

A Frequency Analysis of Light Transport

Frédo Durand
MIT-CSAIL

Nicolas Holzschuch
ARTIS* GRAVIR/IMAG-INRIA

Cyril Soler
MIT-CSAIL

Eric Chan
François X. Sillion
ARTIS* GRAVIR/IMAG-INRIA

Abstract

We present a signal-processing framework for light transport. We study the frequency content of radiance and how it is altered by phenomena such as shading, occlusion, and transport. This extends previous work that considered either spatial or angular dimensions, and it offers a comprehensive treatment of both space and angle.

We show that occlusion, a multiplication in the primal, amounts in the Fourier domain to a convolution by the spectrum of the blocker. Propagation corresponds to a shear in the space-angle frequency domain, while reflection on curved objects performs a different shear along the angular frequency axis. As shown by previous work, reflection is a convolution in the primal and therefore a multiplication in the Fourier domain. Our work shows how the spatial components of lighting are affected by this angular convolution.

Our framework predicts the characteristics of interactions such as caustics and the disappearance of the shadows of small features. Predictions on the frequency content can then be used to control sampling rates for rendering. Other potential applications include precomputed radiance transfer and inverse rendering.

Keywords: Light transport, Fourier analysis, signal processing

1 Introduction

Light in a scene is transported, occluded, and filtered by its complex interaction with objects. By the time it reaches our eyes, radiance is an intricate function, and simulating or analyzing it is challenging.

Frequency analysis of the radiance function is particularly interesting for many applications, including forward and inverse rendering. The effect of local interactions on the frequency content of radiance has previously been described in a limited context. For instance, it is well-known that diffuse reflection creates smooth (low-frequency) light distributions, while occlusion and hard shadows create discontinuities and high frequencies. However, a full characterization of global light transport in terms of signal processing and frequency analysis presents two major challenges: the domain of light rays is intricate (three dimensions for position and two for direction), and light paths can exhibit an infinite number of bounces (i.e. in terms of signal processing, the system has dense feedback).

To address the above challenges, we focus on the neighborhood of light paths [Shinya et al. 1987]. This restriction to local properties is both a price to pay and a fundamental difficulty with the problem we study: characteristics such as reflectance or presence and size of blockers are *non-stationary*, they vary across the scene.

This paper presents a theoretical framework for characterizing light transport in terms of frequency content. We seek a deep understanding of the frequency content of the radiance function in a

* ARTIS is a team of the GRAVIR lab (UMR 5527), a joint unit of CNRS, INPG, INRIA and UJF.

Copyright © 2005 by the Association for Computing Machinery, Inc.
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for 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 ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions Dept, ACM Inc., fax +1 (212) 869-0481 or e-mail permissions@acm.org.
© 2005 ACM 0730-0301/05/0700-1115 \$5.00

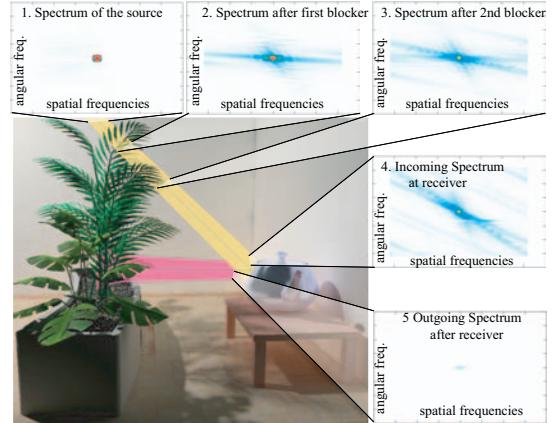


Figure 1: Space-angle frequency spectra of the radiance function measured in a 3D scene. We focus on the neighborhood of a ray path and measure the spectrum of a 4D light field at different steps, which we summarize as 2D plots that include only the radial components of the spatial and angular dimensions. Notice how the blockers result in higher spatial frequency and how transport in free space transfers these spatial frequencies to the angular domain. Aliasing is present in the visualized spectra due to the resolution challenge of manipulating 4D light fields.

scene and how it is affected by phenomena such as occlusion, reflection, and propagation in space (Fig. 1). We first present the two-dimensional case for simplicity of exposition. Then we show that it extends well to 3D because we only consider local neighborhoods of rays, thereby avoiding singularities on the sphere of directions.

Although we perform our derivations in an abstract setting, we keep practical questions in mind. In particular, we strongly believe that a good understanding of frequency creation and attenuation allows for more efficient sampling strategies for stochastic approaches such as Monte-Carlo global illumination. Furthermore, it leads to better sampling rates for light-field rendering, precomputed radiance transfer, and related applications. Finally, our frequency analysis can shed key practical insights on inverse problems and on the field of statistics of natural images by predicting which phenomena can cause certain local-frequency effects.

1.1 Contributions

This paper makes the following contributions:

Frequency study in space and angle. Our framework encompasses both *spatial* and *directional* variations of radiance, while most previous work studied only one of these two components.

Local surface interactions. We describe the frequency effects of local shading, object curvature, and spatially-varying BRDF.

Global light-transport. We provide expressions for the frequency modification due to light transport in free space and occlusion.

Most of the derivations in this paper are carried out in 2D for clarity, but we show that the main characterizations extend to 3D.

1.2 Related work

Radiance exhibits both spatial and angular variations. A wealth of previous work has studied the frequency content along one of these components, but rarely have both space and angle been addressed. We do not discuss all applications of Fourier analysis, but rather focus on studies of frequency modification in light transport.

Filtering and sampling Heckbert's seminal work on texture anti-aliasing [1989] derives *local* bandwidth for texture pre-filtering based on a first-order Taylor expansion of the perspective transform. The effect of perspective is also studied in the contexts of holography and light field sampling [Halle 1994; Isaksen et al. 2000; Chai et al. 2000; Stewart et al. 2003], mostly ignoring visibility and specular effects.

Local illumination as a convolution Recently, local illumination has been characterized in terms of convolution and it was shown that the outgoing radiance is band-limited by the BRDF [Ramamoorthi and Hanrahan 2001b; Ramamoorthi and Hanrahan 2004; Basri and Jacobs 2003]. However the lighting is assumed to come from infinity and occlusion is ignored. Frolova et al. [2004] explored spatial lighting variations, but only for convex diffuse objects. We build on these approaches and extend them by adding spatial dimensions as well as other phenomena such as occlusion and transport, at the expense of first-order approximations and a local treatment. Ramamoorthi et al. [2004] have also studied local occlusion in a textured object made of pits such as a sponge. Our treatment of occlusion considers complex blockers at an arbitrary distance of the blocker and receiver.

Wavelets and frequency bases Wavelets and spherical harmonics have been used extensively as basis functions for lighting simulation [Gortler et al. 1993; Keller 2001] or pre-computed radiance transfer [Sloan et al. 2002; Ramamoorthi and Hanrahan 2002]. They are typically used in a data-driven manner and in the context of projection methods, where an oracle helps in the selection of the relevant components based on the local frequency characteristics of radiance. Refinement criteria for multiresolution calculations often implicitly rely on frequency decomposition [Sillion and Drettakis 1995]. In our framework we study the frequency effect of the *equations* of light transport in the spirit of linear systems, and obtain a more explicit characterization of frequency effects. Our results on the required sampling rate can therefore be used with stochastic methods or to analyze the well-posedness of inverse problems.

Ray footprint A number of techniques use notions related to bandwidth in a ray's neighborhood and propagate a footprint for adaptive refinement [Shinya et al. 1987] and texture filtering [Igehy 1999]. Chen and Arvo use perturbation theory to exploit ray coherence [2000]. Authors have also exploited *on-the-fly* the frequency content of the image to make better use of rays [Bolin and Meyer 1998; Myszkowski 1998; Keller 2001]. Our work is complementary and provides a framework for frequency-content prediction.

Illumination differentials have been used to derive error bounds on radiance variations (e.g. gradients [Ward and Heckbert 1992; Annen et al. 2004], Jacobians [Arvo 1994], and Hessians [Holzschuch and Sillion 1998], but only provide local information, which cannot easily be used for sampling control.

Fourier analysis has also been extensively used in optics [Goodman 1996], but in the context of wave optics where phase and interferences are crucial. In contrast, we consider geometric optics and characterize frequency content in the visible spatial frequencies. The varying contrast sensitivity of humans to these spatial frequencies can be exploited for efficient rendering, e.g. [Bolin and Meyer 1995; Ferwerda et al. 1997; Bolin and Meyer 1998; Myszkowski 1998]. Finally we note that the Fourier basis can separate different phenomena and thus facilitate inverse lighting [Ramamoorthi and Hanrahan 2001b; Basri and Jacobs 2003] depth from focus [Pentland 1987] and shape from texture [Malik and Rosenholtz 1997].

ℓ_R	2D light field (2D) around ray R
x	spatial dimension (distance to central ray)
v	directional dimension in 2-plane parameterization
θ	directional dimension in plane-sphere parameterization
\widehat{f}	Fourier transform of function f
Ω_X	frequency along dimension X
i	$\sqrt{-1}$
$f \otimes g$	convolution of f by g
d	Transport distance
$V(x, v)$	visibility function of the blockers
$\cos_+(\theta)$	clamped cosine term: $\max(\cos \theta, 0)$
dE	Differential irradiance (after cosine term).
ρ	BRDF

Figure 2: Notations.

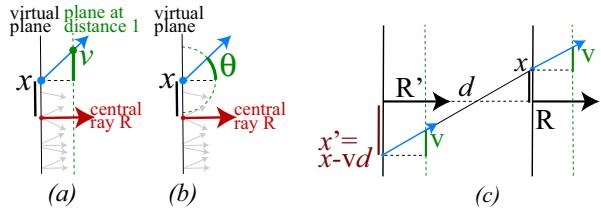


Figure 3: (a-b) The two light field parameterization used in this article. Locally, they are mostly equivalent: we linearize $v = \tan \theta$. (c) Transport in free space: the angular dimension v is not affected but the spatial dimension is reparameterized on v .

2 Preliminaries

We want to analyze the radiance function in the neighborhood of a ray along all steps of light propagation. For this, we need a number of definitions and notations, summarized in Fig. 2. Most of the derivations in this paper are carried out in 2D for clarity, but we shall see that our main observations extend naturally to 3D.

2.1 Local light field and frequency content

We consider the 4D (resp. 2D) slice of radiance at a virtual plane orthogonal to a central ray. We focus on the neighborhood of the central ray, and we call radiance in such a 4D (resp. 2D) neighborhood slice a *local light field* (Fig. 3 left). Of the many parameterizations that have been proposed for light fields, we use two distinct ones in this paper, each allowing for a natural expression of some transport phenomena. Both use the same parameter for the spatial coordinates in the virtual plane, x , but they differ slightly in their treatment of directions. For our *two-plane parameterization*, we follow Chai et al. [2000] and use the intersection v with a parallel plane at unit distance, expressed in the local frame of x (Fig. 3-a). In the *plane-sphere parameterization*, we use the angle θ with the central direction (Fig. 3-b) [Camahort et al. 1998]. These two parameterizations are linked by $v = \tan \theta$ and are equivalent around the origin thanks to a linearization of the tangent.

We study the Fourier spectrum of the radiance field ℓ_R , which we denote by $\widehat{\ell}_R$. For the two-plane parameterization, we use the following definition of the Fourier transform:

$$\widehat{\ell}_R(\Omega_x, \Omega_v) = \int_{x=-\infty}^{\infty} \int_{v=-\infty}^{\infty} \ell_R(x, v) e^{-2i\pi\Omega_x x} e^{-2i\pi\Omega_v v} dx dv \quad (1)$$

Examples are shown for two simple light sources in Fig. 4, with the spatial dimension along the horizontal axis and the direction along the vertical axis. We discuss the plane-sphere parameterization in Section 4.

One of the motivations for using Fourier analysis is the convolution-multiplication theorem, which states that a convolution

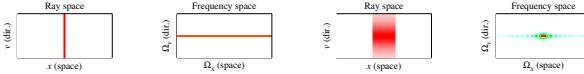


Figure 4: (a) A point light source is a Dirac in space times a constant in angle. (b) Its Fourier transform is a constant in space times a Dirac in angle. (c) A spot light with a finite-size bulb has a smooth falloff in angle. (d) Its Fourier transform is a sinc times a bell curve.

in the primary domain corresponds to a multiplication in the Fourier domain, and vice-versa. As we show in this paper, it affords a compact formulation of frequency modification.

2.2 Overview

When light flows in a scene, phenomena such as transport in free space, occlusion, and shading each modify the local light field in a characteristic fashion. These operations are described (in 2D) in Section 3 as filters operating on the frequency signal $\widehat{\ell}$. In Section 4, we describe the general case of local shading and extend the presentation to 3D in Section 5. Section 6 compares our framework with previous work and shows a simple application.

3 Transport phenomena as linear filters

This section describes the effect on frequency content of successive stages of light transport. All phenomena are illustrated in Fig. 6 for a simple 2D scene, where light is emitted at the source, transported in free space, occluded by obstacles, transported again, and reflected by a surface. At each step, we show a space-direction plot of radiance, in primal and frequency space, as well as space-direction frequency-domain plots obtained in a similar 3D scene (Fig. 9-b). Note the excellent qualitative agreement between the 2D predictions and 3D observations.

3.1 Travel in free space

Travel in free space is a crucial operation because the directional variation also turns into spatial variation. Consider a slide projector: At the source, we have a Dirac in space and the image in the directional domain. At the receiver, the signal of the image is present in a combination of space and angle. When light travels in free space, the value of radiance is unchanged along a ray, and travel in free space is a reparameterization of the local light field (Fig. 3-c). The value of radiance at a point x after transport can be found using:

$$\ell_R(x, v) = \ell_{R'}(x - vd, v), \quad (2)$$

where d is the travel distance. To compute the Fourier transform $\widehat{\ell}_R$, we insert the change of variable $x' = x - vd$ in the integral of Eq. 1:

$$\begin{aligned} \widehat{\ell}_R(\Omega_x, \Omega_v) &= \int_{x'=-\infty}^{\infty} \int_{v=-\infty}^{\infty} \ell_R(x', v) e^{-2i\pi\Omega_x(x'+vd)} e^{-2i\pi\Omega_v v} dx' dv \\ &= \int_{x'=-\infty}^{\infty} \int_{v=-\infty}^{\infty} \ell_R(x', v) e^{-2i\pi\Omega_x x'} e^{-2i\pi(\Omega_v + d\Omega_x)v} dx' dv, \end{aligned}$$

This is a shear in the directional dimension (Fig. 6, steps 2 and 4):

$$\widehat{\ell}_{R'}(\Omega_x, \Omega_v) = \widehat{\ell}_R(\Omega_x, \Omega_v + d\Omega_x) \quad (3)$$

The longer the travel, the more pronounced the shear.

3.2 Visibility

Occlusion creates high frequencies and discontinuities in the radiance function. Radiance is multiplied by the binary occlusion function of the occluders:

$$\ell_{R'}(x, v) = \ell_R(x, v) V(x, v) \quad (4)$$

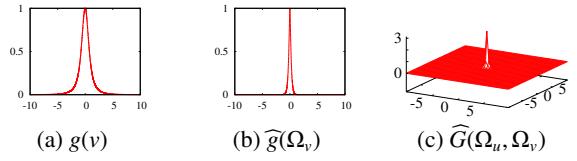


Figure 5: Unit area in lightfield parameterisation.

where $V(x, v)$ is equal to 1 when there is full visibility and to 0 when the occluders are blocking light transport. At the location of occlusion, V mostly depends on x (Fig. 6, step 3).

According to the multiplication theorem, such a multiplication amounts to a convolution in the frequency domain:

$$\widehat{\ell}_{R'}(\Omega_x, \Omega_v) = \widehat{\ell}_R(\Omega_x, \Omega_v) \otimes \widehat{V}(\Omega_x, \Omega_v) \quad (5)$$

If the occluders are lying inside a plane orthogonal to the ray, the occlusion function is a constant in angle, and its Fourier transform is a Dirac in the angular dimension. In the general case of non-planar occluders, their spectrum has frequency content in both dimensions, but the angular frequency content is restricted to a wedge with a span proportional to the depth extent of the blockers. Our formulation also handles semi-transparent blockers by using non-binary occlusion functions.

After occlusion, another transport step usually occurs, shearing the spectrum and propagating the occlusion from the spatial dimension to the angular dimension (Fig. 6, step 4).

3.3 Local diffuse shading

We first treat local shading by a planar Lambertian reflector. Curved and glossy reflectors will be treated in greater details in Section 4.

For a diffuse reflector with albedo ρ , the outgoing radiance ℓ_o has no directional variation; it is simply the integral over all directions of the incoming radiance per unit area on the receiver surface:

$$\ell_o(x) = \rho \int_{\Omega} \ell_i(x, v) dA_{\perp} \quad (6)$$

dA_{\perp} , the differential area on the receiver surface is $\cos_{+} \theta$ times the Jacobian of the lightfield parameterization; with $v = \tan \theta$, we have:

$$dA_{\perp} = \frac{dv}{(1+v^2)^{\frac{3}{2}}} = g(v) dv$$

We introduce $dE(x, v) = \ell_i(x, v) g(v)$, the differential irradiance at point x from the direction v . Since dE is a product of two functions, its spectrum is the convolution of their spectra:

$$\widehat{dE}(\Omega_x, \Omega_v) = \widehat{\ell}_i(\Omega_x, \Omega_v) \otimes \widehat{g}(\Omega_v)$$

The reflected radiance ℓ_o is the integral of dE over all directions v ; it is therefore the value of \widehat{dE} at $\Omega_v = 0$, that is, $\widehat{\ell}_o(\Omega_t) = \rho \widehat{dE}(\Omega_t, 0)$. Putting everything together, we have:

$$\widehat{\ell}_o(\Omega_x) = \rho \left[\widehat{\ell}_i(\Omega_x, \Omega_v) \otimes \widehat{g}(\Omega_v) \right]_{\Omega_v=0} \quad (7)$$

$g(v)$ is a bell-like curve (Fig. 5-a); its Fourier transform is:

$$\widehat{g}(\Omega_v) = 4\pi |\Omega_v| K_1(2\pi |\Omega_v|)$$

where K_1 is the first-order modified Bessel function of the second kind. \widehat{g} is highly concentrated on low frequencies (Fig. 5-b); the effect of convolution by \widehat{g} is a very small blur of the spectrum in the angular dimension (Fig. 6, step 5).

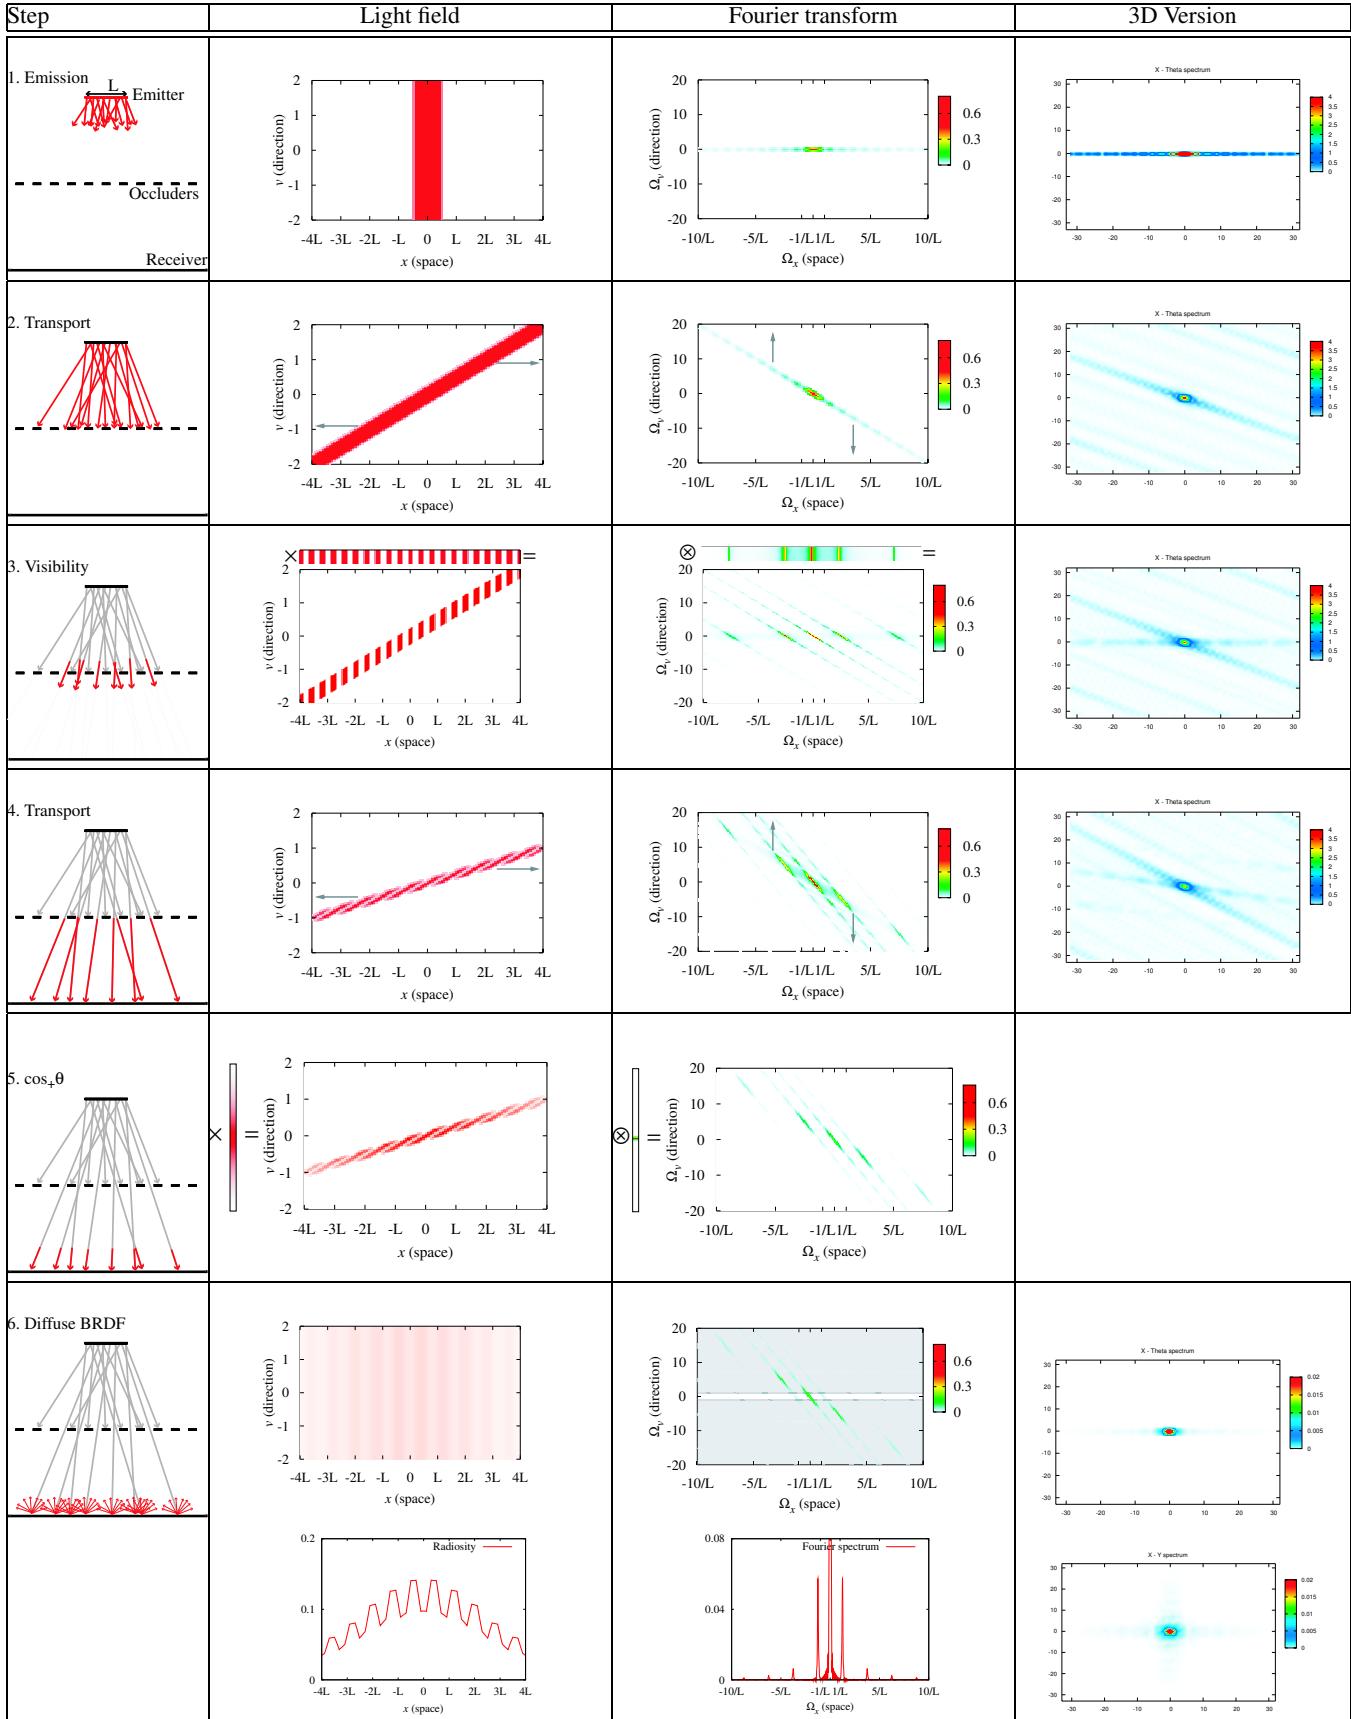


Figure 6: Effects on the spectrum of the various steps of light transport with a diffuse reflector. 2D Fourier transforms for steps 1 to 4 are obtained analytically; step 5 (convolution) is performed numerically. 3D Version spectrums are obtained numerically, *via* a Photon-Mapping algorithm and a FFT of the light field computed.

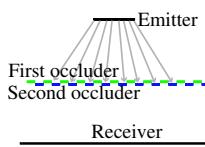
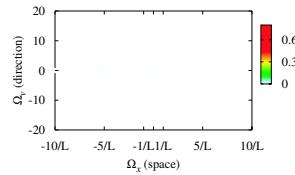
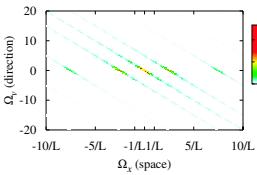


Figure 7: Scene configuration for visibility experiment. Left: spectrum with only one occluder. Right: spectrum with two occluders, computed with full precision and phase.

3.4 Example and discussion

Fig. 6 illustrates the various steps of light transport for a simple scene such as Fig. 9b. The slopes of the transport shears correspond to the travel distance (steps 2 and 4). Visibility increases the spatial frequency content through the convolution by a horizontal kernel in frequency space (step 3). There are only a finite number of blockers in Fig. 6, which explains why their spectrum is not a Dirac comb times a sinc, but a blurry version. The blocker spectrum mostly contains a main central lobe corresponding to the average occlusion and two side lobes corresponding to the blocker main frequency. This results in a replication of the sheared source spectrum on the two sides. The smaller the blocker pattern, the further-away these replicas are in frequency space. The final diffuse integration (step 6) discards all directional frequencies.

The main differences between the 3D and 2D plots of the spectra in Fig. 6 come from aliasing problems that are harder to fix with the 4D light field. Furthermore, in the 3D scene, the position of the blockers is jittered (see Fig. 9), which results in a smoother spectrum.

Feature-based visibility The spectra in Fig. 6 show that the second transport (step 4) pushes the ‘replicas’ to the angular domain. This effect is more pronounced for high-frequency blockers, for which the replicas are farther from the vertical line. Since the final diffuse integration keeps only the spatial line of frequencies (step 5), the main high-frequency lobe of the blockers is eliminated by diffuse shading. This is related to the feature-based approach to visibility [Sillion and Drettakis 1995], where the effect of small occluders on soft shadows is approximated by an average occlusion. However, our finding goes one step further: where the feature-based technique *ignores* high frequencies, we show that, for small-enough blockers, most high-frequencies are effectively removed by integration.

Combining several blockers A difficult scene for visibility is the case of two occluders that individually block half of the light, and together block all the light (Fig. 7). In our framework, if one carries out the computations with full precision, taking phase into account, one gets the correct result: an empty spectrum (Fig. 7, right).

However, for practical applications, it is probably not necessary to compute the full spectrum. Instead, we consider elements of information about the maximal frequency caused by the scene configuration, as we show in Section 6.2. In that case, one can get an overestimation of the frequencies caused by a combination of blockers, but not an underestimation.

4 General case for surface interaction

So far, we have studied only diffuse shading for a central ray normal to a planar receiver (although rays in the neighborhood have a non-normal incident angle). We now discuss the general case, taking into account the incidence angle, arbitrary BRDF, receiver curvature as well as spatial albedo variation. Our framework builds upon

Ramamoorthi and Hanrahan [2001b] and extends it in several ways, which we discuss in Section 6.1.

In a nutshell, local reflection simply corresponds to a multiplication by the cosine term and a convolution by the BRDF. However, a number of reparameterizations are necessary to take into account the incidence and outgoing angles, as well as the surface curvature. We first treat the special case of rotation-invariant BRDFs such as Phong before addressing more general forms as well as texture and spatially-varying BRDFs. Recall that we study frequency content in ray neighborhoods, which means that for local reflection, we consider an incoming neighborhood and an outgoing neighborhood.

Plane-sphere parameterization Since local reflection mostly involves integrals over the directional dimension, it is more naturally expressed in a parameterization where angles are uniform. This is why we use here a plane-sphere parameterization where the directional component θ is the angle to the central ray (Fig. 3-b). The spatial dimension is unaffected.

In the plane-sphere parameterization, the domain of directions is the S^1 circle, which means that frequency content along this dimension is now a Fourier series, not a transform. Fig. 8 shows the effect of reparameterizing angles on the frequency plane. The frequency distribution is very similar, although the spectrum is blurred by the non-linear reparameterization. For bandwidth analysis, this introduces no significant error. Note that for all local interactions with the surface (and thus in this entire section), there is no limitation to small values of θ , the linearization $v = \tan \theta \approx \theta$ will only be used again after light leaves the surface, for subsequent transport.

4.1 Rotation-invariant BRDFs on curved receivers

Local shading is described by the shading equation

$$\ell_o(x_i, \theta_o) = \int_{\theta_i} \ell(x_i, \theta_i) \rho_{x_i}(\theta'_o, \theta'_i) \cos_+ \theta'_i \, d\theta_i \quad (8)$$

where the primed angles are in the local frame of the normal while the unprimed angles are in the global frame (Fig. 10). For now, we assume that the BRDF ρ does not vary with x_i . Local shading is mostly a directional phenomenon with no spatial interaction: the outgoing radiance at a point is only determined by the incoming radiance at that point. However, the normal varies per point.

As pointed out by Ramamoorthi and Hanrahan [2001b], local reflection combines quantities that are naturally expressed in a global frame (incoming and outgoing radiance) and quantities that live in the local frame defined by the normal at a point (cosine term and BRDF). For this, we need to rotate all quantities at each spatial location to align them with the normal. This means that we rotate (reparameterize) the incoming radiance, perform local shading in the local frame, and rotate (reparameterize) again to obtain the outgoing radiance in a global frame. All steps of the local shading process are illustrated in Fig. 10 and discussed below.

Step 1 & 7: Reparameterization into the tangent frame We first take the central incidence angle θ_0 into account, and reparameterize in the local tangent frame with respect to the central normal direction. This involves a shift by θ_0 in angle and a scale in space

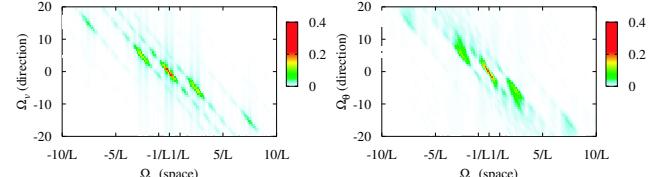


Figure 8: Spectrum arriving at the receiver (step 4 of Fig. 6), before and after the sphere-plane reparameterization. Left: (Ω_x, Ω_v) spectrum. Right: $(\Omega_x, \Omega_\theta)$ spectrum.

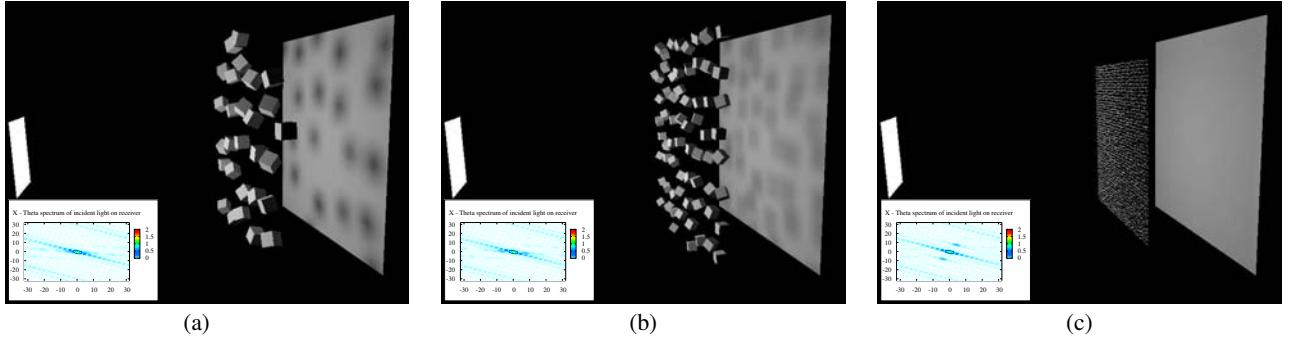


Figure 9: Complex frequency effects in light transport. The three scenes have the same area light and diffuse receiver and differ only by the frequency content of the blockers. (a) Large blockers result in few high frequencies. (b) With smaller (higher frequency) blockers, high frequencies increase on the receiver. (c) For very high-frequency blockers, high frequencies on the receiver nearly disappear.

by $1/\cos \theta_0$. We also flip the directions so that incident rays are pointing up and match the traditional local reflection configuration (Fig. 10), step 1). We omit the full derivation for brevity and provide directly the equations corresponding to step 1 and 7 of Fig. 10:

$$\widehat{\ell}_i(\Omega_x, \Omega_\theta) = e^{-i\Omega_\theta \theta_0} / |\cos \theta_0| \widehat{\ell}(-\Omega_x \cos \theta_0, \Omega_\theta) \quad (9)$$

$$\widehat{\ell}'(\Omega_x, \Omega_\theta) = e^{i\Omega_\theta \theta_1} |\cos \theta_1| \widehat{\ell}_o(\Omega_x / \cos \theta_1, \Omega_\theta) \quad (10)$$

Step 2: Per-point rotation The directional slice corresponding to each point must be shifted to rotate it in the local frame of the normal at that point (Fig. 10 step 2): $\theta'_i = \theta_i - \alpha(x_i)$.

For a smooth surface, we use a first-order Taylor expansion of the angle α of the normal at a point x_i . Given the curvature k , we have $\alpha(x_i) = kx_i$ and the reparameterization is $\theta'_i = \theta_i - kx_i$. This is a shear, but now along the directional dimension, in contrast to the transport shear. Similarly, the Fourier transform is sheared along the spatial dimension (Fig. 10 step 2, last row):

$$\widehat{\ell}'_i(\Omega'_x, \Omega'_\theta) = \widehat{\ell}_i(\Omega'_x + k\Omega'_\theta, \Omega'_\theta) \quad (11)$$

After this reparameterization, our two-dimensional spatio-directional local light field is harder to interpret physically. For each column, it corresponds to the incoming radiance in the frame of the local normal: the frame varies for each point. In a sense, we have unrolled the local surface and warped the space of light ray in the process [Wood et al. 2000]. The direction of the shear depends on the sign of the curvature (concave vs. convex).

Step 3: Cosine term and differential irradiance In the local frame of each point, we compute differential irradiance by multiplying by the spatially-constant clamped cosine function \cos_+ . This multiplication corresponds in frequency space to a convolution by a Dirac in space times a narrow function in angle:

$$\widehat{dE}'(\Omega_x, \Omega_\theta) = \widehat{\ell}'_i(\Omega_x, \Omega_\theta) \otimes \widehat{\cos_+}(\Omega_\theta) \delta_{\Omega_x=0} \quad (12)$$

Over the full directional domain, the spectrum of \cos_+ is:

$$\widehat{\cos_+}(\Omega_\theta) = \cos(\pi^2 \Omega_\theta) \frac{2}{1 - (2\pi \Omega_\theta)^2} \quad (13)$$

Most of the energy is centered around zero (Fig. 12-a) and the $1/\Omega_\theta^2$ frequency falloff comes from the derivative discontinuity¹ at $\pi/2$. Equivalent to the two-plane reparameterization (Section 3.3), the cosine term has only a small vertical blurring effect.

¹A function with a discontinuity in the n th derivative has a spectrum falling off as $1/\Omega^{n+1}$. A Dirac has constant spectrum.

Step 4: Mirror-direction reparameterization Common BRDFs mostly depend on the difference between the mirror reflection and the outgoing direction. This is why we remap the local incoming light using a mirror reflection around the normal (Fig. 10 step 4): $dE'_r(\theta'_r) = dE'(-\theta'_i)$

$$\widehat{dE}'_r(\Omega_x, \Omega'_\theta) = \widehat{dE}'(\Omega_x, -\Omega'_\theta) \quad (14)$$

This is equivalent to reparameterizations of surface light fields and BRDFs [Wood et al. 2000; Ramamoorthi and Hanrahan 2002].

Step 5: BRDF convolution In the mirror parameterization, assuming that the BRDF depends only on the angle difference, the shading equation 8 becomes:

$$\ell'_o(x_i, \theta'_o) = \int_{\theta'_r} dE'_r(x_i, \theta'_r) \rho'(\theta_o - \theta'_r) d\theta'_r \quad (15)$$

Which is a convolution of dE'_r by ρ' for each x_i ; that is, we convolve the 2D function dE'_r by a spatial Dirac times the directional shift-invariant BRDF ρ' (Fig. 10 step 5). In the Fourier domain, this is a multiplication by a spatial constant times the directional spectrum of the BRDF.

$$\widehat{\ell}'_o(\Omega'_x, \Omega'_\theta) = \widehat{dE}'_r(\Omega'_x, \Omega'_\theta) \widehat{\rho'}(\Omega'_\theta) \quad (16)$$

Note, however, that our expression of the BRDF is not reciprocal. We address more general forms of BRDF below.

Step 6: Per-point rotation back to tangent frame We now apply the inverse directional shear to go back to the global frame. Because we have applied a mirror transform in step 4, the shear and inverse shear double their effect rather than canceling each other. Since the shear comes from the object curvature, this models the effect of concave and convex mirror and how they deform reflection. In particular, a mirror sphere maps the full 360 degree field to the 180 degree hemisphere, as exploited for light probes.

4.2 Discussion

The important effects due to curvature, cosine term, and the BRDF are summarized in Fig. 10. Local shading is mostly a directional phenomenon, and the spatial component is a double-shear due to curvature (step 2 and 6). The cosine term results, in frequency space, in a convolution by a small directional kernel (step 3) while the BRDF band-limits the signal with a multiplication of the spectrum (step 5). Rougher materials operate a more aggressive low-pass, while in the special case of mirror BRDFs, the BRDF is a Dirac and the signal is unchanged.

Curvature has no effect on the directional bandwidth of the outgoing light field, which means that previous bounds derived in the

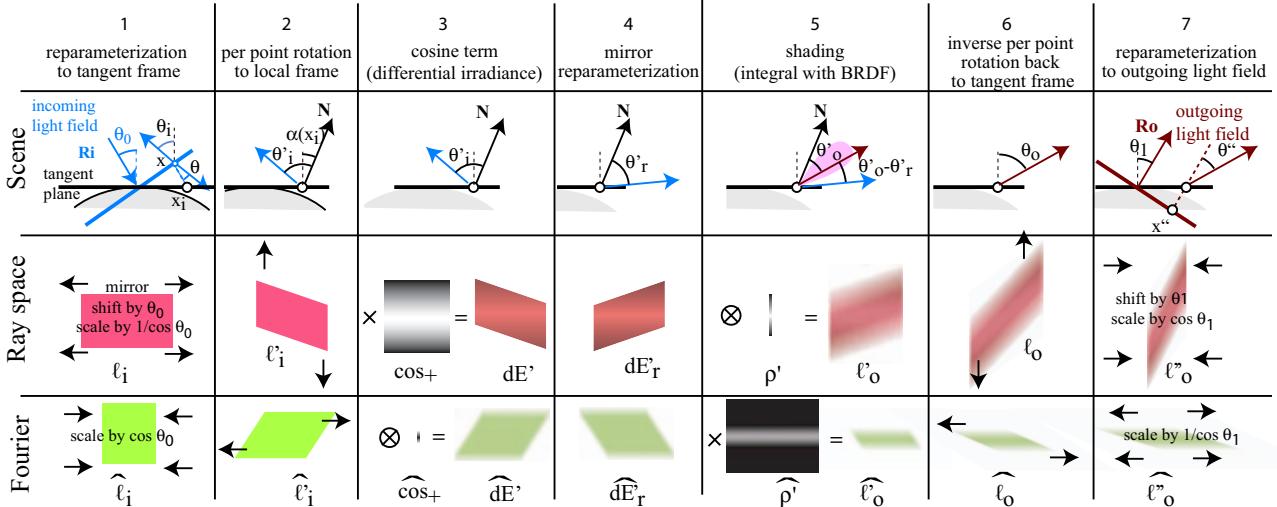


Figure 10: Local shading for a curved receiver with arbitrary BRDF.

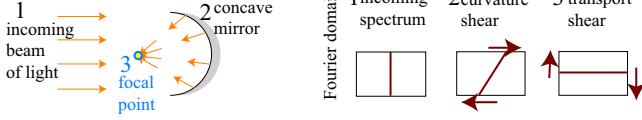


Figure 11: Caustic due to negative curvature. A shear in angle and then in space are combined to result in a transfer from the directional frequencies to the spatial frequencies.

special case of infinite lighting [Ramamoorthi and Hanrahan 2004; Basri and Jacobs 2003; Ramamoorthi and Hanrahan 2001a; Ramamoorthi and Hanrahan 2002] are valid for spatially-varying illumination. However, the spatial frequency content is strongly affected by curvature, which has important practical implications.

The effect of the curvature shear is further increased by the spatial scaling back to the tangent frame in step 7, as described by Eq. 10. We stress that this explains the well-known difficulty in sampling specular lighting in situations such as backlighting on the silhouette of a curved object. This is modeled by the effect of the curvature shear, the BRDF bandwidth, and the angular scale due to rotation into the tangent frame.

A case study: simple caustics Caustics are an example of the interaction between spatial and angular aspects of light transport. We illustrate this effect with a simple case similar to a solar oven (Fig 11). A parallel beam of light hits a surface of negative curvature with a mirror (Dirac) BRDF and converges toward a focal point. This is modeled in our framework by an incoming spectrum that has energy only in the angular domain. The shear due to curvature followed by the shear due to transport result in a signal where the energy is concentrated in space: it is a Dirac at the focal point.

4.3 Rotation-varying BRDFs

Not all BRDFs can be simplified into a term that depends only on the difference to the mirror direction. For example, the Fresnel term depends on the incoming angle. We now derive the effect of shading by a BRDF that is factored into separable terms that depend on the incoming angle θ'_i and the difference between the outgoing angle θ'_o and the mirror direction θ'_r [Ramamoorthi and Hanrahan 2002], that is, $\rho(\theta'_i, \theta'_o) = f(\theta'_i) \rho'(\theta'_o - \theta'_r)$.

Since the term f does not depend on the outgoing angle, it can be applied in the same way as the \cos_+ term, using a multiplication that corresponds to a convolution in frequency space; the rest of

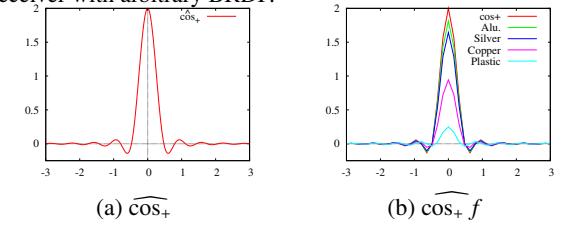
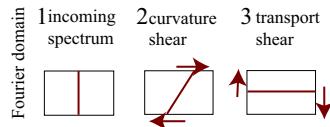


Figure 12: Spectrum of the clamped cosine in the sphere-plane parameterization. Spectrum of $\cos_+ f$ (cosine and Fresnel terms) for different materials.

the shading remains the same with a convolution by ρ' . Combining the multiplication by f with the mirror reparameterization of step 4 and the convolution by ρ' of step 5, we obtain in frequency space a convolution followed by a multiplication:

$$\widehat{\ell}_o(\Omega'_x, \Omega'_\theta) = (\widehat{dE'_r}(\Omega'_x, \Omega'_\theta) \otimes \widehat{f}(-\Omega'_\theta) \delta_{\Omega_x=0}) \widehat{\rho'}(\Omega'_\theta) \quad (17)$$

Fig. 12-b shows the spectra of the cosine term \cos_+ multiplied by the Fresnel term for typical materials; it contains mostly low frequencies. Other approximations with separable functions depending on θ'_o are equally easy, just reversing the order of the multiplication and convolution. BRDFs are often approximated by sums of separable terms, which can be handled easily in our framework because the Fourier transform is linear.

4.4 Texture mapping

When the result of shading is modulated by a texture $T(x)$, this multiplication corresponds to a convolution in the Fourier domain:

$$\widehat{\ell}_T(\Omega_x, \Omega_\theta) = \widehat{T}(\Omega_x) \delta_{\Omega_\theta=0} \otimes \widehat{\ell}_o(\Omega_x, \Omega_\theta) \quad (18)$$

Since the texture has only spatial components, its spectrum is restricted to the line of spatial frequencies. This means that texture mapping only affects frequencies along the spatial dimension.

4.5 Spatially-varying BRDFs

We now extend our model to include spatially-varying BRDFs and revisit step 5 (shading). For each point, shading is still a convolution

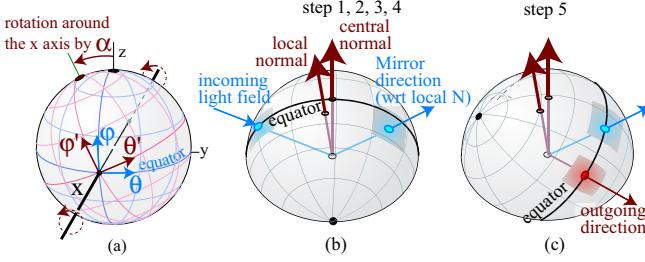


Figure 13: 3D direction parameterizations.

over the directional domain, but the kernel varies spatially.

To model this effect, we exploit the fact that a 2D Fourier transform can be decomposed into two separable 1D transforms, the first one vertically, then horizontally. We consider the intermediate semi-Fourier space $\hat{\ell}(x, \Omega_\theta)$ that represents for each location x the 1D Fourier transform of the directional variation of incoming light. The full Fourier space is then the 1D Fourier transform of the semi-Fourier transform along the x dimension. We have

$$\hat{\ell}_o(x, \Omega_\theta) = \hat{dE}'_r(x, \Omega_\theta) \hat{\rho}(x, \Omega_\theta),$$

which is a multiplication in the semi-Fourier domain, and therefore a convolution *along x only* in full Fourier space:

$$\hat{\ell}(\Omega_x, \Omega_\theta) = \hat{dE}'_r(\Omega_x, \Omega_\theta) \otimes_x \hat{\rho}(\Omega_x, \Omega_\theta)$$

This means that in order to characterize the effect of spatially-varying BRDFs, we consider the spectrum of $\rho(x, \theta)$. We then take the spectrum of the incoming illumination $\hat{\ell}$ and convolve it only horizontally along Ω_x , not vertically. We call this a semi-convolution in Ω_x , which we note \otimes_x .

In the special case of non-varying BRDFs, the spectrum of $\rho(x, \theta)$ is a Dirac times the directional spectrum of the BRDF. The horizontal convolution is a multiplication. If the spectrum of ρ is separable (texture mapping), then the spatially-varying BRDF case is a multiplication followed by a convolution. The special case of a spatially-varying combination of BRDFs [Lensch et al. 2001] can be handled more simply as the superposition of multiple BRDFs with weights encoded as textures.

5 Extension to 3D

We now show how our framework extends to 3D scenes.

5.1 Light-field parameterization phenomena

The derivations presented in Section 3 involve a two-plane light-field parameterization and extend directly to 3D. The only notable difference is the calculation of differential irradiance (Eq. 7), where the projected surface area in 3D becomes:

$$dA_\perp = \frac{du dv}{(1 + v^2 + u^2)^2} = G(u, v) du dv$$

Fig. 5-c presents the spectrum of $G(u, v)$.

5.2 Shading in plane-sphere parameterization

The sphere S^2 of directions is unfortunately hard to parameterize, which prompted many authors to use spherical harmonics as the equivalent of Fourier basis on this domain. In contrast, we have chosen to represent directions using spherical coordinates and to use traditional Fourier analysis, which is permitted by our restriction to local neighborhoods of S^2 . This solution enables a more direct extension of our 2D results, and in particular it expresses well the interaction between the spatial and angular components.

Spherical coordinates We use the spherical coordinates θ, φ where θ , in $[-\pi, \pi]$, is the azimuth and φ , in $[-\pi/2, \pi/2]$, the colatitude. The distortion of this parameterization is $\cos \varphi$, which means that one must remain around the equator to avoid distortion. In this neighborhood, the parameterization is essentially Euclidean, to a first-order approximation.

Local reflection is challenging because it involves four neighborhoods of direction around: the incoming direction, the normal, the mirror direction, and the outgoing direction; in general, we cannot choose a spherical parameterization where they all lie near the equator. Fortunately, we only need to consider two of these neighborhoods at a time (Fig. 4).

For this, we exploit the fact that a rotation around an axis on the equator can be approximated to first order by a Euclidean rotation of the (θ, φ) coordinates: $(\theta', \varphi') = \mathbf{R}_\alpha(\theta, \varphi)$

For brevity, we omit the comprehensive remapping formulas for 3D shading, but we describe the appropriate parameterization for each step as well as the major differences with the 2D case.

Tangent frame We start with a parameterization where the equator is in the incident plane, as defined by the central ray of the incident light field and the central normal vector (Fig. 13-b). If the light field has been properly rotated, only the x spatial dimension undergoes the scaling by $\cos \theta_0$ (Eq. 9)

Curvature In 2D, we approximated the angle with the local normal linearly by $\alpha(x) = kx$; For a surface, the corresponding linearization of the normal direction (θ_N, φ_N) involves a bilinear form [Do Carmo 1976]:

$$(\theta_N, \varphi_N) = \mathbf{M}(x, y) \quad (19)$$

If x and y are aligned with the principal directions of the surface, the matrix \mathbf{M} is an anisotropic scaling where the scale factors are the two principal curvatures. The corresponding remapping of (x, y, θ, φ) is a shear in 4D, with different amounts in the two principal directions. As with the 2D case, the frequency content is sheared along the spatial dimensions.

Differential irradiance and cosine Step 3 is mostly unchanged. Since we placed the equator along the incident plane, the cosine term depends only on θ to a first approximation. The spectrum is convolved with a small 1D kernel in θ (Fig. 12-a).

Rotationally-symmetric BRDFs The mirror reparameterization of step 4 is unchanged, and the angles remain near the equator since the equator also contains the normals. We express the convolution of the mirrored incoming light field by the BRDF in the neighborhood of the outgoing direction. For this, we rotate our spherical coordinates so that the new equator contains both the mirror direction and the direction of the central outgoing ray (Fig. 13-c). Because all the angles are near the equator, the difference angles between an outgoing ray and a mirrored ray can be approximated by $\theta'_o - \theta'_r$ and $\varphi'_o - \varphi'_r$, and Eq. 16 applies.

Recap of rotations In summary, we first need to rotate the light-field parameterization so that the central incidence plane is along one of the axes before reparameterizing from two-plane to sphere-plane (Fig. 13-b). We then need to rotate between the mirror reparameterization and the BRDF convolution to place the central outgoing direction on the equator (Fig. 13-c). Finally we rotate again to put the outgoing plane defined by the normal and central outgoing direction in the equator (not shown).

5.3 Anisotropies in 3D

Our extension to 3D exploits the low distortion of spherical coordinates near the equator, at the cost of additional reparameterization

	Ray space	Fourier	Spectrum formula
Transport			
Travel	shear \square	shear \square	$\widehat{\ell}(\Omega_x, \Omega_\theta + d\Omega_x)$
Visibility	multiplication	convolution	$\widehat{\ell} \otimes \widehat{V}$
Local geometric configuration			
Light incidence	scale spatial	scale spatial	$\frac{e^{-i\Omega_\theta\theta_0}}{ \cos\theta_0 } \widehat{\ell}(-\Omega_x \cos\theta_0, \Omega_\theta)$
Outgoing angle	scale spatial	scale spatial	$e^{i\Omega_\theta\theta_1} \cos\theta_1 \widehat{\ell}_o\left(\frac{\Omega_x}{\cos\theta_1}, \Omega_\theta\right)$
Curvature	shear \square	shear \square	$\widehat{\ell}_i(\Omega'_x - k\Omega'_\theta, \Omega'_\theta)$
Local shading			
Cosine term	multiplication	convolution	$\widehat{\ell}' \otimes \widehat{\cos}$
BRDF	convolution	multiplication	$\rho' \widehat{dE'_r}$
Texture mapping	multiplication	convolution	$\widehat{T} \otimes \widehat{\ell}$
Separable BRDF	multiplication	convolution	$(\widehat{dE'_r} \otimes \widehat{f}) \rho'$
	then convolution	then multiplication	
Space-vary. BRDF	semi-convolution (angles only)	semi convolution (spatial only)	$\widehat{dE'} \otimes_x \widehat{\rho}$

Table 1: Summary of all phenomena

to align the equator with the relevant neighborhoods. Fortunately, these reparameterization act locally like Euclidean rotations along axes that preserve the space-angle separation.

Compared to 2D, the 3D case involves anisotropies both in the directional and spatial components. The *spatial* scale to account for the incident and exitant angle affects only one of the spatial dimensions, along the corresponding plane normal to the tangent frame. Curvature is usually different along the two principal directions. The *directional* cosine term mostly depends on θ , while rotationally-symmetric BRDFs only depend on the spherical distance between mirror and outgoing directions and is more influenced by θ except around the specular peak. These additional anisotropies make the 3D situation more complex, but locally they correspond to linear transforms and preserve the distinction and interaction between spatial and directional effects derived in 2D.

Other local shading effects such as separable BRDFs, texture mapping, and spatially-varying BRDFs can be directly extended from Section 4. While the formulas are complex and are not derived in this paper, the qualitative effects and relevant parameters remain the same as in 2D.

6 Discussion

Table 1 summarizes the building blocks that compose our frequency analysis of light transport. This variety of phenomena can be characterized using simple mathematical operators: scale, shear, convolution and multiplication. Even spatially-varying BRDFs can be handled using a semi-convolution that occurs only along the spatial dimensions.

Some operations such as occlusion are simpler in the original ray space, while others such as shading are more natural in frequency space. Our framework allows us to express them all in a unified way. As discussed above, the 3D case essentially fol-

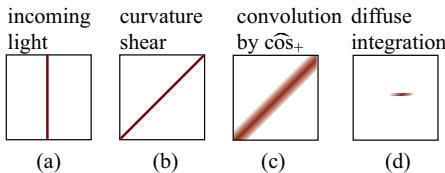


Figure 14: Special case of diffuse shading for an infinite environment map. In this case, the convolution and multiplication are equivalent to a multiplication.

lows the 2D derivation, with additional reparameterizations steps and anisotropies.

In practice, the notion of locality is invoked for three different reasons, whose importance depends on the context and application:

- the use of first-order Taylor series, for example for the curvature or for the $\tan\theta \approx \theta$ remapping,
- the principle of uncertainty, which states that low frequencies cannot be measured on small windows (in which case big neighborhoods are desired), or in other words, that localization is not possible in space and frequency at the same time,
- most real scenes are not stationary, that is, their properties such as the presence and size of blockers vary spatially. Smaller neighborhoods might mean more homogeneous properties and more locally-pertinent conclusions.

We now discuss how our work extends previous frequency characterization, before illustrating how it can be applied, through a proof of concept in the context of ray-tracing.

6.1 Relation to previous work

Light field sampling Our formulation of transport in free space is similar to the derivations of light-field spectra [Isaksen et al. 2000; Chai et al. 2000], and the same relationship between slope and distance is found. Our expression as a transport operator makes it easier to extend these analyzes to arbitrary input signals, and in particular to non-Lambertian objects and occlusion.

Ray footprint Approaches based on ray differentials [Shinya et al. 1987; Igely 1999; Chen and Arvo 2000] capture the shear transforms due to transport and curvature, and our first-order Taylor expansion for curvature corresponds to the same differentials. The approach by Igely [1999] only uses 2D derivatives by considering only ray paths that converge to the viewpoint.

Signal processing framework for local reflection Our framework extends Ramamoorthi and Hanrahan’s signal processing framework for local reflection [2004] with the following key differences:

- we take into account spatial variation of the incoming light and the curvature of the receiver,
- however, we characterize reflection only for a ray neighborhood,
- they parameterize the outgoing radiance by α and θ'_o , while we use a more natural outgoing parameterization in the global frame, at the cost of reparameterization,
- as discussed above, our expression of the cosine term is a convolution in the frequency domain. This cleanly separates the computation of incoming irradiance and BRDF convolution, at the cost of additional steps. It also allows us to express the cosine term for BRDFs such as Phong.

On convolution It might come as a surprise that two phenomena that have been expressed by previous work as convolutions in the primary space, soft shadows [Soler and Sillion 1998] and the cosine term [Ramamoorthi and Hanrahan 2004] correspond in our framework to convolutions *in the frequency domain*. We show here that our formulation in fact extends these previous work and that the primary-space convolution is a special case. The key is that they consider functions that do not vary in one of the domains (space resp. direction). The corresponding spectra are therefore restricted to a line, since the Fourier transform of a constant is a Dirac.

Consider the cosine term for infinitely-distant light sources. The lighting varies only in angle, and its spectrum is restricted to the vertical line of directions (Fig. 14(a)). After the curvature shear, it is a 1D function on the line $k\Omega_\theta = \Omega_x$ (Fig. 14(b)), which we convolve with the vertical kernel $\widehat{\cos}_+$. However, for each spatial frequency Ω_x , there is only one non-zero value of the sheared function. As a result, this convolution is a so-called *outer product* of the two one-dimensional functions, and the result is the pairwise product $\widehat{dE'}(\Omega_x, \Omega_\theta) = \ell'_i(0, \Omega_x/k) \widehat{\cos}_+(\Omega_x/k - \Omega_\theta)$. (Fig. 14(c)). The diffuse integration then restricts the function to the line $\Omega_\theta = 0$, where $\widehat{dE'}(\Omega_x, 0) = \ell'_i(0, \Omega_x/k) \widehat{\cos}_+(\Omega_x/k)$. The convolution followed by

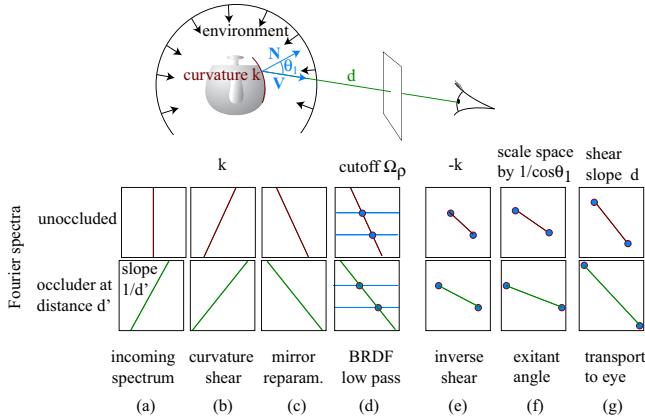


Figure 15: Bandwidth derivation for our adaptive ray tracer.

restriction to the horizontal line turned into a simple product of 1D functions, which corresponds to a convolution in the *primary* space.

The case of soft shadows studied by Soler and Sillion [1998] is similar: the emitter is diffuse and has a spectrum restricted to $\Omega_v = 0$, and the blockers are planar and have the same restrictions. The transport from source to blockers results in slanted lines in frequency space that are convolved together (Fig. 6, step 3). Our framework extends these two cases to arbitrary cases where the spectra are not restricted to lines.

6.2 Sampling rate for glossy rendering

We show how our framework can be used to drive image-space sampling in ray tracing. In particular, we illustrate how our framework can be used to derive sampling rates for algorithms that *do not need to perform computations in the Fourier domain*. While we demonstrate a working implementation, we emphasize that this application is meant only as a proof of concept and that further development is necessary to make it fully general, as we discuss below.

We are motivated by the rendering of glossy materials, which despite effective recent developments [Lawrence et al. 2004] remains computationally expensive. We observe, however, that glossy objects appear blurry, so it should be possible to reduce the image sampling rate. Our framework permits a quantitative expression of the required sampling rate.

Unoccluded environment map We first consider the case of environment-map rendering without occlusion. The incoming light field has only directional content (Fig. 15), and the light incidence angle (Table 1 row 1) has no effect. The shear of curvature (b) results in a line of slope k that gets convolved with the cosine narrow kernel $\widehat{\cos_+}$, which we neglect. After mirror reparameterization (c), a glossy BRDF band-limits this signal (d), which we approximate by a cutoff at an angular frequency Ω_p . This cutoff depends on the BRDF of the object visible in a given region of the image. The upper endpoint of the resulting segment is at coordinate $(-k\Omega_p, \Omega_p)$. We apply the inverse shear (step e) and the scale by $1/\cos\theta_1 = 1/(\mathbf{n} \cdot \mathbf{v})$, where \mathbf{n} is the normal and \mathbf{v} the unit vector towards the viewpoint. We obtain a maximum frequency of $(-\frac{2k}{\cos\theta_1}\Omega_p, \Omega_p)$ for the light leaving an object in the direction of the viewpoint (Fig. 15 step f). A transport shear with distance d yields a bound of $(-\frac{2k}{\cos\theta_1}\Omega_p, \Omega_p - d\frac{2k}{\cos\theta_1}\Omega_p)$

A view computation corresponds to the restriction of the function to the directional domain, and if we assume $d \gg 1$, we obtain the following approximate bound on the directional bandwidth for a region of the image:

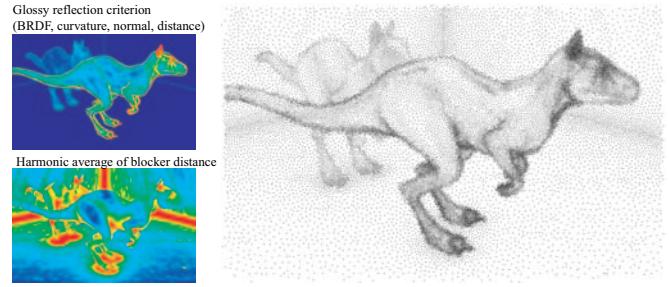


Figure 16: Criteria and sampling pattern used to render Fig. 17. The sampling adapts to curvature, the viewing angle, the BRDF as well as the harmonic average of the distance to potential blockers.

$$B = d \frac{2k}{\mathbf{n} \cdot \mathbf{v}} \Omega_p \quad (20)$$

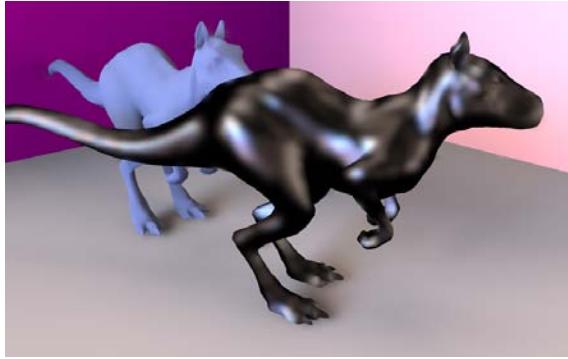
This corresponds to the difficulty in appropriately sampling curved objects at grazing angles, as discussed in Section 4.2. In addition, distant objects are minified and the apparent curvature is increased. In 3D, the curvature and normal angle involve anisotropy, and in practice, we use the absolute value of the larger principal curvature. However, our implementation computes this criterion directly in screen-space with finite-difference approximation to the curvature. As a result, the effect of the normal angle, the distance, and the anisotropies are included for free; see Fig. 16 for a visualization of this criterion. The faceted look is due to the finite-difference curvature and Phong normal interpolation. The BRDF bandwidth for the two dinosaurs and environments were approximated manually based on the BRDF exponent.

Occlusion The above derivation assumes full visibility. We integrate the effect of the blockers using a worst-case assumption on the blocker spectrum, we consider that it has content at all frequency. Based on the effect of the transport shear, we approximate the spectrum due to a blocker at distance d' by a line of slope $1/d'$. Going through the same steps, we obtain an approximate bound of:

$$B' = d \frac{1}{\mathbf{n} \cdot \mathbf{v}} \left(\frac{1}{d'} + 2k \right) \Omega_p \quad (21)$$

To evaluate this criterion, we use the harmonic average of the distance to occluders. This information is computed by sampling a hundred rays for a small set of visible points in the image, in practice 20,000. The criteria is reconstructed over the image using the same reconstruction algorithm as for the final image, which we describe shortly. The blocker criterion is shown Fig. 15. It is similar to Ward et al.’s criterion for irradiance caching [Ward et al. 1988], but expressing it in a unified frequency framework allows us to combine it with other bandwidth considerations such as BRDF roughness.

Algorithm and image reconstruction Our proof-of-concept computes visibility using four samples per pixel, but uses aggressively-sparse samples for shading: on average, 0.05 samples per pixel. We use an edge-preserving reconstruction that exploits the higher-resolution depth and normal to reconstruct the shading, in the spirit of McCool’s filtering of Monte-Carlo ray tracing outputs [1999] but based on a bilateral filter [Tomasi and Manduchi 1998]. As demonstrated in Fig. 17, this results in a smooth reconstruction where needed and on sharp silhouettes. The spatial width of the bilateral filter is scaled according to the bandwidth prediction. Given a bandwidth map, we use the blue-noise method by Ostromoukhov et al. [2004] to generate a set of image samples (Fig. 17, right). In summary, our algorithm is as follows:



Uniform sampling



Using our bandwidth prediction

Figure 17: Scene rendered without and with adaptive sampling rate based on our prediction of frequency content. Only 20,000 shading samples were used to compute these 800×500 image. Note how our approach better captures the sharp detail in the shiny dinosaur’s head and feet. The criteria and sampling are shown in Fig. 16. Images rendered using PBRT [Pharr and Humphreys 2004]

```

Compute visibility at full resolution
  Use finite-differences for curvature criterion
Compute harmonic blocker distance for sparse samples
  Perform bilateral reconstruction
Compute B' based on blocker and curvature
Generate blue noise sampling based on B'
  Compute shading for samples
  Perform bilateral reconstruction

```

Observe how our sampling density is increased in areas of high curvature, grazing angles, and near occluders. The environment map casts particularly soft shadows, and note how the high-frequency detail on the nose of the foreground dinosaur is well captured, especially given that the shading sampling is equivalent to a 200×100 resolution image.

Although these results are encouraging, the approach needs improvement in several areas. The visibility criterion in particular should take into account the light source intensity in a directional neighborhood to better weight the inverse distances. Even so, the simple method outlined above illustrates how knowledge of the modifications of the frequency content through light transport can be exploited to drive rendering algorithms. In particular, similar derivations are promising for precomputed radiance transfer [Sloan et al. 2002] in order to relate spatial and angular sampling.

7 Conclusions and future work

We have presented a comprehensive framework for the description of radiance in frequency space, through operations of light transport. By studying the local light field around the direction of propagation, we can characterize the effect of travel in free space, occlusion, and reflection in terms of frequency content both in space and angle. In addition to the theoretical insight offered by our analysis, we have shown that practical conclusions can be drawn from a frequency analysis, without explicitly computing any Fourier transforms, by driving the sampling density of a ray tracer according to frequency predictions.

Future work On the theory side, we are working on the analysis of additional local shading effects such as refraction, bump-mapping, and local shadowing [Ramamoorthi et al. 2004]. We hope to study the frequency cutoff for micro, meso, and macro-geometry effects [Becker and Max 1993]. The study of participating media is promising given the ability of Fourier analysis to model differential equations. The spectral analysis of light interaction in a full scene is another challenging topic. Finally, the addition of the time dimension is a natural way to tackle effects such as motion blur.

We are excited by the wealth of potential applications encompassed by our framework. In rendering, we believe that many traditional algorithms can be cast in this framework in order to derive

theoretical justification, but also to allow extensions to more general cases (such as from diffuse to glossy). Our preliminary study of sampling rates in ray tracing is promising, and we want to develop new algorithms and data structures to predict local bandwidth, especially for occlusion effects. Precomputed radiance transfer is another direct application of our work.

Our analysis extends previous work in inverse rendering [Ramamoorthi and Hanrahan 2001b; Basri and Jacobs 2003] and we are working on applications to inverse rendering with close-range sources, shape from reflection, and depth from defocus.

Acknowledgments

We thank Jaakko Lehtinen, the reviewers of the Artis and MIT teams, as well as the SIGGRAPH reviewers for insightful feedback.

This work was supported by an NSF CAREER award 0447561 “Transient Signal Processing for Realistic Imagery,” an NSF CISE Research Infrastructure Award (EIA9802220), an ASEE National Defense Science and Engineering Graduate fellowship, the European Union IST-2001-34744 “RealReflect” project, an INRIA équipe associée, and the MIT-France program.

References

- ANNEN, T., KAUTZ, J., DURAND, F., AND SEIDEL, H.-P. 2004. Spherical harmonic gradients for mid-range illumination. In *Rendering Techniques 2004 (Proc. EG Symposium on Rendering 2004)*.
- ARVO, J. 1994. The irradiance Jacobian for partially occluded polyhedral sources. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*, 343–350.
- BASRI, R., AND JACOBS, D. 2003. Lambertian reflectance and linear subspaces. *IEEE Trans. Pattern Anal. Mach. Intell.* 25, 2.
- BECKER, B. G., AND MAX, N. L. 1993. Smooth transitions between bump rendering algorithms. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*, 183–190.
- BOLIN, M. R., AND MEYER, G. W. 1995. A frequency based ray tracer. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*, 409–418.
- BOLIN, M. R., AND MEYER, G. W. 1998. A perceptually based adaptive sampling algorithm. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*, 299–309.
- CAMAHORT, E., LERIOS, A., AND FUSSELL, D. 1998. Uniformly sampled light fields. In *Rendering Techniques '98 (Proc. of EG Workshop on Rendering '98)*, Eurographics, 117–130.

CHAI, J.-X., CHAN, S.-C., SHUM, H.-Y., AND TONG, X. 2000. Plenoptic sampling. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*, 307–318.

CHEN, M., AND ARVO, J. 2000. Theory and application of specular path perturbation. *ACM Trans. Graph.* 19, 4, 246–278.

DO CARMO, M. 1976. *Differential Geometry of Curves and Surfaces*. Prentice Hall.

FERWERDA, J. A., SHIRLEY, P., PATTANAIK, S. N., AND GREENBERG, D. P. 1997. A model of visual masking for computer graphics. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*, 143–152.

FROLOVA, D., SIMAKOV, D., AND BASRI, R. 2004. Accuracy of spherical harmonic approximations for images of Lambertian objects under far and near lighting. In *ECCV 2004, European Conference on Computer Vision*, 574–587.

GOODMAN, J. W. 1996. *Introduction To Fourier Optics*. McGraw-Hill.

GORTLER, S. J., SCHRÖDER, P., COHEN, M. F., AND HANRAHAN, P. 1993. Wavelet radiosity. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*, 221–230.

HALLE, M. 1994. Holographic stereograms as discrete imaging systems. In *SPIE Proc. Vol. 2176: Practical Holography VIII*, S. Benton, Ed., SPIE, 73–84.

HECKBERT, P. 1989. *Fundamentals of Texture Mapping and Image Warping*. Master's thesis, University of California at Berkeley, Computer Science Division.

HOLZSCHUCH, N., AND SILLION, F. X. 1998. An exhaustive error-bounding algorithm for hierarchical radiosity. *Computer Graphics Forum* 17, 4.

IGEHY, H. 1999. Tracing ray differentials. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*.

ISAKSEN, A., McMILLAN, L., AND GORTLER, S. J. 2000. Dynamically reparameterized light fields. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*, 297–306.

KELLER, A. 2001. Hierarchical monte carlo image synthesis. *Mathematics and Computers in Simulation* 55, 1–3 (Feb.), 79–92.

LAWRENCE, J., RUSINKIEWICZ, S., AND RAMAMOORTHI, R. 2004. Efficient BRDF importance sampling using a factored representation. *ACM Transactions on Graphics (Proc. SIGGRAPH 2004)* 23, 3 (Aug.), 496–505.

LENSCH, H. P. A., KAUTZ, J., GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2001. Image-based reconstruction of spatially varying materials. In *Rendering Techniques '01 (Proc. EG Workshop on Rendering 2001)*, Eurographics, 104–115.

MALIK, J., AND ROSENHOLTZ, R. 1997. Computing local surface orientation and shape from texture for curved surfaces. *International Journal of Computer Vision* 23, 2, 149–168.

MC COOL, M. D. 1999. Anisotropic diffusion for monte carlo noise reduction. *ACM Transactions on Graphics* 18, 2, 171–194.

MYSZKOWSKI, K. 1998. The visible differences predictor: applications to global illumination problems. In *Rendering Techniques '98 (Proc. EG Workshop on Rendering '98)*, Eurographics.

OSTROMOUKHOV, V., DONOHUE, C., AND JODOIN, P.-M. 2004. Fast hierarchical importance sampling with blue noise properties. *ACM Transactions on Graphics (Proc. SIGGRAPH 2004)* 23, 3 (Aug.), 488–495.

PENTLAND, A. P. 1987. A new sense for depth of field. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 9, 4 (July).

PHARR, M., AND HUMPHREYS, G. 2004. *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann.

RAMAMOORTHI, R., AND HANRAHAN, P. 2001. An efficient representation for irradiance environment maps. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*.

RAMAMOORTHI, R., AND HANRAHAN, P. 2001. A signal-processing framework for inverse rendering. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*.

RAMAMOORTHI, R., AND HANRAHAN, P. 2002. Frequency space environment map rendering. *ACM Transactions on Graphics (Proc. SIGGRAPH 2002)* 21, 3, 517–526.

RAMAMOORTHI, R., AND HANRAHAN, P. 2004. A signal-processing framework for reflection. *ACM Transactions on Graphics* 23, 4.

RAMAMOORTHI, R., KOUDELKA, M., AND BELHUMEUR, P. 2004. A Fourier theory for cast shadows. In *ECCV 2004, European Conference on Computer Vision*, 146–162.

SHINYA, M., TAKAHASHI, T., AND NAITO, S. 1987. Principles and applications of pencil tracing. *Computer Graphics (Proc. SIGGRAPH '87)* 21, 4, 45–54.

SILLION, F., AND DRETTAKIS, G. 1995. Feature-based control of visibility error: A multi-resolution clustering algorithm for global illumination. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*, 145–152.

SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Trans. on Graphics* 21, 3, 527–536.

SOLER, C., AND SILLION, F. X. 1998. Fast calculation of soft shadow textures using convolution. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*, 321–332.

STEWART, J., YU, J., GORTLER, S. J., AND McMILLAN, L. 2003. A new reconstruction filter for undersampled light fields. In *Proc. EG Symposium on Rendering 2003*, Eurographics, 150–156.

TOMASI, C., AND MANDUCHI, R. 1998. Bilateral filtering for gray and color images. In *Proc. IEEE International Conference on Computer Vision*, IEEE, 836–846.

WARD, G. J., AND HECKBERT, P. 1992. Irradiance gradients. In *Proc. of EG Workshop on Rendering '92*, Eurographics, 85–98.

WARD, G. J., RUBINSTEIN, F. M., AND CLEAR, R. D. 1988. A ray tracing solution for diffuse interreflection. *Computer Graphics (Proc. SIGGRAPH '88)* 22, 4 (Aug.), 85 – 92.

WOOD, D. N., AZUMA, D. I., ALDINGER, K., CURLESS, B., DUCHAMP, T., SALESIN, D. H., AND STUETZLE, W. 2000. Surface light fields for 3D photography. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*, 287–296.