A Precomputed Polynomial Representation for Interactive BRDF Editing with Global Illumination

ANER BEN-ARTZI, KEVIN EGAN, and RAVI RAMAMOORTHI
Columbia University
and
FRÉDO DURAND
MIT CSAIL

The ability to interactively edit BRDFs in their final placement within a computer graphics scene is vital to making informed choices for material properties. We significantly extend previous work on BRDF editing for static scenes (with fixed lighting and view) by developing a precomputed polynomial representation that enables interactive BRDF editing with global illumination. Unlike previous precomputation-based rendering techniques, the image is not linear in the BRDF when considering interreflections. We introduce a framework for precomputing a multibounce tensor of polynomial coefficients that encapsulates the nonlinear nature of the task. Significant reductions in complexity are achieved by leveraging the low-frequency nature of indirect light. We use a high-quality representation for the BRDFs at the first bounce from the eye and lower-frequency (often diffuse) versions for further bounces. This approximation correctly captures the general global illumination in a scene, including color-bleeding, near-field object reflections, and even caustics. We adapt Monte Carlo path tracing for precomputing the tensor of coefficients for BRDF basis functions. At runtime, the high-dimensional tensors can be reduced to a simple dot product at each pixel for rendering. We present a number of examples of editing BRDFs in complex scenes with interactive feedback rendered with global illumination.

1. INTRODUCTION

Recent advances in real-time rendering have improved the ability of designers to interactively specify lighting and materials in computer graphics scenes. While relighting systems have long provided feedback with global illumination in complex scenes [Dorsey et al. 1995], BRDF editing has been limited to simplified settings such as point lights.

Recently, Ben-Artzi et al. [2006] have introduced the ability to edit BRDFs under natural illumination, albeit only with direct lighting. This is a significant limitation since indirect illumination and glossy reflection are essential to the realism of today's renderers and are often critical to correctly perceive and choose material properties.

In this article, we develop a precomputation-based method for interactive BRDF editing with global illumination (see results in Figure 1). The main challenge arises from the fact that final scene radiance (an image) is not even linear in the objects' BRDFs. It is well known that albedo, or more generally a BRDF, has a nonlinear effect because it multiplies the light at each bounce. We develop a higher-order representation of the image as a function of the scene's BRDFs. We precompute a tensor at each pixel, fixing the lighting and view for a static scene but leaving the BRDFs unspecified until runtime when they can be edited.

Our first contribution in Section 3 is a general theoretical framework for BRDF editing, based on a bilinear formulation of the reflection operator, that extends the linear operator formulation of rendering [Arvo et al. 1994]. We show how the precomputed matrix of previous methods must be extended to a multibounce tensor of polynomial coefficients.

The full multibounce tensor is a complete representation of the image as a function of scene BRDFs but is computationally too expensive to treat in full generality. We consider frequency characteristics, developing a tractable approximation that preserves most important

Fig. 1. We simulate an interior design session in which we edit the BRDFs of the couch and floor. The couch's red fabric in (a) is loaded from measured data and edited to a more specular green material in (b). The floor is retextured and made glossy in (b). The reflections of objects in the glossy floor and color bleeding from the couch to the staircase can be seen here as well as in closeups in Figure 10. The scene is lit only from the large windows, using the top half of an exterior environment map (campus). We can see in the comparisons to direct lighting (far right) that most of the room's lighting is a result of indirect illumination.

perceptual effects (Section 4). Specifically, the first bounce from the eye usually uses the full BRDF to capture glossy reflections, while subsequent bounces use a lower-frequency approximation to capture overall shading effects. Within this general framework, we show two possibilities: where further bounces (up to the fourth-bounce) are treated as purely diffuse (Figures 1, 6 and 9), and where additionally, the second bounce from the eye uses a lower-frequency approximation to achieve accurate indirect reflections in glossy surfaces (Figure 4), or even intricate effects like caustics (Figure 11).

For precomputation (Section 5.1), we show how Monte Carlo path tracing can be extended to precompute the multibounce tensors needed. For rendering (Section 5.2), since only one object's BRDF is edited at a time, we show that the tensor can be reduced to a few vector dot products at each pixel whose coefficients can be computed in a very fast runtime preprocess. Our results show a variety of BRDF edits with interactive feedback rendered with global illumination.

2. PREVIOUS WORK

Precomputed Radiance Transfer (PRT). We build on PRT ideas for static scenes [Sloan et al. 2002; Ng et al. 2003]. While these methods focus on lighting design with fixed BRDFs, we focus on BRDF editing with fixed lighting. We are inspired by a body of recent work that underscores the importance of interreflections in relighting [Hasan et al. 2006; Wang et al. 2006; Kontkanen et al. 2006]. All these approaches exploit the linearity of relighting with respect to light intensity, even when global illumination is taken into account. In contrast, BRDF editing is fundamentally nonlinear when global illumination is considered.

Most previous PRT methods precompute a linear light transport vector at each image location, taking advantage of the linearity of light. We extend this concept to a general tensor of coefficients for a high-dimensional polynomial. The idea of going from linear to quadratic or cubic precomputed models has also begun to be explored in physical simulation [Barbic and James 2005] but in the context of differential equations and model dimensionality reduction. In the context of real-time rendering, Sun and Mukherjee [2006] precompute with a larger product of functions, leading to an N-part multiplication at runtime. Each function is still precomputed independently, and the runtime calculations are still linear in any of the individual functions.

While some PRT methods allow for all-frequency effects and view changes [Wang et al. 2006], BRDF editing cannot take advantage of such factorization-based approaches since the BRDF lobe at a pixel is defined over both the dimensions of lighting (ω_i) and view (ω_o) simultaneously and also depends on the surface normal. Therefore, we fix the lighting, view and geometry, but could, in principle, allow a small number of predefined views or lighting conditions to be updated simultaneously (as in Ben-Artzi et al. [2006]).

Global Illumination. Our precomputation method is inspired by offline global illumination algorithms such as Monte Carlo path tracing [Kajiya 1986]. We have also been able to adapt finite element radiosity [Cohen and Wallace 1993], although we found meshing and complexity issues difficult to deal with for our complex scenes and do not discuss it further. Global illumination techniques usually require the BRDF to be fixed, and use it for importance sampling or hierarchy construction. We develop extensions that are independent of the BRDF and allow real-time editing. In effect, we precompute a *symbolic representation* of the image, which can be evaluated at runtime with polynomials in the user-specified BRDF values to obtain the final intensity.

Séquin and Smyrl [1989] also precompute a symbolic representation of the final image for recursive ray tracing but not full global illumination. Phong shading can be evaluated at runtime, allowing later changes to surface parameters, while reflected and refracted contributions are handled with pointers to subexpressions. In contrast, we seek to simulate complex lighting and full global illumination with many possible illumination paths. Therefore, we cannot afford to store or sum all subexpressions. Instead, we show that the final symbolic expression is a polynomial and only precompute its coefficients. We also allow editing of general parametric and measured BRDFs.

BRDF Representations and Editing. Recent work in BRDF editing has begun to allow edits while rendering important visual effects such as environment maps [Colbert et al. 2006] and general complex lighting with cast shadows [Ben-Artzi et al. 2006]. However, they are limited to direct lighting, which neglects many aspects of appearance in realistic settings with global illumination. We utilize existing BRDF representations, giving the user the ability to edit them interactively in a scene with complex lighting, shadows, and interreflections. Our method supports analytic models like Blinn-Phong or Cook-Torrance, measured half-angle distributions [Ashikhmin

Fig. 2. Schematic of the rendering equation. Outgoing radiance from all surfaces in the scene (left) is converted by the **G** operator into a local incident radiance field at each location (middle), which is then reflected via the bilinear reflection operator **K**, that takes as input both the incident lighting and the BRDF.

et al. 2000; Ngan et al. 2005], and variants of data-driven factored or curve-based models [McCool et al. 2001; Lawrence et al. 2006]. We allow users to either edit values corresponding to parameters in standard analytic BRDFs, or 1D curves for measured data.

Interactive Relighting with Dynamic BRDFs. In simultaneous work Sun et al. [2007] have concurrently developed a precomputed approach to interactive relighting with dynamic BRDFs. Both papers observe that the image is nonlinear in the BRDF, a higher-order polynomial representation is needed for indirect light, and that higher bounces can be computed with coarser representations. Sun et al. [2007] enable changing view by limiting themselves to only two interreflection bounces and through the use of an in-out BRDF factorization that is accurate only for very low-frequency materials. Further approximations are used for glossy effects and indirect lighting. The BRDF space also needs to be specified in advance to allow for a tensor decomposition. By contrast, we focus on visually accurate results for BRDF editing. We are limited to fixed view and lighting, but by allowing greater flexibility in terms of the BRDFs (including very specular materials), we can represent and edit continuously. As seen in our results, we can also support three or four bounce interreflections, including high-frequency caustic effects (Figure 11). Other significant contributions include the general theoretical framework for BRDF editing in Section 3, and the discussion of complexity and low-frequency BRDF approximations (Section 4).

3. GENERAL THEORETICAL FRAMEWORK

This section introduces a general theoretical framework for BRDF editing with global illumination, independent of any specific implementation. It builds on the geometric and reflection operators introduced by Arvo et al. [1994] and shows how BRDF editing can be formulated in terms of a new bilinear reflection operator **K**. The operator notation is compact and alleviates the need for complex integrals in the equations. In Section 4, we discuss our approximations within this framework that make the computation tractable.

3.1 Basic Framework Using the Bilinear **K** Operator

The rendering equation [Kajiya 1986] can be written as a linear operator equation [Arvo et al. 1994]. We build on this foundation and write the rendering equation as:

$$B = E + \mathbf{K}(R)\,\mathbf{G}B, \qquad (1)$$

where $B(x, \omega_o)$ is the outgoing surface radiance, $E(x, \omega_o)$ is the emission, and **G** is the linear geometric operator that converts the outgoing radiance from distant surfaces to a local incident radiance field as in Arvo et al. [1994]. Figure 2 shows a schematic, and Table I summarizes notation.

Table I. Table of Notation

$B(x, \omega_o)$	Outgoing radiance (image)
$E(x, \omega_o)$	Emissive radiance of light sources
$L(x, \omega_i)$	Local incident radiance
$R(x, \omega_i, \omega_o)$	BRDFs of all points in the scene
$T^N(x, \omega_o)$	Precomputed multi-bounce tensor
$\rho^m(\omega_i, \omega_o)$	BRDF of object m
$b_j^m(\omega_i, \omega_o)$	Basis function j for the BRDF of object m
$H^m(x)$	Spatial weight map or texture for object m
G	Linear geometric operator
	$\mathbf{G} : B(x, \omega_o) \mapsto L(y, \omega_i)$
	$(\mathbf{G}B)(x, \omega) \equiv B(x'(x, \omega), -\omega)$
$\mathbf{K}(R)$	Reflection operator (equation 2)
c_j^m	BRDF coefficients (equation 4)
d_m	Equivalent albedo of object m (appendix B)
$\overline{d}_z$	Product of albedos d_m
$\vec{X}_i$	Light path with contribution $f(\vec{X}_i)$
F_{jn}^i	Tensor coefficient after freezing BRDFs
J	Number of BRDF bases (usually 64 or 128)
M	Number of objects ($M \sim 5$)
W	Total basis functions ($W = JM \sim 500$)
Z	Number of terms for diffuse $\overline{d}_z$ ($Z \sim 64$)

Arvo et al. [1994] define **K** as a linear reflection operator on the local incident radiance field L. In our first departure from previous representations, we make explicit **K**'s dependence on the BRDFs R of scene objects. We write **K** as a bilinear operator that takes an additional input $R(x, \omega_i, \omega_o)$, which describes the BRDF at each point in the scene and is the kernel of integration in **K**,

$$\mathbf{K} \;:\; L(x, \omega_i),\, R(x, \omega_i, \omega_o) \mapsto B(x, \omega_o)$$

$$(\mathbf{K}(R)\,L)(x, \omega_o) \equiv \int_{\Omega_{2\pi}} R(x, \omega_i, \omega_o)L(x, \omega_i)\cos\theta_i\, d\omega_i. \quad (2)$$

Note that **K** is bilinear or linear with respect to both inputs, incident lighting L, and the BRDFs of objects in the scene R. That is, for scalars a and b,

$$\mathbf{K}(aR_1 + bR_2)\,L = a\mathbf{K}(R_1)\,L + b\mathbf{K}(R_2)\,L$$
$$\mathbf{K}(R)\,(aL_1 + bL_2) = a\mathbf{K}(R)\,L_1 + b\mathbf{K}(R)\,L_2. \quad (3)$$

We now seek to relate R (and hence **K**) to editable BRDFs of individual objects. We assume there are M objects in the scene and, for now, that each object has a single BRDF. Let object m have[1] BRDF ρ^m. We assume the BRDF can be represented as a linear combination of functions over the domain (ω_i, ω_o),

$$\rho^m(\omega_i, \omega_o) = \sum_{j=1}^{J} c_j^m b_j^m(\omega_i, \omega_o). \quad (4)$$

The BRDF basis functions b could be spherical harmonics, wavelets, or any other linear basis. We follow previous BRDF editing methods [Ben-Artzi et al. 2006; Lawrence et al. 2006] that have used box functions over a suitable 1D parameterization such as the half-angle, as described in Appendix A. They have shown that a 1D parameterization is appropriate for most BRDF edits as well as being compatible with parametric edits of most common BRDF models.

Our goal is to use these basis BRDFs to create a method that allows us to alter the kernel of integration in the **K** operator by specifying

---

[1] We use superscripts to denote some property of the function and parentheses for explicitly raising to a power, so ρ^2 is the second in a series, whereas $(\rho)^2$ is ρ-squared.

Fig. 3.   The final value of each pixel is a polynomial in the BRDF coefficients. Here we show an example with 2 surfaces and 2 basis BRDFs shown in yellow (diffuse and specular). Note that the BRDFs we use in practice are different. The combinatorics of multivariate polynomial coefficients make a tensor notation particularly useful.

different weights c_j^m. We first need to use the b_j^ms to describe R over all surfaces. In order to encode per-object BRDFs, we define a surface mask $H^m(x)$ that is 1 if x is on object m, and 0 otherwise.[2]

$$R(x, \omega_i, \omega_o) = \sum_{m=1}^{M} H^m \rho^m = \sum_{m=1}^{M} H^m(x) \sum_{j=1}^{J} c_j^m b_j^m(\omega_i, \omega_o) \quad (5)$$

The super/subscripts in this equation implicitly define basis functions for the full spatially varying R,

$$R = \sum_{m=1}^{M} \sum_{j=1}^{J} c_j^m R_j^m \; ; \; R_j^m(x, \omega_i, \omega_o) = H^m(x) b_j^m(\omega_i, \omega_o). \quad (6)$$

For simplicity, we will often use a single index w (or u or v) to refer to the double script $_j^m$, with $w \in [1, W] : W = MJ$.

3.2   Polynomial Representation for Multibounce

The solution of Equation (1) can be expressed as the expansion

$$B = E + \mathbf{K}(R)\,\mathbf{G}E + \mathbf{K}(R)\,\mathbf{G}\mathbf{K}(R)\,\mathbf{G}E + \dots, \quad (7)$$

where each term N has an intuitive interpretation, as corresponding to N bounces of light from source(s) to viewer.

All current relighting methods rely on the linearity of B with respect to E. Previous BRDF editing methods also take advantage of the linearity of the 1-bounce term in B (i.e., $\mathbf{K}(R)\,\mathbf{G}E$) with respect to $\mathbf{K}$ (and hence with respect to R), requiring them to render using only direct lighting.

However, this linearity no longer applies for BRDF editing with global illumination because the operator $\mathbf{K}(R)$ is applied multiple times. Even for 2-bounce reflections, the system becomes quadratic and must be represented with a quadratic multivariable polynomial in the c_ws. The N-bounce solution is an order N polynomial. We now show how to extend the general PRT approach to these polynomial equations. We start by considering 2-bounce reflection,

$$B^2 \; = \; \mathbf{K}(R)\,\mathbf{G}\mathbf{K}(R)\,\mathbf{G}E \quad (8)$$

$$= \; \mathbf{K}\left(\sum_u c_u R_u\right)\mathbf{G}\mathbf{K}\left(\sum_v c_v R_v\right)\mathbf{G}E \quad (9)$$

$$= \; \sum_u c_u \left(\mathbf{K}(R_u)\,\mathbf{G}\sum_v c_v \left(\mathbf{K}(R_v)\,\mathbf{G}E\right)\right) \quad (10)$$

$$= \; \sum_u \sum_v c_u c_v \mathbf{K}(R_u)\,\mathbf{G}\mathbf{K}(R_v)\,\mathbf{G}E. \quad (11)$$

The bilinear nature of $\mathbf{K}$ is crucial here. We use the linearity of $\mathbf{K}$ with respect to the BRDF to get Equation (10) and the linearity of $\mathbf{K}$ and $\mathbf{G}$ with respect to radiance to get Equation (11).

---

[2]H can also take on nonbinary values to encode spatial weight maps for combining BRDFs [Lensch et al. 2001; Lawrence et al. 2006], and/or for describing textures.

For BRDF editing, the coefficients (c_u and c_v) become the variables of our computation. The fixed quantities for precomputation are the basis BRDF distributions (R_u and R_v) that depend only on the parameterizations of the various BRDFs. $\mathbf{G}$ and E are also known, defined by the geometry and lighting of the scene, respectively. We precompute a 2-bounce transport function T_{uv}^2 and calculate B^2 as

$$T_{uv}^2 \; = \; \mathbf{K}(R_u)\,\mathbf{G}\mathbf{K}(R_v)\,\mathbf{G}E$$
$$B^2 \; = \; \sum_u \sum_v T_{uv}^2 c_u c_v. \quad (12)$$

Most generally, T_{uv}^2 and B^2 are defined over all spatial locations and view directions. In practice, since we fix the view, $T_{uv}^2(x, \omega_o(x))$ is an order 2 tensor (i.e., a matrix for 2-bounce reflection) at each pixel. $B^2(x, \omega_o(x))$ at each pixel is a quadratic polynomial in the variables c, with coefficients given by T^2. When evaluated, it becomes the radiance function for the scene due to light that has reflected off exactly 2 surfaces. More explicitly,

$$B^2 = T_{11}(c_1)^2 + T_{12}c_1 c_2 + \dots T_{uv}c_u c_v + \dots T_{ww}(c_w)^2. \quad (13)$$

Figure 3 illustrates such polynomials for paths of length 1 and 2. All 1- and 2-term combinations of the BRDFs are possible, including repetitions of the same basis function, as in $T_{11}(c_1^1)^2$, since concave objects can reflect onto themselves. Following the same logic, the N-bounce energy is

$$B^N = \sum_{w_N} \sum_{w_{N-1}} \dots \sum_{w_1} T_{w_N w_{N-1} \dots w_1}^N c_{w_N} c_{w_{N-1}} \dots c_{w_1}, \quad (14)$$

where T^N is an order N tensor for each pixel whose size varies with the number of objects and BRDF basis functions per object. We are evaluating a multivariable, degree N polynomial where each w runs over all W values (all basis functions). The variables of this polynomial are the unknown scalar c's. The coefficients are stored in T^N,

$$T_{w_N w_{N-1} \dots w_1}^N = \mathbf{K}(R_{w_N})\,\mathbf{G}\mathbf{K}(R_{w_{N-1}})\mathbf{G}...\mathbf{K}(R_{w_1})\mathbf{G}E. \quad (15)$$

Finally, we construct the image to be displayed by adding the contributions from the different path-lengths:

$$B = E + B^1 + B^2 + \dots + B^{\overline{N}}, \quad (16)$$

where we cut off the expansion in Equation (7) to $\overline{N} + 1$ terms.

4.   MANAGING COMPLEXITY

Section 3 has presented a theoretical framework that is fully general. However, we need to manage complexity in order to develop a tractable algorithm. In particular, the number of terms in Equation (14) is $(W)^N$. Recall that W is already JM, the number of basis BRDFs times the number of objects. For typical values of M (5) and J (64–128), (so $W \sim 500$) the computational and storage

Fig. 4. An evaluation of the accuracy of different two-bounce decreasing-frequency BRDF series, precomputed and rendered in our editing system. (a) Full resolution BRDFs rendered offline in PBRT, hence a series of (∞, ∞). Our method using different series: (b) (128, 8), (c) (128, 8), (d) (16, 8), (e) (128, 1) or diffuse for second bounce from the eye, (f) direct lighting only (Ben-Artzi et al. 2006). We see that a very low-frequency second-bounce BRDF approximation (128, 8) in (b) and (c) is essentially exact. Moreover, even a diffuse second-bounce approximation (128, 1) in (e) provides the correct shiny material perception of the tray and indirect reflections of tray and teapot. By contrast, direct lighting shows only a black shadow of the teapot on the shiny tray, and the spout and handle do not reflect in the teapot.


Fig. 5. Rate of growth for the precomputed data structure T, with and without taking symmetry into account. Left: The base of the exponential growth is effectively reduced due to symmetry (smaller slope). Right: The slope of both growth rates is the same; symmetry offers only a constant offset, corresponding to a scale factor (90 here) for the same asymptotic behavior.

cost for all but the first bounce become prohibitive. In this section, we derive efficient and perceptually accurate approximations.

Some efficiency is obtained by taking advantage of symmetry in the terms of our polynomial. Referring back to Equations (12) and (13), we note that T_{12} and T_{21} both get multiplied by $c_1 c_2$. We can therefore define $\tilde{T}$ by summing all T entries which differ by only a permutation of indices. It can be shown that the number of terms now grows as $\binom{N+W-1}{N}$ which is slower than $(W)^N$ (Figure 5(left)). We can therefore consider a larger number of bounces, in practice, we use up to $N = 4$ bounces which we find provides a good degree of fidelity.

However, as Figure 5 shows, symmetry only reduces the complexity by a constant factor for growth with respect to W. While this factor is about 90 for the case of $N = 4$ shown, the $O((W)^4)$ behavior has not been reduced.

4.1 Low-Frequency BRDF Approximations

To deal with the explosion in the number of bases for R, we make the important observation that Equation (15) does not require us to use the same set of BRDFs R for every occurrence of $\mathbf{K}(R)$. We can define a hierarchy of BRDFs[3] for each object, using fewer bases to represent the BRDF when considering light-surface interactions that

---

[3]This hierarchy is never directly exposed to the user. The user simply edits BRDFs in the usual way, by adjusting parameters or editing high-resolution 1D curves. The system automatically filters these to lower-frequency versions where needed or computes diffuse equivalents as described in Appendix B.

are not directly visible to the viewer. For any bounce n on object m,

$$\rho_n^m(\omega_i, \omega_o) = \sum_{j=1}^{J_n} c_j^m b_j^m(\omega_i, \omega_o)$$

$$J = J_N \geq \cdots \geq J_n \geq \cdots \geq J_1 \geq 1, \qquad (17)$$

where c_j^m and b_j^m correspond to the appropriate hierarchy (and are not the same for different ρ_n.) This creates lower-frequency BRDF approximations, motivated by the often-discussed low-frequency nature of indirect lighting and in experiments by Nayar et al. [2006] and theoretical work that shows that, at each bounce, the BRDFs act as a low-pass filter [Ramamoorthi and Hanrahan 2001; Durand et al. 2005].

We now denote the distribution of BRDFs R with a superscript indicating the number of bases used per object,

$$B^N = \mathbf{K}(R^{J_N})\,\mathbf{G}\mathbf{K}(R^{J_{N-1}})\,\mathbf{G} \cdots \mathbf{K}(R^{J_n})\,\mathbf{G} \cdots \mathbf{K}(R^{J_1})\,\mathbf{G}E. \quad (18)$$

When Equations (14) and (15) use the corresponding hierarchy of R's, the subscripts are modified to run over a smaller domain. Specifically,

$$T^N_{(m_N j_N)\cdots(m_n j_n)\cdots(m_1 j_1)}(x, \omega_o)$$
$$= \mathrm{K}\big(R^{J_N}_{m_N j_N}\big)\mathbf{G} \ldots \mathrm{K}\big(R^{J_n}_{m_n j_n}\big)\mathbf{G} \ldots \mathrm{K}\big(R^{J_1}_{m_1 j_1}\big)\,\mathbf{G}E, \quad (19)$$

where the subscripts $(m_n j_n)$ explicitly denote the object m_n and BRDF basis function j_n. This is the most general form of the multi-bounce tensor T^N.

A variety of decreasing-frequency BRDF series within our editing system are illustrated in Figure 4. The scene is lit by an environment map and modeled after a figure in Ben-Artzi et al. [2006]. For clarity in the comparisons, we use only two bounces, and show series $(J_2, J_1) \equiv (128, 8); (16, 8); (128, 1)$. Figure 4(f) shows direct lighting only as in Ben-Artzi et al. [2006]. This omits important effects for the perception of materials and shininess like the reflection of the teapot in the shiny tray (a black shadow results instead) or the reflection of the spout in the teapot. Figure 4(a) is the reference image, rendered offline (with two bounces) in PBRT [Pharr and Humphreys 2004].

Figure 4 underscores that further bounces can be represented with very low-frequency BRDFs. The reference image in 4(a) is essentially identical to the (128, 8) BRDF series in 4(b) and 4(c) that uses only 8 BRDF bases for the second bounce. In fact, our direct material perception primarily responds to the glossy tray reflecting nearby objects like the teapot. Therefore, even the purely diffuse approximation for further bounces from the eye (128, 1) in 4(e) is usually
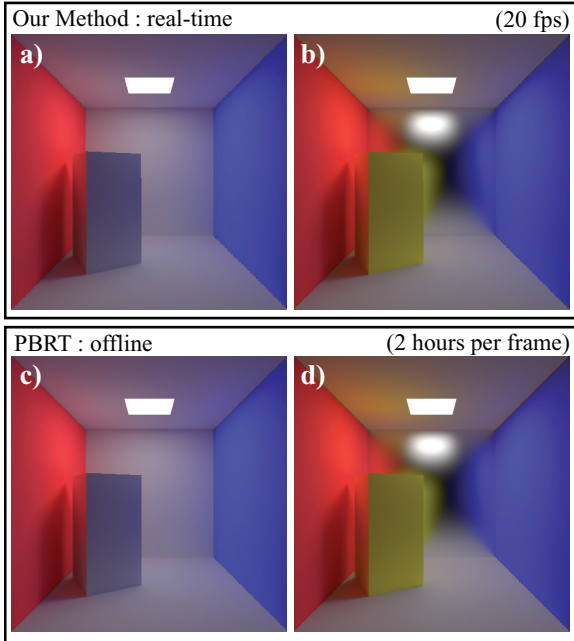
Fig. 6. BRDF edits in Cornell Box. We edit the color of the small box, and also make the back wall shinier, all while rendering interactively with four bounces of global illumination. Notice the correct color bleeding in the scene and correct glossy reflections of the scene in the back wall of (b). We compare our results to offline PBRT reference images in (c) and (d).

adequate. In that case, the teapot appears diffuse in the reflection in 4(e), but this approximation is only for objects seen indirectly through reflections, and not easily noticeable.

Often, our choices are dictated by available computational resources. For a given complexity $J_2 J_1 = 128$, two possible options are $(16, 8)$ in 4(d) and $(128, 1)$ in 4(e). Both images are quite accurate and can be edited within our system. They make different trade-offs. The glossy reflection of the teapot in the tray is slightly more accurate in $(16, 8)$ because of a better BRDF approximation for the second bounce. However, the first bounce is represented at lower frequency than $(128, 1)$, for example, direct reflections on the teapot are somewhat less sharp; this can become more noticeable for measured BRDFs and very shiny materials.

The observations from Figure 4 indicate that using a diffuse equivalent for further bounces is a reasonable and efficient approximation, and we use it for some of our examples (Section 4.2). For more complex effects like caustics, we explore instead a series where the second bounce from the viewer uses a low-frequency approximation (Section 4.3).

4.2 Diffuse Approximation for Further Bounces

In the limit, $J_{N-1} = 1$, and we approximate each object's BRDF using a single basis function that is a diffuse lobe, scaled by the equivalent albedo of the BRDF. See Appendix B for a derivation of the equivalent albedo, d_m. We usually use four bounces $(J, 1, 1, 1)$ (with typically $J = 64$).

$$B^N \approx \mathbf{K}(R^J)\,\mathbf{G}\mathbf{K}(R^1)\,\mathbf{G}...\mathbf{K}(R^1)\,\mathbf{G}E$$
$$T^N_{(j)(m_{N-1})\cdots(m_1)} = \mathbf{K}(R^J_j)\mathbf{G}\mathbf{K}(R^1_{m_{N-1}})\mathbf{G}\ldots\mathbf{K}(R^1_{m_1})\mathbf{G}E, \quad (20)$$
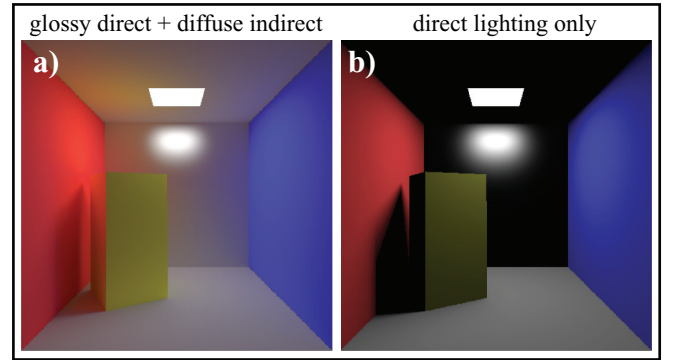
Fig. 7. Qualitative errors in simpler approximations. Compare 7(a) and 7(b) to Figure 6(b). Figure 7(b) is the direct lighting approximation of Ben-Artzi et al. [2006] which fails to capture many important global illumination effects. 7(a) is a sum of 7(b) and diffuse indirect lighting, capturing some global effects but providing an inconsistent material perception for glossy objects like the back wall.

where we have simplified the index pairs $(m_n j_n)$ in the general tensor of Equation (19) as follows. We drop the m_N subscript in the first index pair since only one object, $m_N(x)$, is visible through a pixel. We also drop the j_n subscripts for further bounces since there is only one basis BRDF when using R^1. Thus, we also simply use j instead of j_N for the bounce closest to the eye.

This approximation fully treats the first bounce from the viewer, including glossy reflections of the nearby scene. Bounces further from the viewer (and hence reflections of objects not seen directly) are treated as diffuse. The complexity at each pixel reduces from $O((W)^N)$ to $O(J(M)^N)$. We will later see how this reduces further to $O(J)$ for rendering since we edit only one object or material at a time.

Evaluation. Figures 6(a) and 6(b) are produced with our system. In 6(a), all surfaces are diffuse, while in 6(b) we edit the back wall to make it a glossy material, and change the color of the inner box to yellow. It is clear our method enables perception of material appearance because objects correctly reflect other nearby objects (see the glossy interreflections of the room in the back wall in 6(b)), while also accurately preserving global effects like color bleeding onto the large inner box. We compare to reference images using offline path tracing with PBRT in Figures 6(c) and 6(d), which confirms the accuracy of the approximation.

By contrast, Figure 7(b) shows the direct lighting approximation of Ben-Artzi et al. [2006] for the configuration in Figure 6(b). Not only is it missing a lot of energy, but it also lacks the reflections of the room in the back wall, which makes it difficult to assess the desired glossiness while editing.

Note that our method treats the first bounce from the eye with the full BRDF to get glossy reflections of nearby objects,

$$B \approx E + \mathbf{K}(R^J)\,\mathbf{G}E + \mathbf{K}(R^J)\,\mathbf{G}\sum_{N=2}^{\overline{N}}(\mathbf{K}(R^1)\,\mathbf{G})^{N-1}E.$$

Figure 7(a) compares to an alternative coarser approximation we originally tried using in our framework. This simply adds a diffuse indirect solution to the full direct lighting result. It is essentially a series $(1, 1, 1, 1)$ for the indirect illumination, and therefore the

most efficient technique,

$$B \approx E + \mathbf{K}(R^J)\,\mathbf{G}E + \mathbf{K}(R^1)\,\mathbf{G}\sum_{N=2}^{\overline{N}}(\mathbf{K}(R^1)\,\mathbf{G})^{N-1}E,$$

where the main difference is that $\mathbf{K}(R^1)$ is used instead of $\mathbf{K}(R^J)$ for the leftmost operator of the multiple-bounce terms. Figure 7(a) is clearly better than direct lighting only; some global illumination is usually better than none.

However, a comparison with Figure 6(b) shows that, while further bounces can be approximated as diffuse, the first bounce from the eye does need the full high-frequency BRDF. Unlike our method, Figure 7(a) gives an inconsistent material appearance of the back wall, that may be difficult to interpret while editing. While the direct reflection of the light source is glossy, the indirect reflections of the room appear diffuse.

4.3 Slower Decay of BRDF Series

Treating later bounces as diffuse works well in most scenes (see Figures 1, 4(e), 6, and 9). However, in some configurations like concave curved reflectors, higher frequency indirect effects like caustics are visually important [Durand et al. 2005].

To handle such challenging situations, we need to reduce the BRDF frequencies more slowly, using more (8–16) basis functions for the second bounce from the eye. (Cases where even the third or higher bounces away from the eye need to be high frequency are quite rare, though our general framework does not preclude taking them into account.) We compensate for the extra memory cost either by reducing the number of bases for the first bounce (Figure 4(d)) or by using fewer bounces (Figure 11 has only two bounces).

We have already seen an example of using 8 BRDF basis functions for the second bounce in Figure 4(b)–(d), that gives a more accurate reflection of the teapot in the shiny tray. In practice, our editing system also includes diffuse approximations for the third and fourth bounces to augment the series in Figure 4 (see Table II). An even more challenging example is Figure 11, that involves caustic effects. In this case, we use 16 BRDF basis functions for the second bounce, with a BRDF series of the form $(J, J/4) \equiv (64, 16)$.

5. IMPLEMENTATION

We now describe our implementation, starting with our precomputation method (Section 5.1), followed by the rendering algorithm (Section 5.2), and some practical extensions (Section 5.3).

5.1 Monte Carlo Precomputation

We need to precompute the tensors defined by Equation (19) at each pixel. The important special case for diffuse approximation in further bounces is given by Equation (20). Recall that the $\mathbf{K}$ operators involve integrals over the hemisphere, which means that each $T^N(x)$ requires nested (high-dimensional) integrals over ray paths. This is similar to traditional global illumination. We adapt Monte Carlo path tracing [Kajiya 1986] for precomputation because of its flexibility, ease of implementation, and negligible memory overhead.

Each value in the different tensors can be seen as a separate integral over the space of paths. However, it is easier and more similar to traditional path tracing to sample path space for all integrals at the same time. We generate random paths and, for each path, update the appropriate tensor integrals. We must modify three basic aspects of the path tracer. First, we cannot generate rays by sampling the BRDF (since it is unknown). Second, we must add each path's contribution


Fig. 8.   Consider a path $\vec{X}$ from the light (ℓ) to the eye (e). The light hits objects B, C, D, B before reaching the eye. At the final bounce (x_4), the view direction defines how the basis functions of B's BRDF divide the incoming light directions. Of the three bases ($a, b,$ and c), the configuration of the last bounce places it in c. Therefore, this path contributes to $T^4_{Bc,D,C,B}$.

to the correct tensor element as opposed to simply contributing to the final pixel radiance. Third, we must compute the contribution of each path using basis BRDFs.

Consider a given path $\vec{X}$ from a point on a light source (ℓ) to the eye (e) that passes through points $x_N, x_{N-1}, \ldots, x_1$ on objects $m_N, m_{N-1}, \ldots, m_1$ as illustrated in Figure 8.

Sampling Path Space. According to Monte Carlo theory, any random sampling can be used to generate new directions when building up the path. We follow standard path tracing and generate rays recursively from the eye. We sample the light source at each bounce to generate all path lengths.

Path tracers usually importance sample the BRDF to select a new direction at each intersection. Unfortunately, we have no knowledge of the final BRDF. We cannot sample according to the basis BRDFs either because they will be combined with arbitrary weights at runtime. The simplest approach would be to sample the cosine-weighted hemisphere uniformly, but this would yield high variance when sharp specular lobes are used. Instead, we take advantage of the general form of typical BRDFs and sample according to a mixture of 70% diffuse and 30% high-gloss (Blinn exponent of 200). This places more importance on specular directions and enables low-noise results for the full range of glossy to near-mirror BRDFs in a practical editing session.

Tensor Update. For each random path, we need to accumulate a contribution to the appropriate tensor element. In effect, we are computing coefficients in a symbolic polynomial representation of the basis BRDFs (in the spirit of symbolic rendering by Séquin and Smyrl [1989]). In our case, we have chosen bases that do not overlap, and therefore a given path requires updating exactly one tensor element. The j index in Equation (20) is determined by the basis function that contains the configuration of incoming and outgoing directions at the last intersection point (in Figure 8, this is x_4). The outgoing direction for the bounce to the eye ($x_4 - e$) partitions the space of incoming directions into bands corresponding to our different box-basis functions ($a, b,$ and c in Figure 8). In the example, band c contains the incoming path direction which determines j. More generally, we would choose the j for which $b_j(\omega_i, \omega_o)$ is nonzero. In the case of the diffuse approximation for further bounces, we only care about the objects containing the further bounces (the indices $m_{N-1}\ldots m_1$ in Equation (20)). In Figures 8, these are DCB. The more slowly decaying BRDF series with multiple specular bounces would use a similar band selection for the second bounce from the eye as for the first.

Tensor values involve a standard Monte Carlo sum,

$$T^N_{w_N w_{N-1} \ldots w_1}(x) = \frac{1}{Q} \sum_i \frac{f(\vec{X}_i)}{p(\vec{X}_i)}, \qquad (21)$$

where Q is the number of paths through pixel x, and the sum runs only over paths $\vec{X}_i$ that correspond to the specific subscripts (bands and objects) in T^N. $f(\vec{X}_i)$ is the contribution of $\vec{X}_i$, and $p(\vec{X}_i)$ is the probability of generating it.

Path Contribution. In standard path tracing, the path contribution $f(\vec{X}_i)$ is the direct visible lighting at x_1, multiplied by the product of BRDFs (corresponding to **K**) and cosine terms at each intersection (the visibility in **G** is already considered when creating valid paths). In our case, we must instead multiply by the appropriate basis BRDFs.

For the first bounce from the eye, we use

$$\tilde{b}_j(e, x_N, x_{N-1}) = b^{m_N(x_N)}_j(\omega_i, \omega_o) \cos\theta_i, \qquad (22)$$

where $\omega_i(x_N, x_{N-1})$ and $\theta_i(x_N, x_{N-1})$ depend on the direction of the incident ray and $\omega_o(e, x_N)$ on the outgoing view direction. For the slowly decaying series, a very similar form can be used for the second bounce from the eye, simply considering a lower-frequency $\tilde{b}_j(x_N, x_{N-1}, x_{N-2})$. For the other bounces, we use the single diffuse basis:

$$D(x_n, x_{n-1}) = \frac{1}{\pi} \cos(\theta_i(x_n, x_{n-1})). \qquad (23)$$

Finally, $f(\vec{X})$ is a product of the terms at each bounce. For the diffuse approximation for further bounces, this is

$$f(\vec{X}) = \tilde{b}^{m_N(x_N)}_j(\omega_i, \omega_o) D(x_{N-1}, x_{N-2}) \ldots D(x_1, \ell) E(\ell). \quad (24)$$

Optimizations. Our precomputation is essentially the same complexity as rendering a single image with MCPT. Moreover, many standard path-tracing optimizations can still be applied. For example, we have adapted irradiance caching [Ward et al. 1988]. Instead of generating paths that terminate at the light source, we find the direct lighting at the last surface point x_1. We cache the irradiance in a preprocess that samples visibility on a grid within each triangle.

5.2 Rendering in Real Time

We now focus on the runtime rendering computation for each pixel. For compactness of notation, this section will deal primarily with the diffuse approximation for further bounces, but more slowly decaying series use very similar methods and are discussed briefly at the end.

To simplify notation, we denote the tensor as $T^N_{jz}(x)$, where the single super index z is a shorthand for writing out $m_{N-1} \ldots m_1$ in Equation (20) (viewed as an index, $z \in [1, Z = (M)^{N-1}]$). Similarly, we also denote the product of diffuse equivalents d_m of each object by $\overline{d}_z$,

$$
\begin{aligned}
z &\equiv \{m_{N-1}, \ldots, m_n, \ldots, m_1\} \\
\overline{d}_z &\equiv d_{m_{N-1}} d_{m_{N-2}} \ldots d_{m_n} \ldots d_{m_2} d_{m_1}.
\end{aligned} \qquad (25)
$$

Note that z represents a list of indices, while $\overline{d}_z$ is a single number corresponding to the product of the albedos d_m.

Finally, we can adapt Equation (14) for rendering,

$$B^N(x) = \sum_{j=1}^{J} \sum_{z=1}^{Z} T^N_{jz}(x) \, c_j \overline{d}_z. \qquad (26)$$

During the edit, the user specifies the BRDF coefficients c_j (either directly by editing a curve or implicitly by editing a parametric model). The diffuse equivalents d (and hence $\overline{d}_z$) are then computed as described in Appendix B. Finding the c_j and $\overline{d}_z$ occurs once per frame for the whole image. Using the precomputed $T^N_{jz}(x)$, the double summation in Equation (26) must now be evaluated at each pixel.

Object Freezing. Equation (26) requires $O(JZ)$ operations per pixel. On modern hardware, this is fast but still not real time (requiring a couple of seconds per update). To reduce complexity, we observe that a user edits the BRDF of only one object at a time. We use a runtime precomputation that performs the bulk of the calculations in Equation (26) by temporarily freezing the BRDFs of the remaining objects.

Recall that $\overline{d}_z$ represents a multivariable polynomial in the d's of the objects in the scene. For example, if we have $z = \{1, 3, 2, 1, 5, 3\}$, $\overline{d}_z = (d_1)^2 d_2 (d_3)^2 d_5$. However, if all but one of the d's are fixed, this becomes just a single-variable polynomial in the unfrozen d of the object being edited. If in this example all objects except object 1 are frozen, we can define a constant $A = d_2(d_3)^2 d_5$, so that $\overline{d}_z$ becomes a simple quadratic polynomial, $\overline{d}_z = A \cdot (d_1)^2$ in only the edited variable d_1.

To implement this scheme more formally, we need a helper function $n(z, i)$ that tells us how many times a given edited object i appears in z. In the previous example, $n(z, i) = 2$ (for $i = 1$) tells us that $\overline{d}_z$ is a quadratic polynomial in d_i alone. We can also define the constant A more formally as $A = \overline{d}_z/(d_i)^n$. Finally, we compute at runtime a new tensor of coefficients at each pixel where each row represents a single-variable polynomial in d_i,

$$F^i_{jn}(x) = \sum_N \sum_z \begin{cases} n(z, i) \neq n : 0 \\ n(z, i) = n : T^N_{jz}(x) \frac{\overline{d}_z}{(d_i)^n} \end{cases} \qquad (27)$$

Our real-time rendering step is now a direct evaluation of

$$B(x) = \sum_{n=0}^{\overline{N}-1} (d_i)^n \sum_{j=1}^{J} c_j F^i_{jn}(x). \qquad (28)$$

For each power $(d_i)^n$, we simply evaluate a dot-product $c_j F_j$, essentially as in a standard linear PRT formulation.

The computation in Equation (27) requires $O(JZ)$ operations per pixel, comparable to simply evaluating Equation (26) directly once. This requires a short (usually 5-10 seconds) mode switch each time the user begins editing a different object. The real-time rendering in Equation (28) is now $O(J)$ (the number of bounces $\overline{N}$ is a small constant, usually four.)

Two further optimizations are possible. In a practical editing session, the coefficients c_j change slowly from frame-to-frame, especially if we transform into a wavelet representation. This temporal coherence can be directly exploited using the incremental wavelet rendering method described in Ben-Artzi et al. [2006]. Finally, if we are rendering a pixel of an object that is not being edited, the c_j do not change at all. (Note however, that the object's appearance will still be affected because of global illumination.) This makes it possible to further precompute $V^i_n = \sum_j F^i_{jn} c_j$, reducing the cost to evaluating a simple polynomial in d_i.

Slower Decaying Series. We briefly describe the generalization to more slowly decaying series as in Section 4.3. The general rendering operation of Equation (26) is now best described as a triple summation since we are dealing with three distinct representations

of R: R^{J_N}, $R^{J_{N-1}}$, and R^1.

$$B^N(x) = \sum_{j=1}^{J_N} \sum_{w=1}^{M J_{N-1}} \sum_{z=1}^{Z} T_{jwz}^N(x)\, c_j \hat{c}_w \overline{d}_z, \qquad (29)$$

with $\hat{c}_w$ denoting the lower-frequency BRDF coefficients for the second bounce (J_{N-1} bases on each of the M objects).

Object freezing is a bit more difficult, theoretically requiring the creation of a three-dimensional F_{jkn}^i, where $j \in [1, J_N]$, $k \in [1, J_{N-1}]$, $n \in [0, \overline{N}{-}2]$. In practice, we think of the j as one dimension, and $k' \equiv kn$ as the other.

5.3 Extensions

We briefly describe two important practical extensions.

Objects with Fixed BRDFs. Large scenes can contain many small objects that cause an exponential increase in memory requirements. Such scenes usually do not require editing the BRDFs of all objects. When designing the materials in a room, one typically does not want to change the BRDFs of small placeholder objects like books or toys. We extend our algorithm by implementing the ability to fix the BRDFs of certain objects at precomputation. Note that their shading is still updated, based on global illumination from other surfaces. This should also not be confused with runtime object freezing, which occurs temporarily during an editing session.

In precomputation, instead of using diffuse equivalents D or BRDF bases $\tilde{b}$, we must use the full known BRDF $\rho^m(\omega_i, \omega_o) \cos \theta_i$ for reflections from fixed object m. Rendering is unchanged for editable objects since there are no new BRDF bases. For the fixed object, we still use $T_z^N(x)$ and multiply by the diffuse albedos of editable objects. However, the BRDF bases (and index j) need not be considered. In Figure 1, the table and its legs have fixed BRDFs.

Spatial Weight Maps. So far, we have focused on BRDF effects. Spatial variation can be handled with textures to modulate the BRDF over an object's surface. If we do not seek to edit them, the textures can be directly incorporated into the BRDF basis functions (as the multiplicative $H^m(x)$ terms in Equation (5)). Finally, while we have discussed a single BRDF per object for clarity, our framework and implementation can handle multiple colocated BRDFs for each object. If we also seek to modify the spatial blending of BRDFs, as we do for the floor in Figure 1, we can simply modulate the directly viewed surfaces in image-space by the multiplicative spatial blending weights (weight maps) or texture. For global illumination, we are concerned only with the low-frequency behavior of the weight maps or textures, and we modulate the diffuse equivalent albedos by the average value of the weight map for each BRDF layer.

6. RESULTS

Section 6.1 briefly discusses the types of edit performed on the scenes in Figures 1, 4, 6, 9 and 11 and the global illumination effects involved. Then, Section 6.2 gives performance results for precomputation time and memory usage, while Section 6.3 discusses rendering frame rates.

6.1 Editing and Visual Effects

Cornell Box. Figure 6 shows the Cornell box where we edit the parameters of a Blinn-Phong and diffuse reflectance model. In going from 6(a) to 6(b), we make the back wall glossy with correct interreflections of the nearby scene and change the color of the inner box, demonstrating accurate color bleeding. Note that even the


Fig. 9. (a) original, (b) closeup after anisotropic edit of vase in (a), (c) closeup after fresnel effect strongly increased for (b).

Fig. 10. Closeups for Figure 1. Note the color bleeding of the couch onto the stairs.

simple color adjustment was not possible interactively with global illumination in previous methods.[4]

Teatray. Figure 4 shows a teatray scene. The teapot has a Cook-Torrance BRDF with specular and diffuse components and the handles and tray a Blinn-Phong model. While only a single set of BRDFs for the objects is shown in Figure 4 for brevity, we can freely edit the teapot, tray, and handles in real time, using any of the BRDF series shown in the figure (extended to three or four bounces).

Vase. Figure 9 shows a variety of flexible BRDFs possible within our system. The flowers are diffuse, while the stems are a diffuse+specular BRDF to enable a glossy coating. The table has diffuse and specular BRDFs with the diffuse shown textured in Figure 9. The vase uses a diffuse and a specular layer. The specular BRDF is an Ashikhmin-Shirley BRDF with fixed exponent of 225 but adjustable ratio to adjust the direction and amount of anisotropy (see Appendix A). A second specular layer allows for edits to the Fresnel term. Appendix A describes how to enable Fresnel control for any BRDF via a second additive layer.[5]

---

[4]Since the Cornell Box was designed for verifiable comparison to PBRT, it was precomputed with the floor and back wall using Blinn-Phong with 64 bands, while the other objects use 2 bands to represent a pure diffuse BRDF with editable albedo.

[5]The glossy layers of the stems and table use 64 bands. The Ashikhmin-Shirley layers of the vase use 256 bands. All diffuse layers use 2 bands.

Fig. 11. To obtain accurate indirect reflections, we enable second-bounce glossy reflections. We show results with an implementation that uses $J = 64$ for the first bounce from the eye and $J = 16$ for the second. With this approximation, we can not only get caustics on the floor from the ring but also see the reflection of the floor in the ring accurately as shown in the rightmost image.


Fig. 12. In the first scene, our method using the diffuse approximation (a) produces results that are nearly identical to the three-bounce PBRT reference image (b). In the second scene, the couch and floor have both been made glossy (blinn exponents of 80 for the couch and 250 for the floor), but the diffuse approximation (c) still produces results that are visually close to the three-bounce PBRT reference image (d). The main differences between (c) and (d) are in areas where a large contribution of light interacts with two or more glossy BRDFs (e.g., the couch cushions). The indirect lighting from the couch to the walls and staircase, as well as the indirect lighting from the walls to the couch, are both captured accurately.

Room. Figure 1 shows one potential application of our system to design interiors. In this case, indirect light is critical since a large part of the scene, like the back wall and much of the couch, would be black with direct lighting only. Most of the indirect illumination is low frequency so we use our diffuse approximation for further bounces. For this scene, we use three bounces with a BRDF series (64, 1, 1).

This is a complex scene with 8 objects (6 editable), environment lighting from the windows, and a variety of materials including measured reflectance and textures which can all be edited. Our system renders interactive feedback with global illumination, enabling the user to correctly perceive and edit materials to design the interior of the room.

The effects of global illumination are clearly seen in the closeup views of Figure 10 where the couch color bleeds onto the stairs. Notice also the glossy reflection of the green couch in the glossy floor on the right. The lower portion of the couch is lit only by indirect illumination so this is actually a glossy reflection of indirect light and needs at least 3 bounces to simulate properly.[6]

Ring. Figure 11 shows how our system can be used even to choose materials to create the desired complex global illumination effects like caustics. In this case, we use a slower series decay (64, 16) that includes two specular bounces to obtain accurate indirect specular reflections. Our system interactively updates both the caustics on the

floor from the ring and the reflection of the floor in the ring accurately as the BRDFs are edited. Note that the sharpness of the caustics and indirect reflections are maintained even though the second bounce BRDF is still quite low frequency. Without our system, it would be quite difficult to interactively select the glossiness of the plane and ring to explore possible appearances of the caustics.

Comparison to Reference Images. Our system is designed to produce results that convey the important visual features, including complex highly specular materials and indirect lighting. Figure 4(a) shows the reference PBRT rendering that is seen to be quite close to our low-frequency BRDF approximation for the teapot scene (Figure 4(d) and 4(e)), including reflections of the teapot in the glossy floor. Figures 6(c) and 6(d) compare our method to offline rendering with PBRT for the Cornell box, demonstrating excellent agreement, including correct color bleeding and glossy reflections of the back wall. Finally, Figure 12 compares our result with reference PBRT renderings for the complex room scene in Figure 1 (note that there is some aliasing in the PBRT images since we made the decision to have PBRT sample only the center of image pixels to speed up rendering and more closely match the implementation of our method). In Figure 12(a) and Figure 12(b), with mostly low-frequency materials, our results match the reference PBRT image closely because the low-frequency interreflection effects are correctly handled. In Figure 12(c) and 12(d) that have more specular materials, some differences can be seen where light bounces between two or more glossy BRDFs. In these cases, our approximation of diffuse reflectance for further bounces from the eye leads to some discrepancies with ground truth. Nevertheless, the results give a correct perception of material properties as required for interactive reflectance design.

---

[6]Different objects in the room were computed with different BRDF parameterizations and resolutions. The couch and floor have 64 bands for specular BRDFs. The floor also has a diffuse layer with 2 bands. The stairs and sills use 16 bands for glossy BRDFs. The walls and ceiling have only a diffuse BRDF with 2 bands, allowing color and intensity edits but not glossy.

Table II. Table II. Precomputation Time, Storage of Precomputed Data Structures, and Rendering Time for Each of Our Scenes

scene specifications				precomputation			storage per path-length (MB)					rendering	
name	*resolution*	*M*	*BRDF series*	*samples*	*lights*	*time(h:m)*	*1*	*2*	*3*	*4*	*total*	*freeze*	*fps*
Cornell box	450×450	6	64,1,1,1	12K	81/81	5:58	17	104	365	975	1461	10s	22
Vase	512×512	4	64,1,1,1	4K	10K/400	5:29	65	262	637	—	964	8s	19
Room	640×480	6	64,1,1	8K	14K/800	9:25	27	165	563	—	755	7s	20
Teatrays	320×240	3	128,1,1,1	10K	10K/2K	1:07	13	37	75	125	250	3s	25
			16,8,1,1			1:42	2	37	113	333	485	3s	16
			128,8,1			2:10	13	301	853	—	896	15s	5
Ring	400×400	2	64,16	4K	4K/400	6:15	20	607	—	—	627	16s	6

The scene specification includes image resolution, number of objects (M), and the BRDF series. Precomputation lists the number of samples per pixel for path tracing and the number of point lights used to sample the environment for direct/indirect lighting (we usually use fewer samples for indirect illumination). Storage is shown separately (in MB) for each pathlength or number of bounces. Rendering time indicates the time for object freezing when selecting a single object to edit and for real-time rendering.

6.2 Precomputation Times and Memory Usage

Table II describes precomputation times and memory usage. The rows show the scenes (including multiple BRDF series for the teatray). The image resolutions were chosen primarily to fit the resolution of the accompanying video—the teatrays were computed at lower resolution for faster testing and comparison to PBRT renders (which took hours to generate low-noise still images at these resolutions).

Precomputation Time. The precomputation time (on an Intel Xeon 3.2 GHz 64-bit machine) ranges from one to several hours, depending linearly on the number of path tracing samples used and also varying with the number of point lights to sample the environment. Interestingly, these wall clock times are about as fast (and sometimes faster than) for standard PBRT to render a single (uneditable) image of the scene. This is because the complexity of our precomputation is essentially the same as rendering a single image with path tracing. Thus, our precomputation times, while large, are comparable to those for high-quality global illumination image synthesis methods.

Memory Usage. The memory usage grows for each bounce since there are more polynomial terms in T^N (as shown in Figure 5 (left)). The growth is relatively slow, each bounce increases the size of T^N by a factor of about 3 for higher bounces. Nevertheless, the highest (usually fourth) bounce requires more than half the total memory. This is an interesting direction for future work since it is usually the darkest and most amenable to compression. However, in our experiments with direct quantization and wavelet compression, we were not easily able to find a scheme that was free of image artifacts. The problem is different from compression for relighting since we visualize the BRDF and image directly, making artifacts more apparent. We instead simply drop zero values and use RGBE compression. At the end, our total precomputed data structures are in the range of several hundred MB. While this is large, it is comparable, for instance, to all-frequency relighting approaches on similar resolution images.

Note that the storage sizes shown in Table II are larger than the theoretical runtime memory requirements. Due to our object-freezing step, only the smaller F^i_{jn} in Equation (27) needs to be in main memory. We can avoid preloading all of the data at the cost of higher object-switch times.

In comparing the different BRDF series for the teatray, (128, 8, 1) requires the most memory because of the extra factor of 8 for second-bounce BRDF bases (as opposed to (16, 8, 1, 1) and (128, 1, 1, 1)). To keep memory consumption reasonable, we limit to 3 bounces rather than 4 as with the other series. These numbers are comparable to the ring that also uses 16 bases in the second bounce. (We similarly limit the number of bounces in the ring to two (64, 16)).

6.3 Rendering Speed

For simplicity, we pursue a purely software implementation, although the simple dot-product form of Equation (28) indicates that GPU implementations, similar to those recently developed for relighting, should also be applicable.

As can be seen in the video and the last column of Table II, our software method achieves frame rates of 15–25 fps on most scenes. The time for object-freezing when we switch from one object to another for editing is about 5–10 seconds and isn't disruptive to typical editing sessions, as seen in the video. The rendering and mode switch times are somewhat larger when there are more BRDF bases for the second bounce (one teatray example and ring) on account of the additional complexity. However, they are still interactive, and our video demonstrates interactive editing of the ring scene. In summary, our results and video for the first time show practical real-time BRDF editing as interactive feedback is rendered with global illumination.

7. LIMITATIONS

Besides the design decision to limit ourselves to fixed lighting and view, our method currently has a number of other limitations. In some cases, there are simple ways in which these limitations can be overcome as described in the following. Note that the theoretical framework in Section 3 is completely general, and future implementations can make different complexity and practical trade-offs as appropriate.

Textured Materials. The method in this article depends on writing the BRDF distributions as a linear combination of basis BRDFs. Spatially-varying materials where the BRDF varies significantly at each pixel (such as a texture controlling the Phong exponent) are not addressed in our article.[7]

However, we can consider complex spatially-varying materials that are linear combinations of basis BRDFs as in Lensch et al. [2001] and Lawrence et al. [2006], with the $H(x)$ functions in Section 3 indicating the appropriate weights as discussed in Section 5.3.

If we want to edit the spatial blending weights (texture patterns), as opposed to baking them into the precomputation, we need to make some approximations. For direct lighting, this is simply a modulation and is easily addressed. For indirect lighting, to compute the equivalent albedo, we average the overall intensity of the texture.

---

[7]By using a deep texture to modulate each individual basis BRDF it would be possible to do more complicated edits, such as change the Phong exponent, using a spatially-varying texture. This would require texture editing techniques and other changes that we do not address.

This does have the limitation that the surface would appear uniform and untextured as seen in the reflection on a nearby glossy object. Nevertheless, our system allows textures to contribute for the first bounce and average albedo to material perception. In Figure 1, we load in different spatial weight maps (textures) for the room's floor to demonstrate runtime texture changes.

Scalability. The nature of global illumination requires accounting for an increasing number of events for higher bounces. This complexity, and our techniques for dealing with it, are discussed in detail in Section 4. One consequence is that our method is currently not scalable to scenes with tens to hundreds of objects, where each has an editable BRDF. We can get around this by using fixed BRDFs for small objects that need not be edited as discussed in Section 5.3. In the more general case, however, an iterative approach would need to be employed. The user needs to edit a few materials at a time (such as a single couch and table), followed by a new precomputation to enable edits to other objects, while fixing the BRDFs of the materials already edited.

Complex BRDFs. A general editing system should be able to handle almost any analytic or measured BRDF. In practice, memory and computation limit us to 1D functions of a given BRDF parameterization (often the half-angle). Nevertheless, our representation does capture the full 4D-nature of the BRDF. Appendix A describes the most general form amenable to editing within our system. As described in Ben-Artzi et al. [2006] and recapped in Appendix A, the quotient BRDF can be arbitrarily complex. It is only the editable behavior that must be describable as a function of some 1D-parameterization. Figure 9 shows a single material (the vase) for which we edit both the direction of anisotropy, and then the Fresnel effect, always rendering with the Geometric Attenuation Factor as part of the BRDF model.

Sampling Noise and Banding Artifacts. The specific implementation choices for our precomputation may, in some cases, lead to banding and sampling artifacts. Both the environment and area lights are sampled with static point samples. This is an implementation choice made in the ray tracer for precomputation and not a limitation of our approach. Environment map samples were chosen with a method similar to Agarwal et al. [2003]. This is the largest contributor to the banding artifacts, which are most noticeable when using materials that are near-mirror.

For indirect lighting, we sample, as described in Section 5.1, with 70% diffuse and 30% Blinn exponent 200. While this greatly reduces the noise for most materials, near-mirrors still show some noise. We note that near-mirrors are the ultimate stress test for our system since they effectively promote each of the light bounces by 1. However, we still handle them rather well, and, in other realistic materials, these effects are much less noticeable. Note that the quality of the images can be chosen by the user, depending only on the CPU time allocated to the preprocess. (Indeed, some previous precomputation-based methods have allocated 100's of CPU hours to produce very high-quality images.)

8. CONCLUSIONS AND FUTURE WORK

This article has described a complete theoretical analysis and practical implementation of a real-time rendering method, which enables interactive editing of BRDFs with global illumination effects. We expect that our method will have significant applications to design in computer graphics where the artist can now interactively specify material properties in the final scene with complex lighting, shadows and interreflections.

In the process, we develop a new precomputation-based framework that can handle nonlinear effects involving multivariable poly-

nomials for multiple bounces. Our contributions include a general theoretical framework for expressing global illumination as a precomputation-based interactive rendering process based on reflection and geometric operators, an analysis of computational complexity to develop tractable approximations, and effective precomputation and rendering methods and extensions. It is likely that many of these insights could be used in other contexts, like PRT for relighting, or even offline global illumination.

More generally, we have proposed a new and efficient method for symbolically rendering an image. Instead of accumulating each path in the color of a single pixel, that path is effectively stored in symbolic form, including the product of all BRDF terms encountered along it. Paths involving the same BRDF terms are accumulated in the appropriate tensor coefficient. This symbolic approach is likely to have broad applications in other domains where we seek to interactively edit or explore the space of material parameters, such as animation, simulation, and geometric modeling.

APPENDIXES

Appendix A: BRDF Basis Functions

Our framework is general for any choice of linear BRDF basis functions but benefits from those tailored to the natural space in which material edits occur. In practice, we use the 1D-reparameterized box basis functions of Ben-Artzi et al. [2006],

$$\rho(\omega_i, \omega_o) = \rho_q(\omega_i, \omega_o) f(\gamma); \quad f = \sum_j^J c_j b_j(\gamma(\omega_i, \omega_o)), \quad (30)$$

where f is the 1D-editable factor, and ρ_q is the quotient BRDF that captures more complex but uneditable behavior like normalization constants and the $\mathcal{GAF}$. f can be set directly by the user by editing a 1D-curve or computed by the system based on the user setting parameters of analytic BRDFs. The previous form involves an appropriate 1D-parameterization γ of the BRDF's 4D-domain. Some examples of γ are: θ_{half}, θ_{diff}, θ_{in}, and θ_{out}.

Some BRDFs have two factors such as Cook Torrance with $\gamma = \theta_{half}$ to control specular behavior and $\gamma = \theta_{diff}$ to control the Fresnel effect. We use only a single factor for practical reasons, to keep memory requirements manageable. However, we show in the following that many of the important two-curve edits can be achieved with a single factor, by careful use of the quotient BRDF.

The Fresnel effect is the most common use of the θ_{diff} factor. If we use the Schlick [1994] approximation, a BRDF that includes a Fresnel term (e.g. Cook-Torrance) becomes

$$\rho = \rho_q f(\theta_h)(F + (1 - F)(1 - \cos \theta_d)^5). \quad (31)$$

F is a function of the wavelength (color channel) and index of refraction only. This allows us to define ρ as the sum of two BRDFs, each with just one editable factor but different ρ_q,

$$\rho = \rho_q f_1 + \rho_{q_2} f_2; \quad f_1 = F f(\theta_h)$$
$$\rho_{q_2} = (1 - \cos \theta_d)^5 \rho_q; \quad f_2 = (1 - F) f(\theta_h) \quad (32)$$

At runtime, F is evaluated based on the user's choice of index of refraction, and f_1 and f_2 are set via the user interface. The two BRDFs are computed and summed (just as a specular and diffuse layer would be summed) to yield an accurate composite.

Anisotropy is the result of an elongated highlight. As presented in Ben-Artzi et al. [2006], two factors can be used to adjust the width of the highlight along the tangent and binormal directions. If we know the overall specularity of the material, we can separate the Ashikhmin-Shirley BRDF into a quotient that captures the width of

the highlight and a 1D-factor that controls the ratio of the elongation in the two perpendicular directions:

$$\rho_{AS} = \rho_q(\cos\theta_h)^{n_u\cos\phi_h}(\cos\theta_h)^{(rn_u)\sin\phi_h}; \quad r \equiv n_v/n_u \quad (33)$$

$$\rho_{AS} = \rho_{qu}(\omega_i, \omega_o)(\gamma_r(\omega_h, \omega_o))^r \quad (34)$$

$$\gamma_r = (\cos\theta_h)^{n_u\sin\phi_h}; \quad \rho_{qu} = \rho_q(\cos\theta_h)^{n_u\cos\phi_h} \quad (35)$$

A similar single-factor form can be obtained if the amount and direction of anisotropy (r) is known and only the width of the highlight needs to be adjusted.

Appendix B: Equivalent Albedo

We choose the equivalent albedo to match the average BRDF value, or more exactly, the output power for a uniform incident radiance field. This also corresponds formally to choosing the best perturbed **K** operator as in Arvo et al. [1994]. In other words,

$$
\begin{aligned}
d &= \frac{1}{\pi}\int\int \rho(\omega_i, \omega_o)\cos\theta_i\cos\theta_o\, d\omega_i d\omega_o \\
&= \frac{1}{\pi}\sum_j c_j \int\int \rho_q(\omega_i, \omega_o)b_j(\gamma(\omega_i, \omega_o))\cos\theta_i\cos\theta_o\, d\omega_i, d\omega_o.
\end{aligned}
\tag{36}
$$

The term in the integral now depends only on known quantities—the quotient BRDFs and the basis functions, and can therefore be evaluated by dense Monte Carlo sampling (this needs to be done only once for a given parameterization, not even for each scene). Call this e_j. Finally, at runtime, we simply need to compute

$$d = \frac{1}{\pi}\sum_j c_j e_j, \tag{37}$$

with the predetermined e_j and the dynamically chosen c_j.

REFERENCES

AGARWAL, S., RAMAMOORTHI, R., BELONGIE, S., AND JENSEN, H. 2003. Structured importance sampling of environment maps. *ACM Trans. Graph. (SIGGRAPH'03) 22*, 3.

ARVO, J., TORRANCE, K., AND SMITS, B. 1994. A framework for the analysis of error in global illumination algorithms. In *SIGGRAPH'94*, 75–84.

ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. 2000. A mirofacet-based BRDF generator. In *SIGGRAPH'00*, 65–74.

BARBIC, J., AND JAMES, D. 2005. Real-time subspace integration of St. Venant-Kirchoff deformable models. *ACM Trans. Graph. (SIGGRAPH'05) 24*, 3, 982–990.

BEN-ARTZI, A., OVERBECK, R., AND RAMAMOORTHI, R. 2006. Real-time BRDF editing in complex lighting. *ACM Trans. Graph. (SIGGRAPH'06) 25*, 3, 945–954.

COHEN, M., AND WALLACE, J. 1993. *Radiosity and Realistic Image Synthesis*. Academic Press.

COLBERT, M., PATTANAIK, S., AND KRIVNEK, J. 2006. Brdf-shop: Creating physically correct bidirectional reflectance distribution functions. *IEEE Comput. Graph. Appl. 26*, 1.

DORSEY, J., ARVO, J., AND GREENBERG, D. 1995. Interactive design of complex time dependent lighting. *IEEE Comput. Graph. Appl. 15*, 2 (March), 26–36.

DURAND, F., HOLZSCHUCH, N., SOLER, C., CHAN, E., AND SILLION, F. X. 2005. A frequency analysis of light transport. *ACM Trans. Graph. (SIGGRAPH'05) 24*, 3, 1115–1126.

HASAN, M., PELLACINI, F., AND BALA, K. 2006. Direct to indirect transfer for cinematic relighting. *ACM Trans. Graph. (SIGGRAPH'06) 25*, 3, 1089–1097.

KAJIYA, J. 1986. The rendering equation. In *SIGGRAPH'86*.

KONTKANEN, J., TURQUIN, E., HOLZSCHUCH, N., AND SILLION, F. 2006. Wavelet radiance transport for real-time indirect lighting. In *Proceedings of the EuroGraphics Symposium on Rendering*.

LAWRENCE, J., BEN-ARTZI, A., DECORO, C., MATUSIK, W., PFISTER, H., RAMAMOORTHI, R., AND RUSINKIEWICZ, S. 2006. Inverse shade trees for non-parametric material representation and editing. *ACM Trans. Graph. 25*, 3, 735–745.

LENSCH, H. P. A., KAUTZ, J., GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2001. Image-based reconstruction of spatially varying materials. In *12th Eurographics Workshop on Rendering*, 103–114.

MCCOOL, M., ANG, J., AND AHMAD, A. 2001. Homomorphic factorization of BRDFs for high-performance rendering. In *SIGGRAPH '01*, 171–178.

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 Trans. Graph. (SIGGRAPH'06) 25*, 3, 935–944.

NG, R., RAMAMOORTHI, R., AND HANRAHAN, P. 2003. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Trans. Graph. (SIGGRAPH'03) 22*, 3, 376–381.

NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of BRDF models. In *Proceedings of the Europe Graphics Symposium on Random*, 117–126.

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

RAMAMOORTHI, R., AND HANRAHAN, P. 2001. A signal-processing framework for inverse rendering. In *SIGGRAPH'01*, 117–128.

SCHLICK, C. 1994. An inexpensive BRDF model for physically-based rendering. *Comput. Graph. For. 13*, 3, 233–246.

SÉQUIN, C. H., AND SMYRL, E. K. 1989. Parameterized ray tracing. In *SIGGRAPH'89*, vol. 23.

SLOAN, P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Trans. Graph. (SIGGRAPH'02) 21*, 3.

SUN, W., AND MUKHERJEE, A. 2006. Generalized wavelet product integral for rendering dynamic glossy objects. *ACM Trans. Graph. (SIGGRAPH'06) 25*, 3, 477–487.

SUN, X., ZHOU, K., CHEN, Y., LIN, S., SHI, J., AND GUO, B. 2007. Interactive relighting with dynamic BRDFs. *ACM Trans. Graph. (SIGGRAPH'07) 26*, 3, 27.

WANG, R., TRAN, J., AND LUEBKE, D. 2006. All-frequency relighting of glossy objects. *ACM Trans. Graph. 25*, 2, 293–318.

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