

Axis-Aligned Filtering for Interactive Physically-Based Diffuse Indirect Lighting

Soham Uday Mehta¹

Brandon Wang¹

Ravi Ramamoorthi¹

Fredo Durand²

¹University of California, Berkeley ²MIT CSAIL




Figure 1: (a) We render the Sponza Atrium with 262K triangles, textures and 1-bounce physically-based global illumination at about 2 fps on an NVIDIA GTX 690 graphics card, with an average of 63 Monte Carlo (adaptive) samples per pixel (spp) raytraced on the GPU with Optix, followed by adaptive image filtering. (b) Adaptive sampling rates and filter widths (in pixels) derived from our novel frequency analysis of indirect illumination. (c) Insets of the unfiltered result. Adaptive sampling produces lower noise in high-frequency regions with small filter sizes (see right of bottom inset), with greater noise in low-frequency regions, that will be filtered out. Compare to (d) uniform standard stratified Monte Carlo sampling with uniformly distributed noise. Our method (e) after adaptive sampling and filtering is accurate at 63spp. (f) Equal error at 324 spp, which is still noisy. Overhead in our algorithm is minimal, and we provide a speedup vs equal error of $5\times$. Readers are encouraged to zoom into the PDF for this and all subsequent figures, to more clearly see the noise and image quality.

Abstract

We introduce an algorithm for interactive rendering of physically-based global illumination, based on a novel frequency analysis of indirect lighting. Our method combines adaptive sampling by Monte Carlo ray or path tracing, using a standard GPU-accelerated raytracer, with real-time reconstruction of the resulting noisy images. Our theoretical analysis assumes diffuse indirect lighting, with general Lambertian and specular receivers. In practice, we demonstrate accurate interactive global illumination with diffuse and moderately glossy objects, at 1-3 fps. We show mathematically that indirect illumination is a structured signal in the Fourier domain, with inherent band-limiting due to the BRDF and geometry terms. We extend previous work on sheared and axis-aligned filtering for motion blur and shadows, to develop an image-space filtering method for interreflections. Our method enables $5-8\times$ reduced sampling rates and wall clock times, and converges to ground truth as more samples are added. To develop our theory, we overcome important technical challenges—unlike previous work, there is no light source to serve as a band-limit in indirect lighting, and we also consider non-parallel geometry of receiver and reflecting surfaces, without first-order approximations.

CR Categories: Computing Methodologies [Computer Graphics]: Rendering—Ray Tracing

Keywords: Fourier, sampling, filtering, global illumination

Links: [DL](#) [PDF](#) [WEB](#) [VIDEO](#)

1 Introduction

Interactive rendering of indirect illumination is one of the grand challenges of computer graphics. In this paper, we take an important step towards solving this problem for diffuse interreflections, with both Lambertian and Phong receivers, based on physically-accurate Monte Carlo ray or path tracing, followed by image-space filtering. Monte Carlo integration at each pixel has long been regarded as the gold standard for accuracy—but not suitable for interactive use, with hundreds of samples needed and slow render times. This has led to a number of real-time but approximate alternatives, such as point-based gathering [Wang et al. 2009; Maletz and Wang 2011] or voxel-based cone tracing [Crassin et al. 2011]. We seek to obtain the best of both worlds; physically accurate and interactive.

We are inspired by recent work on sheared filtering for motion blur and soft shadows by Egan et al. [2009; 2011a; 2011b], which has demonstrated dramatically reduced sample counts. Most recently, Mehta et al. [2012] developed an axis-aligned filtering method for area light soft shadows on diffuse surfaces (axis-aligned or sheared refers to the pixel-light space, rather than the image domain, although the method also uses an axis-aligned image filter). This approach trades off a somewhat increased sample count for a much simpler filter, that reduces to an adaptive 2D image-space gaussian blur, does not require storage or search over an irregular 4D domain, and allows for adaptive sampling and adjustment of filter sizes to

ACM Reference Format

Mehta, S., Wang, B., Ramamoorthi, R., Durand, F. 2013. Axis-Aligned Filtering for Interactive Physically-Based Diffuse Indirect Lighting. ACM Trans. Graph. 32, 4, Article 96 (July 2013), 11 pages.
DOI = 10.1145/2461912.2461947 <http://doi.acm.org/10.1145/2461912.2461947>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

2013 Copyright held by the Owner/Author. Publication rights licensed to ACM.
0730-0301/13/07-ART96 \$15.00.

DOI: <http://dx.doi.org/10.1145/2461912.2461947>

guarantee convergence with more samples. However, none of these methods is designed for global illumination.

In this paper, we develop an axis-aligned filtering method, based on an analysis of the frequency-domain structure of the indirect light field. Our theory considers diffuse interreflections, but receiving surfaces can have general BRDFs. In practice, since our approach is based on accurate path tracing and always converges in the limit, our method works for diffuse and moderately glossy objects. Specific theoretical and practical contributions include:

Fourier Analysis of Indirect Illumination: Our main theoretical contribution (Sec. 3) is a frequency analysis of indirect lighting. We provide an exact derivation (Secs. 3.1, 3.2), without first-order assumptions inherent in previous works. This is essential for handling arbitrarily oriented surfaces—unlike the constant velocity assumption in motion blur [Egan et al. 2009], we cannot assume parallel receivers and reflecting surfaces. Oriented reflectors involve a nonlinear transformation in the Fourier analysis, but surprisingly the spectrum of the indirect light field still lies in a double wedge, bounded by the minimum and maximum depths of the reflector; this result enables us to leverage much of the filtering theory in previous work. We also extend the theory to glossy receivers (Sec. 3.3).

Bandlimits for Axis-Aligned Filtering of Indirect Lighting: Earlier work [Egan et al. 2011b; Mehta et al. 2012] used the light source (and its size) to bandlimit the occlusion function. However, there is no single light source in global illumination; we show that the geometry/parameterization and BRDF play the role of the bandlimit instead, and derive axis-aligned filter sizes, as well as adaptive sampling rates for diffuse and glossy surfaces (Sec. 4).

Interactive Sampled Global Illumination: We demonstrate interactive global illumination with one or more indirect bounces—with adaptive sampling and accurate Monte Carlo path tracing using NVIDIA's Optix GPU raytracer, followed by adaptive image filtering; an example is shown in Fig. 1, and later in Figs. 6, 7, 8, 11.

2 Previous Work

Our method builds on Monte Carlo ray and path tracing, introduced in seminal papers of Cook [1984] and Kajiya [1986]; these are still usually regarded as the gold standards for physically accurate rendering; we leverage their accuracy, while achieving interactivity.

Interactive Global Illumination: A number of brute-force and approximate methods exist for interactive global illumination [Ritschel et al. 2012], but do not usually guarantee physical accuracy or convergence. These include interactive raytracing [Wald et al. 2002; Wald et al. 2007]; we use the fast GPU raytracer in NVIDIA's Optix, but focus on reducing sample count and filtering. Our approach is orthogonal to GPU raytracer accelerations [van Antwerpen 2011]. Another method is approximate voxel-based cone tracing [Crassin et al. 2011] on the GPU. Point-based approaches include micro-rendering [Ritschel et al. 2009]. [Wang et al. 2009] (with refinements in [Maletz and Wang 2011]) raytrace shading points and partition them into coherent shading clusters using adaptive seeding followed by k-means, and then apply final gather to evaluate the irradiance using GPU-based photon mapping. In contrast, we sample every pixel using Monte Carlo.

Precomputation: Precomputation-based methods [Sloan et al. 2002] can be used for indirect illumination [Hasan et al. 2006] in relighting static scenes. Our approach does not require precomputation, and can be used with dynamic geometry.

Adaptive Sampling: Irradiance caching (IC) and gradients [Ward et al. 1988; Ward and Heckbert 1992] attempt to ex-

trapolate irradiance on diffuse surfaces from neighboring pixels, tracing a pixel only if the error is high. An extension to general low-frequency radiance is given by [Krivaneck et al. 2005]. [Guo 1998] provides a more general adaptive sampling heuristic.

Note that IC is used to determine where to put caches, not to compute a sampling rate, and is not an interactive technique. Some IC heuristics approximate our results, but our bandlimits are based on fundamental Fourier analysis. Closer to our approach, [Kontkanen et al. 2004] use the harmonic mean of occluder distances as a heuristic to set a filter width at each pixel, and use the total filter weight around each pixel to estimate its sampling rate. Like IC, they do not guarantee convergence either. In contrast, our adaptive filtering and sampling follows directly from the novel Fourier analysis derivations. We also do not require their 3D tree of pixel locations to search for potential contributing pixels at each point. Avoiding this search, and providing a simple gaussian filter, dramatically reduces overhead and enables interactivity.

Adaptive Filtering: Recently, several adaptive filtering and reconstruction methods have been proposed, but they are all designed for offline use. Building on [Hachisuka et al. 2008; Overbeck et al. 2009], Lehtinen et al. [2011; 2012] demonstrate GPU-accelerated reconstruction for temporal effects and indirect light fields, but their methods take several minutes. Similarly, we are inspired by recent work on iterative filtering of [Sen and Darabi 2012], and the anisotropic statistical filtering of [Li et al. 2012] and [Rouselle et al. 2012], as well as simple image denoising in graphics [Rushmeier and Ward 1994; McCool 1999]. But these are all offline methods, which enables more complex filtering and adaptive sampling.

In contrast, past work on fast global illumination has involved simple depth-space heuristic filters [Shirley et al. 2011], edge-avoiding wavelets [Dammertz et al. 2010], or filtering secondary scene attributes [Bauszat et al. 2011]. We differ in using frequency analysis to develop a spatially-varying image-space gaussian filter.

Frequency Analysis: We are inspired by Chai et al. [2000] and Durand et al. [2005], who introduce the basic space-angle and pixel-light Fourier theory on which we build, as do many previous works in this area [Egan et al. 2009; Soler et al. 2009; Egan et al. 2011a; Egan et al. 2011b; Belcour et al. 2013]. We build most directly on the axis-aligned filtering approach of [Mehta et al. 2012], which reduces to simple 2D image-space filtering rather than irregular reconstruction from the 4D light field, and can therefore be implemented very efficiently, with minimal time or memory overhead. We extend it non-trivially from soft shadows to global illumination, where there is no single light source, and we consider non-parallel geometry of receivers and reflectors as well as glossy receivers.

3 Fourier Analysis of Indirect Illumination

In this section, we perform a Fourier analysis of the indirect illumination light field (the incoming light from other objects, as a function of spatial location and incident angle). We assume that the direct lighting is computed separately, and focus on the global component, using the geometry in Fig. 2(a). We derive the wedge shape of the Fourier spectrum (Figs. 3b,c), which allows us (Sec. 4) to apply a suitable axis-aligned filter, with adaptive sampling rates.

We first consider a receiver, which is the surface seen at a pixel, illuminated by a reflector, which is the nearest surface in a particular direction. We will later see that the theory extends naturally to multiple reflectors at a range of distances and orientations. In essence, we are considering final gather, which adds up the full indirect light (including multiple bounces) from reflectors.

Indirect bounces of light will generally be lower frequency, since




Figure 2: (a) The (x, v) parameterization for the indirect illumination light field, and associated notation. (b) Showing how the indirect illumination is reflected towards the camera.

the BRDF acts as a low-pass filter at each bounce. This diffuse property of global illumination has been exploited in previous works [Nayar et al. 2006; Ben-Artzi et al. 2008]. We also utilize it in our theory, letting the receiver have a general BRDF, but assuming reflected light from the reflector is diffuse—we set filter sizes and sampling rates based on this analysis. The practical algorithm always uses radiance values from accurate path tracing and converges in the limit (but may slightly under or over-blur for non-diffuse indirect light with fewer samples, as in Fig. 9). Our results show we can handle diffuse and moderately glossy objects (Fig. 8).

We first derive the equations in flatland or 2D, using the (x, v) light field parameterization, as in [Durand et al. 2005], and then show the extension of bandlimits to 3D in Sec. 4. x is measured along the local plane of the receiver (globally, the receiver could of course be curved, as could the reflectors) and v is measured on a plane parallel to the receiver plane, and a unit distance from it. We make minimal assumptions on spatial properties of the global illumination; reflectors (and receivers) could have complex textures or high-frequency lighting. As we will see, the shape of the Fourier spectrum is determined by key features of the geometry and incoming light. The distance(s) of the reflector will determine the slope (or range of slopes) of the indirect light field. The BRDF of the receiver, and the parameterization, determine bandlimits of the Fourier spectrum.

3.1 Indirect Light from the Reflector

We first consider the indirect light from the reflector, and then show how this is integrated for global illumination at the receiver.

Assume the reflector slope is s relative to the receiver, as shown in Fig. 2a. We parameterize the indirect light field as $L_i(x, v)$ for a receiver point x in a direction v . Similarly, we parameterize a reflector point by x' , so that its coordinates are $(x', z') = (x', z_0 + sx')$, where z_0 is the intercept of the reflector line at $x' = 0$. The reflected light is then given by $L_r(x')$ —this can include high-frequency illumination, multiple bounces, and texture on the reflector, but has no directional information; we focus on diffuse inter-reflections. However, the receiver can be glossy as discussed later in Sec. 3.3. The reflector lies between (x_1, z_1) and (x_2, z_2) , so that $L_r(x') = 0$ unless $x_1 \leq x' \leq x_2$.

For a receiver point x in a direction v , the reflecting point is given by $x' = x + zv$, where z is the distance perpendicular to the receiver, as shown in Fig. 2(a). Simple geometry¹ dictates that

¹These relations do not apply when $x_1 = x_2$, so that $s = \infty$. Parameterizing L_r by z instead of x' in that case will lead to a similar result.

$z = z_0 + sx'$ from which it follows that

$$\begin{aligned} z = z_0 + s(x + zv) &\Rightarrow z = \frac{z_0 + sx}{1 - sv} \\ x' = x + zv &\Rightarrow x' = \frac{x + z_0v}{1 - sv}, \end{aligned} \quad (1)$$

which implies that the indirect light field $L_i(x, v) = L_r(x')$ is

$$L_i(x, v) = L_r\left(\frac{x + z_0v}{1 - sv}\right). \quad (2)$$

This is a simple relation of the indirect light field to the outgoing reflected light. In the special case that $s = 0$, when receiver and reflector are parallel, it reduces to $L_i(x, v) = L_r(x + z_0v)$. In that case, the 2D indirect light field L_i is the 1D reflected or outgoing light L_r , sheared by an amount proportional to the distance z_0 of the reflecting surface. This first-order simplification is analogous to the free space light propagation discussed in [Durand et al. 2005], and also mathematically similar to the relation for the visibility function in [Egan et al. 2011a]. Note however that there is no “visibility” term in our case; rather we are considering the indirect illumination field. There is also no separate lighting term, as there is for shadows. In essence, we consider the rendering equation, rather than the reflection equation, and integrate over the reflector surface. Moreover, we generalize many previous light field analyses, in explicitly considering a general slope s for the reflector, which leads to the more general rational form above (with denominator of $1 - sv$).

3.2 Fourier Spectrum of Incident Light Field

We now conduct a frequency-space analysis.² We denote Fourier domain quantities with a hat, and arguments using Ω , and with the symbol $j = \sqrt{-1}$. Equation 2 now becomes

$$\hat{L}_i(\Omega_x, \Omega_v) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} L_r\left(\frac{x + z_0v}{1 - sv}\right) e^{-j(x\Omega_x + v\Omega_v)} dx dv. \quad (3)$$

We first do the integral along x to compute the partial Fourier transform (denoted with a tilde on top). If v is held fixed, the argument to L_r is simply a scale and shift of x , with the Fourier transform being given by standard Fourier scale³ and shift theorems,

$$\begin{aligned} \tilde{L}_i(\Omega_x, v) &= e^{jz_0v\Omega_x} (1 - sv) \hat{L}_r[(1 - sv)\Omega_x] \\ \hat{L}_i(\Omega_x, \Omega_v) &= \int e^{-jv(\Omega_v - z_0\Omega_x)} (1 - sv) \hat{L}_r[(1 - sv)\Omega_x] dv. \end{aligned} \quad (4)$$

First, consider the special case of a parallel reflector with $s = 0$. In this case, $\hat{L}_r(\Omega_x)$ comes out of the integral, which reduces to a delta function with $\hat{L}_i(\Omega_x, \Omega_v) = \hat{L}_r(\Omega_x) \delta(\Omega_v - z_0\Omega_x)$. The Fourier spectrum is compact, essentially given by a shear of $\hat{L}_r$, and restricted to the line $\Omega_v = z_0\Omega_x$ (Fig. 3a). This has the same mathematical form as the visibility function in [Egan et al. 2011a; Mehta et al. 2012]. In those works, they address a range of depths simply by extending the frequency spectrum to a double wedge bounded by minimum and maximum depths. However, that extension is heuristic and not formally justified. Equation 4 is more general, explicitly considering a general sloped reflector; we proceed to formally derive the wedge spectrum in this case. We also explicitly consider the finite extent of the reflector, thereby treating (dis)occlusion (directions where there is no indirect light).

²For readability, we omit numerical constant factors to normalize the Fourier transforms, that do not affect the insights or the final results.

³Technically, the Fourier scale theorem requires the multiplicative factor $|1 - sv|$. Since the reflector has finite extent and is in front of the receiver to reflect on to it, we take the positive sign without loss of generality.




Figure 3: (a) Spectrum of the indirect light field for a diffuse reflector at a single depth z_0 is a line, (b) Spectrum of the indirect light field for a diffuse reflector in the depth range $[z_1, z_2]$ is a double wedge. Because of discontinuities (high derivatives) at reflector end-points, significant energy is concentrated at extreme slopes z_1 and z_2 . (c) shows a schematic of the double wedge and our axis-aligned filter determined by the BRDF/transfer bandlimit $\Omega_h^{\max}$. (d) Power spectral density of our flatland diffuse receiver transfer function h_d . $\Omega_h^{\max} = 2.0$ captures 95% energy and (e) the same PSD in 3D, where we need a slightly higher bandlimit of 2.8 (the blue box) to capture the same fraction.

In equation 4, we make the substitutions $u = (1 - sv)\Omega_x$ so that $v = (1 - u/\Omega_x)/s$ and $dv = -du/(s\Omega_x)$ so that,

$$\begin{aligned} \hat{L}_i(\Omega_x, \Omega_v) &= \int e^{-j(\Omega_v - z_0\Omega_x)(1-u/\Omega_x)/s} u \hat{L}_r(u) \frac{-du}{s\Omega_x^2} \\ &= \frac{-\exp[-j(\Omega_v - z_0\Omega_x)/s]}{s\Omega_x^2} \int u \hat{L}_r(u) e^{ju(\Omega_v - z_0\Omega_x)/(s\Omega_x)} du. \end{aligned} \quad (5)$$

The integral is now simply the inverse Fourier transform of $u \hat{L}_r(u)$, evaluated at $(\Omega_v - z_0\Omega_x)/(s\Omega_x)$. Recall the Fourier transform of the derivative L'_r is $\hat{L}'_r(u) = ju \hat{L}_r(u)$, so that,

$$\hat{L}_i(\Omega_x, \Omega_v) = \frac{j \exp[-j(\Omega_v - z_0\Omega_x)/s]}{s\Omega_x^2} L'_r\left(\frac{\Omega_v - z_0\Omega_x}{s\Omega_x}\right). \quad (6)$$

This is a general result, with the Fourier transform of the indirect light field expressed in terms of the (derivative of) spatial content of the reflected light, and evaluated at a sheared and scaled argument. The first term is simply a phase offset and an Ω_x^{-2} falloff. The second term is more interesting. Since $L_r(x)$ only takes non zero values for $x_1 \leq x \leq x_2$ (the extent of the reflector), the same holds for $L'_r(x)$. Then, $\hat{L}_r(\Omega_x, \Omega_v)$ takes non-zero values only when⁴

$$x_1 \leq \frac{\Omega_v - z_0\Omega_x}{s\Omega_x} \leq x_2 \Rightarrow (z_0 + sx_1)\Omega_x \leq \Omega_v \leq (z_0 + sx_2)\Omega_x. \quad (7)$$

By definition $z_0 + sx_1 = z_1$ and $z_0 + sx_2 = z_2$. Therefore,

$$z_1\Omega_x \leq \Omega_v \leq z_2\Omega_x. \quad (8)$$

In other words, the frequency spectrum lies in a double wedge, with slopes z_1 and z_2 bounded by the minimum and maximum depths of the reflector, as shown in Fig. 3(b). (Note also the significant energy at the extreme slopes z_1 and z_2 , since the derivative L'_r in equation 6 is large at the end-points of the reflector).

For the special case of a parallel reflector ($s = 0$), we have that $z_1 = z_2 = z_0$, and the spectrum is restricted to the single line $\Omega_v = z_0\Omega_x$, as shown in Fig. 3(a). By taking the limit of $s \rightarrow 0$ in equation 6, one can derive⁵ $\hat{L}_i(\Omega_x, \Omega_v) = L_r(\Omega_x) \delta(\Omega_v - z_0\Omega_x)$.

⁴The inequalities hold for $s > 0$. For $s < 0$, we must reverse the inequalities, but the same formula for the wedge in equation 8 is obtained as long as we adopt the convention that $z_1 < z_2$.

⁵Clearly, as $s \rightarrow 0$, the L'_r term becomes a delta function of the

Finally, for multiple reflectors, we combine the spectra for individual reflectors⁶, and use the double wedge bounded by the minimum and maximum depths of all reflectors, a schematic of which is shown in Fig. 3(c). Figure 4(c) verifies this numerically for a flatland scene with multiple reflectors that also occlude each other.

3.3 Outgoing Light from the Receiver

We now consider the actual image, corresponding to the outgoing light after it is reflected from the receiver. Let $f(v, v_c)$ be the receiver's BRDF. We do not explicitly consider texture in this section, which will simply modulate the reflected light. The reflected outgoing radiance towards camera c from x can be written (see Fig. 2(b)):

$$L_o(x, v_c(x)) = \int_{H^2} L_i(x, v) f(v, v_c) \cos \theta_i d\omega. \quad (9)$$

In flatland, $\cos \theta_i d\omega = (1 + v^2)^{-3/2} dv$, as is derived in [Durand et al. 2005] and elsewhere (from Fig. 2, $\cos \theta_i = 1/\sqrt{1 + v^2}$ and flatland 'solid angle' $d\omega = dv \cos \theta_i / \sqrt{1 + v^2}$). Hence,

$$L_o(x) = \int_{-\infty}^{\infty} L_i(x, v) h(v, v_c) dv \quad (10)$$

$$h(v, v_c) = \frac{f(v, v_c)}{(1 + v^2)^{3/2}}. \quad (11)$$

Here we have combined the BRDF, the cosine and the solid angle terms into a single transfer function h . The 3D extension of these results is straightforward, and discussed briefly in Sec. 4.2; we will see that we can use almost the same band-limits as derived from the flatland analysis. For a diffuse receiver with coefficient k_d ,

$$h_d(v) = k_d \frac{1}{(1 + v^2)^{3/2}}. \quad (12)$$

form $\delta(\Omega_v - z_0\Omega_x)$. To find the multiplier, by definition of the delta function, we must integrate equation 6 over Ω_v . First, substitute that $w = (\Omega_v - z_0\Omega_x)/(s\Omega_x)$ with $d\Omega_v = s\Omega_x dw$. Then, we simply want $\int (j \exp[-jw\Omega_x] \Omega_x^{-1}) L'_r(w) dw$. Simply integrating by parts, this reduces to $\int L_r(w) \exp[-jw\Omega_x] dw$ which is simply the Fourier transform $\hat{L}(\Omega_x)$. Thus, as $s \rightarrow 0$, $\hat{L}_i(\Omega_x, \Omega_v) = \hat{L}(\Omega_x) \delta(\Omega_v - z_0\Omega_x)$.

⁶As in most previous work on light field filtering [Chai et al. 2000; Egan et al. 2011b; Egan et al. 2011a; Mehta et al. 2012], the spectra do not strictly combine because of occlusion between different reflectors, and there can be some leakage outside the wedge for the multiple reflector case. As in previous work, the extent of leakage is empirically small and we neglect it.




Figure 4: Schematic for a flatland scene: (a) The configuration, where blue shows reflecting surface and black shows receiving surface, and the white point is the light source; (b) The incident indirect light field L_i in false color; (c) The numerical double-wedge spectrum of (b); (d) The spectrum in (c) is multiplied by the Fourier transform of h_d ; the black box is our filter, which captures 99% of the energy; (e) shows the Fourier transform of L_o obtained by integrating (d) corresponding to equation 14; the red bars mark our spatial domain filter.

An analytic formula for the Fourier transform of h_d is known in terms of Bessel functions, and is plotted in Fig. 3(d).

Now, consider a receiver Blinn-Phong BRDF (with exponent m and specular coefficient k_s). The half-angle is $(\tan^{-1} v + \tan^{-1} v_c)/2$,

$$h_s(v, v_c) = k_s \frac{\cos^m \left(\frac{1}{2} [\tan^{-1} v + \tan^{-1} v_c] \right)}{(1 + v^2)^{\frac{3}{2}}}. \quad (13)$$

In summary, Sec. 3 has analyzed the indirect light field in both the primal and Fourier domain. A key result is in formally deriving the wedge spectrum for an arbitrarily oriented reflector. This result is mathematically similar to (but more general than) previous analyses of spectra for soft shadows [Egan et al. 2011a; Egan et al. 2011b; Mehta et al. 2012]. Our derivation is more accurate, applying to general configurations of reflector and receiver. Moreover, there is no single light source per se, unlike shadows. In Sec. 4, we build on these results to develop the Fourier space bandlimits and filters for the transfer function, and to select the adaptive sampling rates.

4 Axis-Aligned Filtering

We now develop the bandwidths for an axis-aligned filter (which reduces to image-space blurring) in a fashion similar to [Mehta et al. 2012]. Our filter in the Fourier domain is shown in Fig. 3(c).

We first write the Fourier equivalent of equation 10, which is simply a frequency domain integral for $\hat{L}_o$, the spectrum of outgoing light:

$$\hat{L}_o(\Omega_x) = \int \hat{L}_i(\Omega_x, \Omega_v) \hat{h}(-\Omega_v) d\Omega_v, \quad (14)$$

where we keep the camera direction v_c argument in h and $\hat{h}$ implicit (or alternatively, since $v_c(x)$ is a function only of the spatial pixel x , we can include it in the BRDF term).

4.1 Fourier and Spatial Reconstruction Filter

In the expression above, the transfer function $\hat{h}$ (which includes the BRDF and geometry terms) acts as a low-pass band-limiting filter on the indirect illumination light field. It plays much the same role as the low-frequency lighting in [Egan et al. 2011a] or the area light width in [Mehta et al. 2012]. However, note the different representation that is necessary, compared to those papers. In essence, our indirect illumination field L_i is analogous to their visibility term, while our BRDF and geometry term is analogous to their illumination. Also note that we do not need to explicitly consider visibility.

Of course, there will be directions where no reflector is present and we receive no indirect light contribution; indeed, our Fourier spectrum derivation in Sec. 3.2 explicitly handles finite extent reflectors.

Consider the frequency bandlimits from equation 14. Assume $\hat{h}$ has a frequency bandlimit $\Omega_h^{\max}$. Then, frequencies in $\hat{L}_i$ higher than this value in Ω_v need not be considered during integration. This in turn induces a limit on the maximum spatial frequency of $\hat{L}_o$, as seen in Fig. 3(c). The width of our reconstruction filter Ω_x^r is:

$$\Omega_x^r = \mu \cdot \min \left\{ \frac{\Omega_h^{\max}}{z_{\min}}, \Omega_x^{\max} \right\}, \quad (15)$$

where μ is a scale factor (for now $\mu = 1$) that enables over-sampling and convergence, as discussed in Sec. 4.3, $z_{\min}$ is the minimum world-space distance to any reflector, and $\Omega_x^{\max}$ is the maximum spatial frequency that is always a limit (even if $z_{\min}$ is small). Similar to previous work, we define $\Omega_{\text{pix}}^{\max}$ as the maximum frequency in pixel space, with $\Omega_x^{\max} = \alpha \Omega_{\text{pix}}^{\max} = \alpha d^{-1}$, where d is the projected distance per pixel⁷ and $0 < \alpha < 1$ is a constant. We use $\alpha = 0.3$ in most of our renderings.

We integrate over the v dimension (or in Fourier space Ω_v). Therefore, axis-aligned filtering of the indirect light field in x - v space reduces to a spatial filter over the noisy indirect illumination, which is effectively a screen-space filter. In other words, we first compute the standard noisy Monte Carlo global illumination result $\bar{L}_o(x)$ using relatively few samples to evaluate equation 10. We then filter,

$$L_o(x) = \int N(x - y; \beta) \bar{L}_o(y) dy, \quad (16)$$

where N is the primal domain gaussian filter with standard deviation $\beta = 2/\Omega_x^r$, since the Fourier space standard deviation is taken to be $0.5 \times \Omega_x^r$ (in practice, we use Gaussian filters in spatial and frequency domains as in previous work, instead of sines or boxes).

4.2 Bandlimit of the Transfer Function

The remaining question is to find the numerical bandlimit $\Omega_h^{\max}$ for the transfer function. In these calculations, we can ignore the diffuse and specular coefficients k_d and k_s . For a Lambertian receiver, given by equation 12, Fig. 3(d) shows the power spectral density (PSD) of the transfer function, $|\hat{h}_d(\Omega_v)|^2$. $\Omega_h^{\max} = 2.0$ captures

⁷The lateral size in terms of actual distance that pixel covers. This term also naturally accounts for perspective effects and foreshortening.




Figure 5: (a) Compact packing of spectra of the indirect light field in the Fourier domain. Ω_x^* and Ω_y^* are the minimum sampling rates. (b) the per-pixel minimum sampling rate vs. fourier-space filter scaling μ (image space filter size is inversely proportional to μ) for two different occluder depth ranges.

95% of the energy in $\hat{h}$. Moreover, as shown in Fig. 4(d,e), we see that $\Omega_h^{\max} = 2.0$ is usually sufficient to capture approximately 99% energy in $\hat{L}_o$ since $\hat{L}_i$ also usually decays with frequency.

For the 3D case (4D light field), the transfer function depends on the $\mathbf{v} = (v_1, v_2)$ coordinates, and we denote this 3D extension as $H(\mathbf{v}, \mathbf{v}_c)$. We still have $\cos \theta_i = 1/\sqrt{1 + \mathbf{v} \cdot \mathbf{v}}$ and the solid angle is given by $d\omega = dv_1 dv_2 \cos \theta_i / (1 + v_1^2 + v_2^2)$. Therefore,

$$H(\mathbf{v}, \mathbf{v}_c) = \frac{f(v_1, v_2, v_{c1}, v_{c2})}{(1 + v_1^2 + v_2^2)^2}. \quad (17)$$

Note that our final filter is only along the x dimension, and the PSD of H is symmetric in Ω_{v1} and Ω_{v2} . Figure 3(e) shows the PSD for the diffuse transfer function—we can use a bandlimit slightly higher than the flatland case, $\Omega_h^{\max} = 2.8$, to capture 95% energy.

For the Blinn-Phong BRDF, the bandlimit $\Omega_h^{\max}$ depends on $\mathbf{v}_c$ and the exponent m . There is no closed form analytic expression for the Fourier transform, but we can still obtain a simple numerical linear fit between the exponent m and the band-limit for moderately glossy BRDFs, as discussed in the Appendix for the 3D case. In summary, we use the numerical values,

$$\Omega_{h,d}^{\max} = 2.8 \quad \Omega_{h,s}^{\max}(m) = 3.6 + 0.084m. \quad (18)$$

The appendix also gives bandlimits for the Phong BRDF.

4.3 Adaptive Sampling Rates

Discrete sampling of a continuous signal (here, the indirect light field) can cause aliasing if the sampling rate is not sufficient, even if we subsequently use the proper axis-aligned reconstruction filter. The minimum sampling rate is that which just prevents adjacent copies of spectra from overlapping. As in [Mehta et al. 2012], the most compact packing is that shown in Fig. 5(a), and the minimum sampling rates in the Ω_x and Ω_y axes are:

$$\begin{aligned} \Omega_x^* &= \Omega_x^r + \Omega_x^{\max} = \frac{\Omega_h^{\max}}{z_{\min}} + \alpha \Omega_{\text{pix}}^{\max} \\ \Omega_y^* &= \Omega_h^{\max} + z_{\max} \Omega_x^r = \Omega_h^{\max} \left(1 + \frac{z_{\max}}{z_{\min}} \right). \end{aligned} \quad (19)$$

The per-pixel sampling rate is then given as:

$$\begin{aligned} n &= [(\Omega_x^*)^2 \times A_p] \times [(\Omega_y^*)^2] \\ &= \left(\Omega_h^{\max} \frac{\sqrt{A_p}}{z_{\min}} + \alpha \right)^2 \times (\Omega_h^{\max})^2 \left(1 + \frac{z_{\max}}{z_{\min}} \right)^2 \end{aligned} \quad (20)$$

where A_p is the spatial area in world-space subtended by a pixel $A_p = d^2$, with $\sqrt{A_p} \Omega_{\text{pix}}^{\max} = 1$.

As in [Mehta et al. 2012], we can also increase the sampling rate beyond the minimum required, and simultaneously reduce the filter size, so that our algorithm converges to Monte Carlo ground truth in the limit. In the Fourier domain, we increase the filter size by a factor of $\mu > 1$, such that $\Omega_x^r = \mu \Omega_{x0}^r$, where Ω_{x0}^r is the critical size given by equation 15. Then, the sampling rate increases to

$$n(\mu) = \gamma \cdot \left(\mu \Omega_h^{\max} \frac{\sqrt{A_p}}{z_{\min}} + \alpha \right)^2 \times (\Omega_h^{\max})^2 \left(1 + \mu \frac{z_{\max}}{z_{\min}} \right)^2, \quad (21)$$

where γ is a scaling factor for importance sampling discussed shortly. We can also use equation 21 with $\mu < 1$; in practice, we find $\mu = 0.9$ adequate in most cases. Figure 5(b) shows how the effective samples per pixel vary with μ .

Importance Sampling Adjustment: The theory is derived assuming uniform angular sampling. In rendering, more efficient importance sampling is used, where samples in the important regions are already closer together. This enables us to adjust $\gamma < 1$ in equation 21 for lower sampling rates, where γ depends on the receiver BRDF. For the diffuse cosine lobe, the incident angle may vary in $[0, \pi/2]$ but about 80% of energy is in $[0, \pi/4]$, enabling $\gamma \approx 1/2$. In practice, we use $\gamma = 0.4$, which we find provides adequate quality with lower sampling counts. For the specular lobe, setting $c = \cos(\pi/4) = 1/\sqrt{2}$, about 80% of energy is concentrated in the range $[0, \cos^{-1}(c^{1/m})]$, and so we set $\gamma = \cos^{-1}(c^{1/m})/(\pi/2)$. For Blinn-Phong exponent $m = 20$, this sets $\gamma = 0.11$.

5 Implementation

Given our derivation, the final implementation is simple. Our code uses the NVIDIA OptiX real-time raytracing framework [Parker et al. 2010] on a GTX 690 GPU. We only need to implement a simple extension for multi-bounce path tracing where desired, as well as the core of our adaptive sampling and filtering passes in the OptiX GPU framework. Our method runs entirely on the GPU in three passes, as described below:

Initial Raytracing for Filter Sizes and Adaptive Sampling:

We first trace 16 stratified rays over the visible hemisphere from each pixel/shading point, to compute the per-pixel $z_{\min}$ and $z_{\max}$, which are simply min/max world-space distances to geometry (reflectors).⁸ Example outputs are shown in Fig. 6(f). Direct lighting is also computed in the standard way in this pass (for this paper, we used a single point light source in all examples), and we create buffers for pixel k_d and k_s values so we can separate textures from radiance computation. We then use $z_{\min}$ and $z_{\max}$ to compute the filter sizes and sampling rates according to equations 15 and 21 respectively, separately for the diffuse and glossy components of the pixel. We use $\mu = 0.9$ and $\alpha = 0.3$, with γ set as discussed earlier.

For efficiency, we also use radiance values from the initial 16 samples, so we actually only need $n - 16$ samples in the second pass. Finally, we clamp the maximum sampling rate (the minimum must

⁸As in most adaptive sampling approaches, using only 16 initial rays can give noisy or incomplete estimates of $z_{\min}$ and $z_{\max}$, but we never visualize them directly, and only use these to set filter sizes and sampling rates. As seen in our results, we found the final images to be high quality, with any additional accuracy not worth the cost of more rays. We also tried filtering the z estimates, but found this actually worsened image quality/speed, leading to overly conservative or aggressive sampling rates and filter sizes.




Figure 6: (a) The diffuse Conference scene with 331K triangles is rendered at 3 fps with our method and 60spp, using 1-bounce of indirect illumination. Insets show (b) Equal time uniform MC at 64 spp is still very noisy; (c) our method compared to (d) equal error uniform MC at 440 spp and (e) ground truth at 1024 spp; (f) minimum and maximum reflector distances $z_{\min}$, $z_{\max}$ (g) screen-space filter radius in pixels, and per pixel sampling rates. Our algorithm obtains accurate results $7\times$ faster than basic path tracing.




Figure 7: A Cornell Box scene (a) 3-bounce Indirect illumination, our method using an average 56 spp; insets showing (b) our method compared, to equal error (c) at 298 spp, a $5\times$ speedup.

already be at least 16 because of initial sampling). We use a maximum of 100 samples, scaled up when $\mu > 1$, i.e., 100μ . Correspondingly, we also clamp the minimum $z_{\min}$ to about 2% of maximum scene dimension, since a $z_{\min} = 0$ implies no filtering and leaves corners and edges (such as wall intersections) noisy.

We visualize the resulting sampling rates and filter widths in Figs. 1(b) and 6(g). It is clear that smaller filter sizes and higher sampling rates are used for close-by geometry, while regions with further reflectors, such as the floor in Fig. 1, can use lower sampling rates and wider filters.

Path Tracing: We use OptiX to (stratified and importance) sample each pixel’s hemisphere with the adaptive sampling rates. If more than one bounce is desired, we use standard path tracing for higher bounces. We compute the incoming radiance values and multiply with the BRDF (excluding textures). This gives a noisy estimate $\bar{L}_o$ at each pixel, as shown in Fig. 1(c). Notice that adaptive sampling ensures almost accurate results in high-complexity areas, such as the bottom of the curtain (right of second inset), while showing considerable noise (many fewer samples) on the floor. This is as desired with sampling optimized for filtering; larger filter widths on the floor ensure an accurate final image (Fig. 1(e)).

Screen-Space Filtering: In the final pass, we actually do Gaussian filtering on $\bar{L}_o$, using the spatial filter (equation 16). As

in [Mehta et al. 2012] we use a 2D screen-space filter, with a depth buffer giving the world-space coordinates and distances required. Moreover, we perform the standard separation of the Gaussian filter into two 1D filters along the image dimensions for efficiency.⁹ Finally, we modulate filtered radiance by the textures for k_d and k_s .

In summary, the method consists of path-tracing (the core optimized OptiX raytracer), coupled with initial estimates to set filter sizes, adjust sampling rates, and do screen-space filtering. The latter steps are simple to implement and extremely fast. Overhead is minimal compared to Optix (Table 1), and we achieve interactive speeds of 1-3 frames per second, often with nearly an order of magnitude fewer samples than basic GPU path tracing.

6 Results

We show results of interactive global illumination on five scenes in Figs. 1, 6, 7, 8 and 11. The accompanying video shows animations and screen captures with moving light source, viewpoint and examples of dynamic geometry (no precomputation is required; each frame is rendered independently). To focus on the global component, all paper images show indirect light only; direct illumination is added separately in the first step (initial raytracing in Sec. 5).

Visual Fidelity and Speedups: Our method is accurate in a range of different scenarios, with consistent reductions in sample counts over basic path tracing. Figure 1 shows an example of the Sponza scene (262K triangles) with 1-bounce diffuse indirect illumination (surfaces are Lambertian) and textures. Our image (Fig. 1a,e) is accurate with 63 average samples per pixel. Equal error with stratified Monte Carlo is only achieved for 324 samples, and is visually still noisy. Our method adds minimal overhead, so this is a speedup of $5\times$. Figure 6 shows similar results for the diffuse conference scene with 331K triangles. Our method requires only 60 samples per pixel. The overhead is minimal, with equal time Monte Carlo being only 64 samples. Equal error (still noisy) is only obtained for 440 samples, a speedup of about $7\times$. We demonstrate the extension to multi-bounce diffuse indirect illumination by path tracing in Fig. 7. Note the curved surfaces of the dragon, and complex structures on the elephant’s trunk in the insets, that our method renders accurately at 56 samples per pixel (close examination will show a slight overblur on the trunk).

⁹As in [Mehta et al. 2012], this separation is strictly accurate only if the gaussian width β is the same across all pixels in the filter but in practice we find almost no noticeable difference, since β usually varies slowly.




Figure 8: (a) 2-bounce indirect illumination on a Cornell Box scene with glossy vase and teapots ($m = 12$), with average 85 spp. A close comparison to ground truth is shown in (b) and (c); in (d) we show 2-bounce indirect illumination on the Sibenik Cathedral scene, with a glossy textured floor and red carpet ($m = 20$), with average 86 spp. Insets in (e) and (f) show the accuracy of our filtering on glossy receivers.




Figure 9: Evaluating the theoretical approximation of diffuse indirect light. (a) Shows over-blurring (and hence darkening) of glossy-to-diffuse and glossy-to-glossy in an inset from the Sibenik Cathedral scene, when using $\mu = 1$ at 86 spp. However, we match ground truth (b) closely using $\mu = 2$ at 226 spp, shown in (c).

Figure 8 shows that we can handle moderately glossy receivers for a modified Cornell box (vase and teapots are glossy with Blinn-Phong exponent 12), and for the Sibenik scene with textures and 75K triangles (floor and red carpet have Blinn-Phong exponent 20). Nevertheless, the method matches accurately, capturing the glossy reflections, and color-bleeding near edges. Again, a speedup of about $5\times$ is achieved, over simply using path-tracing in OptiX.

Figure 9 evaluates our use of the “diffuse interreflection” approximation to set sampling rates in Sec. 3, which does not fully consider high frequencies in glossy-to-diffuse or glossy-to-glossy interactions (however, receiver glossiness is fully handled as noted above). While our filter sizes depend on this approximation, note that our method operates on accurate path tracing input, and will therefore always converge in the limit with more samples. In Fig. 9 we see that we slightly overblur (and hence darken) for $\mu = 1$, but glossy-to-diffuse transfer is almost fully accurate for $\mu = 2$.

Quantitative Accuracy: We also evaluated our method quantitatively in Fig. 10, showing RMS errors vs average number of samples for the conference scene from Fig. 6. The error of our method (blue curve) is significantly below stratified Monte Carlo at all sample counts. Moreover, even just using filtering on uniform sampling (green curve) gives a substantial improvement. Note that our method is physically-based and consistent. As we increase the number of samples (higher μ), we do converge to ground truth and error decreases (also shown visually in insets). This is in contrast to most other solutions for fast, approximate global illumination.




Figure 10: (a) Shows a log-log plot of the RMS pixel error vs average sampling rate for the Conference scene of Fig. 6. In (b) we show insets from the conference scene (showing error relative to ground truth magnified $10\times$). It can be seen that our method converges to ground truth both numerically and visually with more samples.

Timings and Overhead: In table 1, we show timings for steps of our algorithm on different scenes, rendered at a resolution of 640×480 . We obtain most of the benefits of axis-aligned filtering, first shown in the soft shadow work of [Mehta et al. 2012], even though our algorithm is somewhat more complex (involving additional texture buffers, and passes to handle diffuse and specular components). The total overhead in a frame is about 60-70 ms, which is small compared to the cost of OptiX path tracing, and results in only a marginal decrease in the performance of the real-time raytracer. Note that additional bounces in path tracing do cause a slowdown in the Optix raytracer, essentially linear in the number of bounces as expected, but do not affect our algorithm.

Note that our filter operates only in image-space and therefore has limited memory requirements (buffers for the noisy image, depth, and textures). We achieve a sample count reduction and speedup of $5 - 8\times$ on most scenes, with interactive frame rates of 1 – 3 fps. Note that we are limited only by the speed of the real-time raytracer, and using further GPU raytracing accelerations would provide further speedups. To our knowledge, this is one of the first

scene	triangles	avg. spp	num. bounces	Optix Ray/Path tracing			Our algorithm			Total	
				1st pass (ms)	2nd pass (ms)	total (ms)	compute spp (ms)	adaptive filtering (ms)	total overhead (ms)	total time (ms)	fps (with/without alg)
Sponza (Fig. 1)	262 K	63	1	181	580	761	10	55	65	826	1.21 / 1.31
Conference (Fig. 6)	331 K	60	1	87	274	361	9	46	55	416	2.40 / 2.77
Cor. Box (Fig. 7)	145 K	56	1	70	291	361	9	63	72	433	2.31 / 2.77
			2	70	802	872	10	62	72	944	1.06 / 1.14
			3	72	1423	1495	10	64	74	1569	0.63 / 0.67
Sibenik (Fig. 8d)	75 K	86	1	110	440	550	10	50	60	610	1.64 / 1.82
			2	108	1402	1510	12	52	64	1574	0.64 / 0.67
Cor. Box (Fig. 11)	16.7 K	59	1	70	248	318	12	49	61	379	2.64 / 3.14
			2	68	622	690	10	57	67	757	1.32 / 1.44

Table 1: Timings of our scenes rendered at 640×480 . Our filtering overhead is small compared to the sampling time, and our impact on the fps is small.




Figure 11: A Cornell Box scene (a) 1-bounce indirect illumination, our method using an average 59 spp at 2.7 fps; insets showing (b) our method unfiltered, (c) our method, filtered, (d) Equal error image with uniform MC 361 spp, (e) Importance Point Projection (IPP, arrows show high frequency details are lost and some regions are over-darkened) (f) Image denoising (BM3D, arrows show blurred geometry edges and artifacts due to noise in input) (g) Reference image with 1024 spp.

demonstrations of accurate interactive global illumination, based on principled Monte Carlo sampling. Alternative methods, discussed next, either sacrifice accuracy, or are offline, adding overheads of seconds to minutes.

Comparisons: In Fig. 11, we include comparisons to two alternative approaches. First, Fig. 11(e) considers the importance point-projection (IPP) method for global illumination ([Maletz and Wang 2011], building on [Wang et al. 2009]). We use the authors’ code on their scene. Because only a few shading points are used, errors can occur in regions of complex interreflections, especially near edges, as shown in the insets. Figure 11(f) compares to standard image denoising using a state of the art denoiser bm3d [Dabov et al. 2007]; this still takes a few seconds, and can overblur or underblur (retain noise) in some regions, since it is not informed of the precise filter size from our frequency analysis. Newer denoisers for rendering [Li et al. 2012; Rouselle et al. 2012] have impressive results, but often report render times in the minutes, as does the indirect light field reconstruction approach of [Lehtinen et al. 2012]. Thus, previous approaches aim for high quality or high accuracy; we differ in providing *accurate* results for *interactive* global illumination.

Extension—Ambient Occlusion: Our theory applies to any outgoing function on the reflectors. If we simply set it to 1, then we will directly compute the complement of ambient occlusion (rays that hit a reflector return 1, rays that miss return 0). This is an accurate interactive ambient occlusion calculation, several orders of magnitude faster than [Egan et al. 2011a]. We show an example result in Fig. 12, computed at interactive rates of 1-2 fps.




Figure 12: Ambient Occlusion on the Sponza scene using an average of 63 spp runs at 1.2 fps.

Limitations: As shown in our numerical simulations and real experiments, the double wedge model is a good representation of the Fourier spectrum for indirect illumination. However, highly curved receivers and reflectors can lead to ‘leaking’ spectra outside theoretical predictions. This is not usually a significant practical problem (see accurate dragon and elephant insets in Fig. 7), and we also al-

leviate it in a number of ways. We implement a normal threshold (set to 10°), beyond which radiance values are not used in the filtering. This avoids filtering across very different orientations.¹⁰ Our axis-aligned filter is looser than previous sheared filters to capture more of the spectrum leaks. And, we can always increase sample count (and μ) to converge to ground truth.

Our average sample counts are low (~ 60) for global illumination, but still high enough to prevent fully real-time performance on complex scenes. However, we are 1-3 orders of magnitude faster than equally accurate Monte-Carlo based rendering methods, providing *interactive* frame rates. Finally, the theory is limited to diffuse interreflections on general BRDF receivers; in practice (Figs. 8, 9) it works for Lambertian and moderately glossy receivers/reflectors.

7 Conclusions and Future Work

A large body of recent work has shown that physically-based sampling followed by adaptive reconstruction, sharing samples between pixels, can be effective at rendering complex effects like motion blur, depth of field and soft shadows. However, the challenge of accurate interactive global illumination has remained elusive. This paper bridges that gap, presenting an interactive method for sampled physically-based interreflections. We show how to map the indirect illumination field into the double wedge Fourier spectrum even for non-parallel reflectors and receivers, how to band-limit this spectrum based on the BRDF and geometry (rather than lighting) terms, and how to handle general Blinn-Phong receivers. We show a variety of results with 1 to 3-bounce interreflections involving diffuse and Phong surfaces rendered at 1-3 fps, and our method guarantees convergence as more Monte Carlo samples are added. While we have focused on interreflections in this paper, similar ideas could be applied in future to other global illumination effects like animated scenes, caustics, multiple specular transports, or Metropolis sampling, filtering over path space. We believe the “sample and denoise” paradigm to be quite powerful as a unifying approach to achieve real-time physically-based global illumination.

Acknowledgements: We thank the reviewers for their helpful suggestions. This work was supported in part by NSF grant CGV # 1115242, 1116303, the Intel Science and Technology Center for Visual Computing, and a Siebel Scholarship. We thank NVIDIA for a last-minute donation of graphics cards, and acknowledge additional equipment and funding from Adobe and Pixar.

References

- BAUSZAT, P., EISEMANN, M., AND MAGNOR, M. 2011. Guided image filtering for interactive high quality global illumination. *Computer Graphics Forum (EGSR 11)* 30, 4, 1361–1368.
- BELCOUR, L., SOLER, C., SUBR, K., HOLZSCHUCH, N., AND DURAND, F. 2013. 5D covariance tracing for efficient defocus and motion blur. *ACM Transactions on Graphics (to appear) [MIT-CSAIL-TR-2012-034]*.
- BEN-ARTZI, A., EGAN, K., RAMAMOORTHY, R., AND DURAND, F. 2008. A Precomputed Polynomial Representation for Interactive BRDF Editing with Global Illumination. *ACM Transactions on Graphics* 27, 2, Article 13, 1–13.
- CHAI, J., CHAN, S., SHUM, H., AND TONG, X. 2000. Plenoptic Sampling. In *SIGGRAPH 00*, 307–318.
- COOK, R., PORTER, T., AND CARPENTER, L. 1984. Distributed Ray Tracing. In *SIGGRAPH 84*, 137–145.
- CRASSIN, C., NEYRET, F., SAINZ, M., GREEN, S., AND EISEMAN, E. 2011. Iterative indirect illumination using voxel cone tracing. *Computer Graphics Forum* 30, 7, 1921–1930.
- DABOV, K., FOI, A., KATKOVNIK, V., AND EGIAZARIAN, K. 2007. Image denoising by sparse 3D transform-domain collaborative filtering. *IEEE Transactions on Image Processing* 16, 8, 2080–2095.
- DAMMERTZ, H., SEWITZ, D., HANIKA, J., AND LENSCH, H. 2010. Edge-avoiding a-trous wavelet transform for fast global illumination filtering. In *High Performance Graphics (HPG)*, 67–75.
- DURAND, F., HOLZSCHUCH, N., SOLER, C., CHAN, E., AND SILLION, F. 2005. A Frequency Analysis of Light Transport. *ACM Transactions on Graphics (Proc. SIGGRAPH 05)* 25, 3, 1115–1126.
- EGAN, K., TSENG, Y., HOLZSCHUCH, N., DURAND, F., AND RAMAMOORTHY, R. 2009. Frequency analysis and sheared reconstruction for rendering motion blur. *ACM Transactions on Graphics* 28, 3.
- EGAN, K., DURAND, F., AND RAMAMOORTHY, R. 2011. Practical filtering for efficient ray-traced directional occlusion. *ACM Transactions on Graphics (SIGGRAPH Asia 11)* 30, 6.
- EGAN, K., HECHT, F., DURAND, F., AND RAMAMOORTHY, R. 2011. Frequency analysis and sheared filtering for shadow light fields of complex occluders. *ACM Transactions on Graphics* 30, 2.
- GUO, B. 1998. Progressive radiance evaluation using directional coherence maps. In *SIGGRAPH 98*, 255–266.
- HACHISUKA, T., JAROSZ, W., WEISTROFFER, R., DALE, K., HUMPHREYS, G., ZWICKER, M., AND JENSEN, H. 2008. Multidimensional adaptive sampling and reconstruction for ray tracing. *ACM Transactions on Graphics* 27, 3.
- HASAN, M., PELLACINI, F., AND BALA, K. 2006. Direct to Indirect Transfer for Cinematic Relighting. *ACM Transactions on Graphics (Proc. SIGGRAPH 06)* 25, 3, 1089–1097.
- KAJIYA, J. 1986. The Rendering Equation. In *SIGGRAPH 86*, 143–150.
- KONTKANEN, J., RÄSÄNEN, J., AND KELLER, A. 2004. Irradiance filtering for monte carlo ray tracing. In *Monte Carlo and Quasi-Monte Carlo Methods 2004*, Springer, 259–272.
- KRIVANEK, J., GAUTRON, P., PATTANAIK, S., AND BOUA-TOUCH, K. 2005. Radiance caching off efficient global illumination computation. *IEEE Transactions on Visualization and Computer Graphics* 11, 5, 550–561.
- LEHTINEN, J., AILA, T., CHEN, J., LAINE, S., AND DURAND, F. 2011. Temporal light field reconstruction for rendering distribution effects. *ACM Transactions on Graphics* 30, 4.
- LEHTINEN, J., AILA, T., LAINE, S., AND DURAND, F. 2012. Reconstructing the indirect light field for global illumination. *ACM Transactions on Graphics* 31, 4.
- LI, T., WU, Y., AND CHUANG, Y. 2012. SURE-based optimization for adaptive sampling and reconstruction. *ACM Transactions on Graphics (SIGGRAPH Asia 2012)* 31, 6.

¹⁰This reduced filter size may however lead to fewer samples in the filter footprint than theory recommends. It is possible to adjust the adaptive sampling for this, but we found the overhead added by this additional pass not worth the marginal improvement in quality.

- MALETZ, D., AND WANG, R. 2011. Importance point projection for GPU-based final gathering. *Computer Graphics Forum (EGSR 11)* 30, 4, 1327–1336.
- MCCOOL, M. 1999. Anisotropic diffusion for monte carlo noise reduction. *ACM Transactions on Graphics* 18, 2, 171–194.
- MEHTA, S., WANG, B., AND RAMAMOORTHY, R. 2012. Axis-aligned filtering for interactive sampled soft shadows. *ACM Transactions on Graphics (SIGGRAPH Asia 12)* 31, 6.
- NAYAR, S., KRISHNAN, G., GROSSBERG, M., AND RASKAR, R. 2006. Fast separation of direct and global components of a scene using high frequency illumination. *ACM Transactions on Graphics (SIGGRAPH 2006)* 25, 3.
- OVERBECK, R., DONNER, C., AND RAMAMOORTHY, R. 2009. Adaptive Wavelet Rendering. *ACM Transactions on Graphics (SIGGRAPH ASIA 09)* 28, 5.
- PARKER, S., BIGLER, J., DIETRICH, A., FRIEDRICH, H., HOBEROCK, J., LUEBKE, D., MCALLISTER, D., MCGUIRE, M., MORLEY, K., ROBISON, A., AND STICH, M. 2010. OptiX: A general purpose ray tracing engine. *ACM Transactions on Graphics* 29, 4, 66:1–66:13.
- RITSCHER, T., ENGELHARDT, T., GROSCH, T., SEIDEL, H., KAUTZ, J., AND DACHSBACHER, C. 2009. Micro-rendering for scalable, parallel final gathering. vol. 28.
- RITSCHER, T., DACHSBACHER, C., GROSCH, T., AND KAUTZ, J. 2012. The state of the art in interactive global illumination. *Computer Graphics Forum* 31, 1, 160–188.
- ROUSELLE, F., KNAUS, C., AND ZWICKER, M. 2012. Adaptive rendering with non-local means filtering. *ACM Transactions on Graphics (SIGGRAPH Asia 2012)* 31, 6.
- RUSHMEIER, H., AND WARD, G. 1994. Energy preserving non-linear filters. 131–138.
- SEN, P., AND DARABI, S. 2012. On filtering the noise from the random parameters in monte carlo rendering. *ACM Transactions on Graphics* 31, 3.
- SHIRLEY, P., AILA, T., COHEN, J., ENDERTON, E., LAINE, S., LUEBKE, D., AND MCGUIRE, M. 2011. A local image reconstruction algorithm for stochastic rendering. In *ACM Symposium on Interactive 3D Graphics*, 9–14.
- SLOAN, P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed Radiance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments. *ACM Transactions on Graphics (Proc. SIGGRAPH 02)* 21, 3, 527–536.
- SOLER, C., SUBR, K., DURAND, F., HOLZSCHUCH, N., AND SILLION, F. 2009. Fourier depth of field. *ACM Transactions on Graphics* 28, 2.
- VAN ANTWERPEN, D. 2011. Improving SIMD efficiency for parallel monte carlo light transport on the GPU. In *High Performance Graphics*.
- WALD, I., BENTHIN, C., SLUSALLEK, P., KOLLIG, T., AND KELLER, A. 2002. Interactive global illumination using fast ray tracing. In *Rendering Techniques (EGWR 02)*.
- WALD, I., MARK, W. R., GÜNTHER, J., BOULOS, S., IZE, T., HUNT, W., PARKER, S. G., AND SHIRLEY, P. 2007. State of the Art in Ray Tracing Animated Scenes. In *STAR Proceedings of Eurographics 07*, The Eurographics Association, D. Schmalstieg and J. Bittner, Eds., 89–116.




Figure 13: (a) Geometry for Phong and Blinn-Phong BRDFs in our parameterization. (b) Our linear approximation for the transfer function bandlimit $\Omega_h^{\max}$ for Blinn-Phong and Phong BRDFs, as a function of the exponent m , is a close match to the numerical data.

WANG, R., WANG, R., ZHOU, K., PAN, M., AND BAO, H. 2009. An efficient GPU-based approach for interactive global illumination. *ACM Transactions on Graphics* 28, 3.

WARD, G., AND HECKBERT, P. 1992. Irradiance Gradients. In *Eurographics Rendering Workshop 92*, 85–98.

WARD, G., RUBINSTEIN, F., AND CLEAR, R. 1988. A ray tracing solution for diffuse interreflections. In *SIGGRAPH 88*, 85–92.

A Appendix: Bandlimits for Glossy BRDFs

The 2D Blinn-Phong BRDF in the $\mathbf{v}$ parameterization was described in equation 13. Here, we discuss the extension to 3D and resulting bandlimit $\Omega_h^{\max}$. Referring to Fig. 13(a), the half vector in the $\mathbf{v}$ plane is

$$\mathbf{v}_h = \frac{l\mathbf{v}_c + l_c\mathbf{v}}{l + l_c} \quad (22)$$

where $l = \sqrt{1 + v_1^2 + v_2^2}$ and $l_c = \sqrt{1 + v_{c1}^2 + v_{c2}^2}$. Then the 3D Blinn-Phong transfer function is

$$H_s(\mathbf{v}, \mathbf{v}_c) = \frac{\cos^m \theta_h}{l^4} = \frac{1}{l^4(1 + v_{h1}^2 + v_{h2}^2)^{m/2}} \quad (23)$$

The Fourier transform of H_s will depend on the direction $\mathbf{v}_c$ and has no simple analytic form. However, since the choice of the v_1, v_2 axes is arbitrary (and we never actually compute using this parameterization), we simply use $\Omega_h^{\max}$ as an upper bound on the bandlimit of H_s for any value of $\mathbf{v}_c$ (given fixed m). For 95% energy, we can plot a curve of $\Omega_h^{\max}$ vs m , as shown in Fig. 13(b), which can be fit with a simple linear approximation in the range $4 < m < 50$,

$$\Omega_h^{\max}(m) = 3.6 + 0.084m. \quad (24)$$

Similarly, in the Phong case, the angle of interest θ_{rc} is between the reflection direction $(1, -\mathbf{v})$ and the camera direction $(1, \mathbf{v}_c)$. Hence, the 3D Phong transfer function is

$$H_s(\mathbf{v}, \mathbf{v}_c) = \frac{\cos^m \theta_{rc}}{l^4} = \frac{(1 - v_1 v_{c1} - v_2 v_{c2})^m}{l^4 l_c^m} \quad (25)$$

The bandlimit now closely fits the approximation,

$$\Omega_h^{\max}(m) = 4.27 + 0.15m. \quad (26)$$

