEE Comp. Graph. Appl.* 16, 3 (1996).
- [LRR04] LAWRENCE J., RUSINKIEWICZ S., RAMAMOORTHI R.: Efficient BRDF importance sampling using a factored representation. *ACM Trans. Graph.* 23, 3 (2004), 496–505.
- [LRR05] LAWRENCE J., RUSINKIEWICZ S., RAMAMOORTHI R.: Adaptive numerical cumulative distribution functions for efficient importance sampling. In *Proc. of Eurographics Symposium on Rendering* (2005), pp. 11–20.
- [LSSS04] LIU X., SLOAN P.-P., SHUM H.-Y., SNYDER J.: All-frequency precomputed radiance transfer for glossy objects. In *Proc. of Eurographics Symposium on Rendering* (2004).
- [MLH02] MCALLISTER D. K., LASTRA A., HEIDRICH W.: Efficient rendering of spatial bi-directional reflectance distribution functions. In *Proc. of Graphics Hardware* (2002).
- [MSW04] MEI C., SHI J., WU F.: Rendering with spherical radiance transport maps. *Comp. Graph. Forum* 23, 3 (2004).
- [NRH04] NG R., RAMAMOORTHI R., HANRAHAN P.: Triple product wavelet integrals for all-frequency relighting. *ACM Trans. Graph.* 23, 3 (2004), 477–487.
- [ODJ04] OSTROMOUKHOV V., DONOHUE C., JODOIN P.-M.: Fast hierarchical importance sampling with blue noise properties. *ACM Trans. Graph.* 23, 3 (2004), 488–495.
- [OSB99] OPPENHEIM A. V., SCHAFFER R. W., BUCK J. R.: *Discrete-Time Signal Processing*, 2nd ed. Prentice Hall, 1999.
- [PH04] PHARR M., HUMPHREYS G.: *Physically Based Rendering: from Theory to Implementation*. Morgan Kaufmann, 2004.
- [QSS00] QUARTERONI A., SACCO R., SALERI F.: *Numerical Mathematics*. Springer-Verlag, New York, NY, 2000.
- [RGKM07] RITSCHEL T., GROSCH T., KAUTZ J., MÜLLER S.: Interactive illumination with coherent shadow maps. In *Proc. of Eurographics Symposium on Rendering* (2007).
- [RH01] RAMAMOORTHI R., HANRAHAN P.: An efficient representation for irradiance environment maps. In *Proc. of SIGGRAPH* (2001).
- [RH02] RAMAMOORTHI R., HANRAHAN P.: Frequency space environment map rendering. *ACM Trans. Graph.* 21, 3 (2002).
- [RH04] RAMAMOORTHI R., HANRAHAN P.: A signal-processing framework for reflection. *ACM Trans. Graph.* 23, 4 (2004), 1004–1042.
- [RWS*06] REN Z., WANG R., SNYDER J., ZHOU K., LIU X., SUN B., SLOAN P.-P., BAO H., PENG Q., GUO B.: Real-time soft shadows in dynamic scenes using spherical harmonic exponentiation. *ACM Trans. Graph.* 25, 3 (2006), 977–986.
- [SA07] SHANMUGAM P., ARIKAN O.: Hardware accelerated ambient occlusion techniques on GPUs. In *Proc. of the Symposium on Interactive 3D graphics and games* (2007).
- [SKS02] SLOAN P.-P., KAUTZ J., SNYDER J.: Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Trans. Graph.* 21, 3 (2002), 527–536.
- [SW01] SUYKENS F., WILLEMS Y. D.: Path differentials and applications. In *Proc. of Eurographics Workshop on Rendering* (2001).
- [SZC*07] SUN X., ZHOU K., CHEN Y., LIN S., SHI J., GUO B.: Interactive relighting with dynamic BRDFs. *ACM Trans. Graph.* 26, 3 (2007).
- [TCE05] TALBOT J. F., CLINE D., EGBERT P.: Importance resampling for global illumination. In *Proc. of Eurographics Symposium on Rendering* (2005), pp. 139–146.
- [TS06] TSAI Y.-T., SHIH Z.-C.: All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. *ACM Trans. Graph.* 25, 3 (2006).
- [War92] WARD G. J.: Measuring and modeling anisotropic reflection. In *Proc. of SIGGRAPH* (1992), ACM Press.
- [WTL04] WANG R., TRAN J., LUEBKE D.: All-frequency relighting of non-diffuse objects using separable BRDF approximation. In *Proc. of Eurographics Symposium on Rendering* (2004).
- [ZHL*05] ZHOU K., HU Y., LIN S., GUO B., SHUM H.-Y.: Precomputed shadow fields for dynamic scenes. *ACM Trans. Graph.* 24, 3 (2005).