

Interactive Relighting with Dynamic BRDFs

Xin Sun* Kun Zhou† Yanyun Chen† Stephen Lin† Jiaoying Shi* Baining Guo†

*State Key Lab of CAD&CG, Zhejiang University

†Microsoft Research Asia



Figure 1: Image sequence of Venus rendered with a dynamic BRDF. Ground truth renderings of the corresponding materials are shown on spheres at the bottom-right of the images. The changes in Venus’ reflectance properties produce corresponding indirect lighting effects on the surrounding walls. Here, precomputation of 12.3 hours with 389 Mbytes memory is used for the 39.7K vertices. The rendering performance is 2.62 fps when the viewpoint, BRDF, and lighting are all simultaneously changing; 10.94 fps when only the viewpoint and BRDF are dynamic; and 57.09 fps when only the viewpoint changes.

Abstract

We present a technique for interactive relighting in which source radiance, viewing direction, and BRDFs can all be changed on the fly. In handling dynamic BRDFs, our method efficiently accounts for the effects of BRDF modification on the reflectance and incident radiance at a surface point. For reflectance, we develop a BRDF tensor representation that can be factorized into adjustable terms for lighting, viewing, and BRDF parameters. For incident radiance, there exists a non-linear relationship between indirect lighting and BRDFs in a scene, which makes linear light transport frameworks such as PRT unsuitable. To overcome this problem, we introduce *precomputed transfer tensors* (PTTs) which decompose indirect lighting into precomputable components that are each a function of BRDFs in the scene, and can be rapidly combined at run time to correctly determine incident radiance. We additionally describe a method for efficient handling of high-frequency specular reflections by separating them from the BRDF tensor representation and processing them using precomputed visibility information. With relighting based on PTTs, interactive performance with indirect lighting is demonstrated in applications to BRDF animation and material tuning.

Keywords: precomputed radiance transfer, bidirectional reflectance distribution function, relighting, global illumination, tensor factorization

† e-mail: {kunzhou, stevelin, bainguo}@microsoft.com

1 Introduction

The visual quality of a computer-generated scene can be appreciably enhanced by simulating the natural propagation of light. Before reaching the viewer, light emitted from illumination sources may undergo a number of physical interactions in a scene, giving rise to rich visual effects such as soft shadows and interreflections. Accounting for the myriad paths of light through a scene and their complex interactions with objects, however, involves a substantial expense in computation.

For interactive relighting, much of this global illumination computation may be performed in advance. Although arbitrary environment lighting is possible at run time, the exit radiance at a surface point may be efficiently evaluated using precomputed radiance transfers (PRT) [Sloan et al. 2002]. Different transfer functions are precomputed for different levels of reflectance effects. For view-independent diffuse reflections, source lighting is transformed by a *transfer vector* into exit radiance from a surface point. For view-dependent glossy reflections, a *transfer matrix* is needed to convert source illumination into a distribution of exit radiance that can be evaluated in the direction of the viewer. These radiance transfer functions are precomputed with respect to given BRDFs in the scene, such that these material properties are essentially baked into the transfer mechanism and cannot be changed at run time.

In this work, we introduce an interactive relighting technique that more generally accommodates changes in reflectance in addition to lighting conditions and viewing directions. The consideration of dynamic BRDFs with indirect lighting brings special challenges. The BRDFs at viewed surface points need to be efficiently evaluated with respect to changing incident light distributions, viewing directions, and reflectance parameters. For this, we present a BRDF tensor representation that can be factorized into adjustable components for each of these dynamic quantities.

A more challenging issue is that incoming radiance at a surface point has a non-linear relationship to BRDFs in a scene [Forsyth and Zisserman 1991]. This non-linearity arises from the fact that distinct light rays may interact with a surface different numbers of times, such that they become transformed dissimilarly by a sur-

face’s BRDF. Because of this non-linearity, dynamic BRDFs cannot be incorporated into PRT or other linear precomputed light transport frameworks in the same manner as dynamic source lighting. Nevertheless, we show that it is possible to decompose the incoming radiance at a surface point into components that can be pre-computed as a function of source radiance and BRDFs in the scene. These components specifically represent indirect lighting along different light paths that arrive at a given surface point. We introduce *precomputed transfer tensors* (PTT) to relate the incident radiance components along each of these paths to source illumination and scene BRDFs. By precomputing PTTs and then evaluating and combining these components at run time, our method efficiently determines the non-linear changes in indirect lighting that result from modifications of BRDF parameters. With this PTT-based formulation, scenes can be interactively rendered with dynamic lighting conditions, viewing directions, and BRDFs.

Aside from the challenges of computing the appearance effects of dynamic BRDFs, the factorization of BRDFs used to facilitate pre-computation also presents a problem. For BRDFs with sharp specular components, numerous basis terms are required for accurate reconstruction, at a considerable expense in performance. To avoid the inefficiencies associated with large bases, our method separates specular reflectance lobes from the BRDFs in the tensor representation to obtain a lower-frequency remainder that can be well modeled with a small basis. The separated high-frequency lighting is processed independently using an efficient procedure based on pre-computed visibility information, and then combined with the PTT results to yield the final rendering.

Our relighting technique has a variety of applications. One is real-time rendering of scenes that contain objects with animated BRDFs. With its interactive performance, our technique also facilitates tuning of material properties rendered with indirect illumination.

2 Related Work

Relighting with Indirect Illumination Early works on efficient relighting with indirect illumination were designed for scenarios with limited user manipulation. For a fixed lighting configuration, viewpoint, and scene geometry, Dorsey *et al.* [1995] utilize the linearity property of light transport to relight scenes with time-variant intensities by adjusting a linear combination of basis images. More recently, several ray tracing techniques have offered greater flexibility in relighting scenes, but without interactive performance. Tabelion and Lamorlette [2004] employed ray tracing with acceleration techniques such as geometry simplification and irradiance caching. An offline ray tracing method was proposed by Meyer and Anderson [2006] for animated sequences, which utilizes statistics from multiple frames rendered with a small number of sampling rays.

Another approach to relighting is to use deep framebuffers as described in [Saito and Takahashi 1990] to cache partially shaded geometry, which is then quickly transferred to the GPU for final shading. For example, given a fixed camera, Gershbein and Hanrahan [2000] precompute geometric information for each pixel such that expensive geometric operations such as visibility computation need not be performed at run time. Also along the same line, Pellacini *et al.* [2005] recently developed a relighting engine that achieves interactive performance while supporting arbitrary appearances. Hašan *et al.* [2006] computed the indirect illumination at a set of view samples from the direct illumination calculated at many gather samples in the scene. All these works capitalize on the assumption that camera movement occurs much less frequently than shading changes. In our relighting system we do not make this assumption. Moreover, in [Pellacini *et al.* 2005] it would be difficult to change surface reflectance because the light shaders need to be manually sim-




Figure 2: Relighting with indirect illumination.

plified and translated into hardware shading languages. In [Hašan *et al.* 2006], dynamic materials are also not well-supported since the distribution of gather samples is dependent on reflectances.

Precomputed Radiance Transfer Methods based on precomputed radiance transfer (PRT) [Sloan *et al.* 2002; Ng *et al.* 2003; Kristensen *et al.* 2005] efficiently render illumination effects such as shadows and interreflections that an object casts onto itself. Although these works generally focus on global illumination with respect to source lighting, changes in BRDFs are possible for direct lighting in [Sloan *et al.* 2002; Ng *et al.* 2004]. However, such modifications would invalidate transfer functions precomputed with the original BRDFs.

BRDF factorization has been employed in a few PRT techniques for efficiently rendering view-dependent reflectance effects with all-frequency lighting [Liu *et al.* 2004; Wang *et al.* 2004; Wang *et al.* 2006]. In these methods, a given BRDF is decomposed into lighting-dependent and view-dependent components, where the lighting-dependent component is included in the precomputed transfer matrix, and the view-dependent component is applied to the PRT surface signal to obtain the outgoing radiance. Factorization and precomputation are done with respect to a fixed BRDF that cannot later be changed. Our technique also separates reflectance factors by BRDF decomposition, but it utilizes a tensor BRDF representation whose factorization enables BRDF modifications at run time. Additionally, we employ a specular lobe separation scheme that elevates the effectiveness of the tensor factorization and significantly improves rendering of shiny BRDFs.

A method for BRDF manipulation with complex lighting was presented in [Ben-Artzi *et al.* 2006]. Outgoing radiance is computed in a PRT framework in which changes to a reflectance parameter are mapped to linear coefficients of precomputed responses. This technique enables a broad range of BRDF editing capabilities, but deals only with a fixed viewpoint and static direct illumination.

BRDF Models BRDFs are 4D functions of incident and outgoing light directions that can vary significantly from material to material. Because of the optical properties of real materials and the physical laws of reflectance, such as energy conservation and reciprocity, real-world BRDFs do not have an arbitrary structure, but lie within a subspace of 4D functions. To represent a subspace of natural BRDFs, many analytic BRDF models, both isotropic [Cook and Torrance 1981; Oren and Nayar 1994] and anisotropic [Kajiya 1985; Ward 1992; Ashikhmin *et al.* 2000], have been proposed. This subspace may be more accurately modeled using sets of measured BRDF data from a large number of real materials [Dana *et al.* 1999; Matusik *et al.* 2003a]. In [Matusik *et al.* 2003b], it was empirically shown that real isotropic BRDFs can be accurately represented using a linear combination of the 100 BRDFs in their acquired data set. Our method also represents BRDFs using a linear basis, which is derived from tensor factorization to facilitate run-time changes in reflectance properties.

Tensor Approximation Recently, tensor approximations [Lathauwer et al. 2000] have been applied for data compression in computer graphics [Vasilescu and Terzopoulos 2004; Wang et al. 2005]. Our method also employs tensor approximation, but it is utilized primarily for efficient rendering, i.e., fast indirect illumination with precomputed tensor data.

3 Relighting with PTTs

The problem of relighting with indirect illumination is illustrated in Fig. 2. A scene is illuminated by distant source radiance described by an environment map. This source radiance propagates through the scene, undergoing global illumination effects such as interreflection and shadowing. Given a fixed spatial configuration of the scene, the spread of illumination is determined by the BRDFs of the objects within. We refer to the direct and indirect illumination that arrives at each viewed point as *transferred incident radiance*. When this transferred incident radiance reaches a surface point, it is reflected according to the point’s BRDF. Evaluating this exit radiance with respect to the BRDF, transferred incident radiance, and view direction yields the rendered appearance of the point.

The goal of relighting is to render the scene with changes in source radiance L and viewing angle ω_o , or more generally, also with changes in scene reflectances H , which we represent as a $n_R \times n_P$ matrix where n_R is the number of distinct dynamic BRDFs in the scene and n_P is the size of a basis used to represent various reflectances. From the rendering equation [Kajiya 1986], the exit radiance of a scene point x can be formulated as the following *relighting equation*:

$$B_x(L, H, \omega_o) = \int_{\Omega} I_x(L, H, \omega_i) f(H_x, \omega_o, \omega_i) (\omega_i \cdot \mathbf{n}_x) d\omega_i. \quad (1)$$

This relighting equation consists of two major components. One is the transferred incident radiance $I_x(L, H, \omega_i)$ arriving at x , which is dependent on the BRDFs in the scene. The other component $f(H_x, \omega_o, \omega_i)$ is the BRDF at x , which relates transferred incident radiance from direction ω_i to exit radiance toward ω_o according to the reflectance H_x of the surface point.

In our algorithm, we precompute the transferred incident radiance I_x with respect to source lighting, scene BRDFs, and incident direction. We also form a tensor representation of the BRDF at surface points x , which is factorized into adjustable terms for incident lighting, viewing direction, and reflectance. These two components of the relighting equation can then be used at run time to rapidly evaluate Eq. 1 for determining the exit radiance at each point toward the viewer. Each of these steps for interactive relighting is described in the following.

3.1 Transferred Incident Radiance

The incident illumination at a surface point may arrive from along many different paths, as illustrated in Fig. 3. Along each path, the total radiance is determined by the BRDF of each surface that interacts with the light. Here, we represent these BRDFs in terms of n_P basis functions computed by PCA on a given BRDF dataset. The BRDF of a surface region r can then be dynamically changed by adjusting the coefficients $H(r, p)$ of each basis function p .

With this basis function representation of BRDFs in the scene, we can further divide the incident radiance along light paths into radiance that corresponds to sequences of BRDF basis functions along each path, e.g., incident radiance that reflects first from a surface region r_1 according to a single BRDF basis function p_1 with coefficient $H(r_1, p_1)$, and then from a region r_2 according to a basis function p_2 with coefficient $H(r_2, p_2)$. For efficient run-time evaluation of indirect lighting, we propose to divide transferred inci-




Figure 3: Paths of indirect illumination with sequences of BRDF basis functions, as modeled by PTTs. The yellow rays show examples of one-bounce interreflections, where $H(r, p)$ denotes an adjustable BRDF coefficient of region r and basis function p . The blue rays illustrate paths with two bounces.

dent radiance into such components, which can be precomputed as a function of source lighting and the BRDF basis function coefficients of the sequence. These precomputed quantities are stored in PTTs, from which the transferred incident radiance components can be rapidly evaluated and combined at run time according to current values of scene BRDFs H and source lighting L .

In processing sequences of BRDF basis functions, we group light paths by the number of interreflections they include, such that the incident radiance is expressed as

$$I_x = I_x^0 + I_x^1 + I_x^2 + \dots$$

where I_x^k is the total transferred incident radiance of light paths with k interreflections. In most scenes, there exist innumerable illumination paths. However, due to surface absorption the magnitude of indirect lighting tends to decrease rapidly for larger values of k , so in practice we consider transferred incident radiance only for $k \leq 2$. Additionally, the frequency characteristics of indirect illumination become considerably lower after each interreflection due to scattering of light, so for further computational efficiency we model the transferred incident radiance of direct lighting I_x^0 in terms of wavelets, of I_x^1 by fourth-order spherical harmonics, and of I_x^2 as a scalar.

Direct lighting I_x^0 can be simply evaluated as a product of the environment map and a PTT $T_x^0(\omega)$ that represents precomputed visibility of lighting from angular direction ω .

As illustrated in Fig. 3, one-bounce interreflection I_x^1 consists of n_R separate light paths with n_P different basis functions for each path. To expedite computation of I_x^1 at run time, we precompute a PTT T_x^1 that relates source radiance to transferred incident radiance with respect to each of these sequences of BRDF basis functions. Specifically, for one-bounce interreflection, we have a 4D tensor $T_x^1(l, r, p, s)$ in terms of source radiance wavelet l , region r , BRDF basis function p , and spherical harmonic s of the transferred incident radiance. This tensor is precomputed by sampling l, r, p values, ray tracing the resultant radiance that arrives at point x , then projecting the radiance onto each spherical harmonic basis function s .

For two-bounce interreflection, we similarly consider light paths that involve $n_R \times n_R$ ordered pairs of regions with $n_P \times n_P$ possible basis function pairs for each path. Since I_x^2 is a scalar, we precompute a 5D PTT $T_x^2(l, r_1, p_1, r_2, p_2)$, where the subscripts index the first and second interreflections.

Local radiance transfers along sequences of three small patches with fixed BRDFs has previously been described in [Aupperle and Hanrahan 1993]. In contrast, global radiance from light source

to scene interreflections to surface point are fully precomputed in PTTs with respect to arbitrary BRDFs for regions in the scene.

3.2 BRDF

For efficient evaluation of BRDFs at a surface point, we present a tensor BRDF representation that can be factorized into adjustable lighting, viewing, and reflectance terms. We model BRDFs from a set of reflectance data $f(h, \omega_o, \omega_i)$ that may be generated from a given reflectance model or taken from measurements of real materials [Matusik et al. 2003a], where h indexes the BRDFs, ω_o is the view direction and ω_i is the incident direction. In representing BRDFs at a surface point, we discretize $f(h, \omega_o, \omega_i)$ as a 3rd-order tensor $F \in \mathbb{R}^{N_h \times N_{\omega_o} \times N_{\omega_i}}$ where N_h , N_{ω_o} and N_{ω_i} are the sampling resolutions of h , ω_o and ω_i respectively. By decomposing it into factors that depend only on h , ω_o and ω_i using tensor factorization [Lathauwer et al. 2000], F can be approximated as

$$F \approx D \times_h U^h \times_{\omega_o} U^{\omega_o} \times_{\omega_i} U^{\omega_i},$$

from which a given BRDF value can then be calculated as

$$f(h, \omega_o, \omega_i) \approx \sum_{i=1}^{n_h} \sum_{j=1}^{n_{\omega_o}} \sum_{k=1}^{n_{\omega_i}} D(i, j, k) U^h(h, i) U^{\omega_o}(\omega_o, j) U^{\omega_i}(\omega_i, k). \quad (2)$$

Here, $D \in \mathbb{R}^{n_h \times n_{\omega_o} \times n_{\omega_i}}$ is the BRDF tensor core, where $n_h \leq N_h$, $n_{\omega_o} \leq N_{\omega_o}$, and $n_{\omega_i} \leq N_{\omega_i}$ are the resolutions of the reflectance parameters, view direction, and incident direction after factorization. The reflectance factor U^h , the view factor U^{ω_o} , and the light factor U^{ω_i} are $N_h \times n_h$, $N_{\omega_o} \times n_{\omega_o}$, and $N_{\omega_i} \times n_{\omega_i}$ matrices respectively.

With this tensor approximation, we gain significant compression of the BRDF set F , and reflectance values can be readily computed for arbitrary incident light distributions, reflectance functions, and viewpoints. In contrast to the usual view-light SVD factorization [Kautz and McCool 1999], which would need to be done for each BRDF, this trilinear factorization leads to greater savings in memory because of the additional approximation along the BRDF dimensions. Moreover, the utilized tensor approximation method of [Lathauwer et al. 2000] provides the best trilinear factorization in terms of RMS error, and can generate compression results of high perceptual quality as shown in [Vasilescu and Terzopoulos 2004; Wang et al. 2005].

We note that the separate use of a factorized BRDF representation for surface shading and a PCA representation for indirect lighting in PTTs is motivated by efficiency considerations. For shading, computation of the exit radiance involves an integration over the incident lighting directions, which can be evaluated more efficiently with the reduced lighting dimension n_L of the tensor BRDF than with the full dimension N_L of a BRDF in a PCA basis. For PTTs, the tensor BRDF would be computationally inefficient because of the need to reconstruct the BRDF for many outgoing directions. Furthermore, since lighting and viewing need not be decoupled in the indirect lighting computation, we can take advantage of the greater accuracy of a PCA basis, which does not compress and degrade lighting and viewing information.

3.3 Relighting Computation

With PTTs for transferred incident radiance and the tensor decomposition for BRDFs, the relighting equation can be rapidly evaluated at run time for arbitrary distant source radiance, scene BRDFs, and viewing directions. Let the source radiance L be represented by a Haar wavelet basis $\{L_l \mid 1 \leq l \leq N_L\}$ such that $L = \sum C_l^L L_l$ [Ng et al. 2003]. We rewrite the relighting equation of Eq. 1 using the

BRDF tensor approximation of Eq. 2 to obtain the exit radiance

$$B_x(L, H, \omega_o) \approx \sum_{i=1}^{n_h} \sum_{j=1}^{n_{\omega_o}} \sum_{k=1}^{n_{\omega_i}} D(i, j, k) U^h(H, i) U^{\omega_o}(\omega_o, j) \sum_{l=1}^{N_L} C_l^L U_x^l(l, k, H). \quad (3)$$

The *incidence component* U_x^l represents the transferred incident radiance at x after modulation by the BRDF lighting factor U^{ω_i} . We compute it at run time as separate terms in correspondence with PTTs:

$$U_x^l = U_x^{I0} + U_x^{I1} + U_x^{I2} \quad (4)$$

where U_x^{Iq} denotes the incidence component for lighting that has undergone q interreflections before arriving at x . Each of these terms is evaluated as follows.

Direct Lighting In the computation of direct lighting, transferred incident radiance is determined by the light source visibility in PTT T_x^0 . For computational convenience, we in practice precompute a product $\tilde{T}_x^0(\omega, k) = U^{\omega_i}(\omega, k)(\omega \cdot \mathbf{n}_x) T_x^0(\omega)$ of visibility and light factors as in [Liu et al. 2004; Wang et al. 2004], such that the incidence component can be more efficiently computed as

$$U_x^{I0}(l, k) = \int_{\Omega} L_l(\omega) \tilde{T}_x^0(\omega, k) d\omega,$$

where L_l represents the l -th wavelet basis function of source illumination. Since direct lighting does not interact with the scene before arriving at x , it is independent of the scene's BRDFs.

One-Bounce Interreflection In the run-time evaluation of the incidence component for one-bounce interreflection, products between the transfer tensor T^1 and the BRDF coefficients of each region are computed and added together to obtain the transferred incident radiance, which is then mapped onto the BRDF light factors through an inner product with a matrix C^{ω_i} :

$$U_x^{I1}(l, k, H) = \sum_{s=1}^S C^{\omega_i}(s, k) \sum_{r=1}^{n_R} \sum_{p=1}^{n_P} H(r, p) T_x^1(l, r, p, s)$$

where

$$C^{\omega_i}(s, k) = \int_{\Omega} L_s^{SH}(\omega) U^{\omega_i}(\omega, k) (\omega \cdot \mathbf{n}_x) d\omega.$$

$H(r, p)$ denotes the adjustable BRDF coefficient of region r and basis function p . Like T^1 , matrix C^{ω_i} is a quantity precomputed in terms of spherical harmonics.

The one-bounce transferred incident radiance could in principle be directly integrated with the lighting-dependent factors of the BRDF tensor, rather than being first projected onto the SH basis. However, the lighting-dependent factors are computed specifically for BRDF representation and do not provide an adequate reconstruction of the incident radiance distribution, which is needed for handling specular lobes as described in Sec. 4.1. In addition, use of the SH basis allows for a clean decoupling of the radiance transfer and surface reflection components of relighting, such that either could be replaced if desired.

Two-Bounce Interreflection For two-bounce interreflection, we process radiance transferred through ordered pairs (r_1, r_2) of scene regions. Since transferred incident radiance after two bounces is considered to arrive at x in a uniform distribution, the incidence component is

$$U_x^{I2}(l, k, H) = C^{\omega_i}(k) \sum_{r_1=1}^{n_R} \sum_{p_1=1}^{n_P} \sum_{r_2=1}^{n_R} \sum_{p_2=1}^{n_P} H(r_1, p_1) H(r_2, p_2) T_x^2(l, r_1, p_1, r_2, p_2)$$

where

$$C^{\omega_i}(k) = \int_{\Omega} U^{\omega_i}(\omega, k) (\omega \cdot \mathbf{n}_x) d\omega.$$




Figure 4: *Specular lobe separation. From left to right: Original; Without specular lobe; 0° , 7° , 14° , and 21° specular lobe components.*

4 Implementation Issues

4.1 Specular Lobe Separation

The high-frequency specular lobes of many BRDFs necessitate a large number of basis terms for accurate reconstruction, but the use of a large basis precludes interactive performance. For improved handling of specular effects, we separate high-frequency specular lobes from the BRDF tensor and rapidly process them using pre-computed visibility, which is represented in the spatial domain by $64 \times 64 \times 6$ cubemaps. This separation divides BRDFs into a specular lobe f_s and the remainder f_r :

$$f(h, \omega_o, \omega_i) = f_s(h, \omega_o, \omega_i) + f_r(h, \omega_o, \omega_i).$$

The specular lobe f_s is modeled as a sum of Gaussians, each with different neighborhood support:

$$f_s(h, \omega_o, \omega_i) = \sum_{k=0^\circ, 7^\circ, 14^\circ, 21^\circ} f_s^{(k)}(h, \omega_o, \omega_i)$$

where $f_s^{(k)}$ is a univariate Gaussian function of which three standard deviations covers a k -degree solid angle centered around the mirror reflection direction of ω_o . The values of k are sampled according to the cubemap resolution of the incident radiance, such that 0° represents perfect mirror reflection and each successively broader Gaussian covers an incrementally larger neighborhood of cubemap pixels, as illustrated in Fig. 4. The separation of $f_s^{(k)}$ components from the BRDF may be computed using an Expectation-Maximization algorithm [Dempster et al. 1977] under the constraint that each Gaussian is centered around the mirror reflection direction and has a fixed neighborhood support, i.e., $k = 0^\circ, 7^\circ, 14^\circ, 21^\circ$. The removal of f_s from the BRDFs leaves a relatively low-frequency remainder f_r that can be more effectively compressed by tensor approximation.




Figure 5: *BRDF tensor approximation errors for different n_h and n_{ω_i} .*

The BRDF tensor is then constructed from f_r , while the effects of f_s are separately evaluated. For direct illumination, we multiply each




Figure 6: *Reconstruction results from the tensor BRDF with different numbers of specular lobe components. From left to right: Ground Truth; Without specular lobe separation; With specular lobe separation for one (0°), two ($0^\circ, 7^\circ$), three ($0^\circ, 7^\circ, 14^\circ$), and four ($0^\circ, 7^\circ, 14^\circ, 21^\circ$) components.*

term $f_s^{(k)}$ of the specular lobe with a corresponding environment map $E^{(k)}$, prefiltered with a Gaussian of solid angle $k = 0^\circ, 7^\circ, 14^\circ$, or 21° . The products are aggregated to obtain the specular lobe component. For indirect lighting, the specular lobe is computed with respect to transferred incident radiance from one-bounce inter-reflection, which can be rapidly converted into the spatial domain using prefiltered SH basis functions. The uniform incident radiance from two-bounce interreflection is disregarded. With the prefiltered environment maps, we utilize the fractional visibility over the corresponding solid angles, which is computed from the visibility maps of each vertex.

BRDF tensor approximation errors with specular lobe separation are shown in Fig. 5 for a set of 100 measured BRDFs [Ngan et al. 2005] represented using the Cook-Torrance model [Cook and Torrance 1981]. For the finest sampling of $n_h = 6$ and $n_{\omega_i} = 32$, the squared error of tensor approximation without specular lobe separation is 0.952465, which is substantially higher than for coarse sampling with specular lobe separation. Rendering comparisons are exhibited in Fig. 6 for BRDFs with various levels of shininess. As shown in the figure, a greater number of lobe components leads to better approximations of BRDFs.

In related work, [Green et al. 2006] approximates high-frequency transfer components by a sum of Gaussians, which are computed separately from low-frequency components. These transfer functions are precomputed with respect to fixed scene BRDFs, and are therefore not intended for relighting with dynamic BRDFs. In our tensor representation, we do not use the half-angle parameterization of [McCool et al. 2001], because lighting and viewing must be decoupled for real-time surface shading. Since a half-angle parameterization cannot be employed in our technique, we utilize specular lobe separation to reduce error in the tensor approximation.

4.2 BRDF Navigation

For modifying BRDFs, we employ navigation schemes tailored to the BRDF model being used. We have utilized the Blinn-Phong [Blinn 1977] and the Cook-Torrance models to represent specular BRDF components, with the diffuse components represented separately. For the Blinn-Phong model, the set of BRDFs is generated by sampling the shininess parameter at integer values from 1 to 200. At run time, a BRDF can be selected for an arbitrary shininess value by linear interpolation of the sampled BRDFs.

For the Cook-Torrance model, we form the set of BRDFs by fitting




Figure 7: Measured BRDFs represented using the Cook-Torrance model.

to measured data of 100 real materials [Ngan et al. 2005], plotted in Fig. 7 with respect to the Fresnel and surface roughness parameters. After triangulation of these samples, BRDF navigation is done by selecting an arbitrary point within the convex hull, and then determining the BRDF by barycentric interpolation of the measured BRDFs according to the position of the point in a triangle. For both the Cook-Torrance and Blinn-Phong models, we also utilize adjustable weighting coefficients for the specular and diffuse components in the navigation.

4.3 Implementation Settings and Computational Costs

Due to the large sampling space of BRDFs, illuminations, viewing directions, and light paths, a number of tradeoffs between accuracy and performance need to be made. We chose our implementation settings primarily based upon those used in previous works, e.g., a one-bounce interreflection approximation was employed in [Tabellion and Lamorlette 2004] for computer generated films, and a 4th-order SH lighting basis was used in [Kristensen et al. 2005] to approximate indirect lighting.

Tensor BRDFs For light reflected at the surface point, specular components of BRDFs are represented by the BRDF tensor, while diffuse reflection is handled independently of the tensor decomposition. In our implementation, the BRDF tensor parameters are set to $n_h = 4$ and $n_{\omega_i} = n_{\omega_o} = 8$. The BRDF lighting and viewing factors U^{ω_i} and U^{ω_o} are parameterized according to an orthographic projection from the upper hemisphere onto the unit disk as in [Wang et al. 2004].

The BRDF tensor approximation parameters in our implementation work well for the two isotropic BRDF sets used in this paper: the Blinn-Phong model, and the 100 fitted Cook-Torrance BRDFs which sample various real-world reflectances [Ngan et al. 2005]. In an analysis of the BRDF factorization for the Cook-Torrance BRDFs, we quantitatively tested all combinations of $n_{\omega_i} \times n_h = \{4, 8, 12, 16, 20, 24, 28, 32\} \times \{1, 2, 3, 4, 5, 6\}$. The squared error of each configuration is reported in Fig. 5. We set n_h to 4 because the error curves converge at 4, and because a 4-term approximation has been shown to be visually adequate for materials with moderate specular components [Wang et al. 2004]. Based on a similar analysis, we set n_h to 3 for the Blinn-Phong BRDF set, which has lower variability. Since the two BRDF data sets are isotropic, we set n_{ω_o} to be equal to n_{ω_i} . In our implementation, we chose $n_{\omega_i} = n_{\omega_o} = 8$ mainly for performance considerations and its reasonable level of accuracy. For materials with highly specular components, our specular lobe decomposition significantly improves the approximation accuracy as shown in Fig. 6. The tensor compression time for the examples in our paper is less than one hour.

Indirect-lighting BRDFs For the PCA basis used for BRDF representation in PTTs, we utilized five basis functions (four specular and one diffuse) for both the Cook-Torrance and Blinn-Phong




Figure 8: One-bounce interreflections with different lighting resolutions. From left to right, n_L of $6 \times 4 \times 4$, $6 \times 8 \times 8$, and $6 \times 16 \times 16$. We consider a resolution of $6 \times 4 \times 4$ to be too low, since the caustic on the table is too dim. For caustics results that also include direct lighting and two-bounce interreflection, please refer to Fig. 13.

BRDFs, which give a squared reconstruction error of $6.23e-4$ and $1.61e-7$, respectively. This error level is similar to that of the tensor approximation, even though the PCA basis includes specular lobe components.

Direct Lighting For each vertex, the incidence component U_x^{I0} is an $N_L \times n_{\omega_i}$ matrix. To handle high-frequency lighting, N_L is very large, $6 \times 64 \times 64$ in our implementation. For efficient storage and rendering, we compress this data by maintaining for each ω_i only $n_L = 96$ wavelet coefficients with the largest area-weighted magnitude [Ng et al. 2003], which provides an optimal approximation in the L^2 -norm. The resulting light transport matrix is then compressed as in [Liu et al. 2004; Wang et al. 2004] with little loss in rendering quality.

The visibility map of each vertex is also sampled at $6 \times 64 \times 64$ directions, and is stored uncompressed.




Figure 9: Clustered tensor approximation errors for different reduced lighting dimensions and eigenvectors per cluster.




Figure 10: Rendering results with various CTA compression levels (number of eigenvectors / lighting dimension). Top row: 8/64, 16/64, 32/16, 32/32. Bottom row: 32/64, 32/96, 64/64, ground truth.

One-bounce Interreflection For each vertex, the one-bounce PTT is a $n_R \times n_P \times N_L \times S$ tensor where $n_R \leq 3$ and $n_P = 5$ in our imple-




Figure 11: Ground truth comparison of PTT-based rendering. Left: ray tracing result; Middle: direct lighting result using our method; Right: our result including interreflections.

mentation. To reduce its large size, we first decrease N_L by using a lower resolution ($n_L = 6 \times 8 \times 8$) representation of the environment map that gives a fair approximation, as illustrated in Fig. 8. Even with this reduction in N_L , the tensor data is still large, so we employ clustered tensor approximation for further compression. In contrast to the method presented in [Tsai and Shih 2006], which iterates between clustering and tensor approximation, we instead employ a faster scheme that applies CPCA [Sloan et al. 2003] to cluster the PTTs of vertices into a number of groups, and then collects the PTTs of each group into a single tensor by adding a vertex dimension. Each of these group tensors is then compressed by tensor approximation along only the vertex and lighting dimensions, while the other tensor dimensions are not compressed because of their already small size.

With this technique, a significant savings in precomputation time and storage of one-bounce PTTs is achieved beyond the clustered tensor approximation in [Tsai and Shih 2006]. Although one-bounce interreflections are represented in a low-frequency SH basis, some high-frequency effects are nevertheless captured, since interreflections are raytraced using a PCA basis for BRDFs, and a specular lobe component is modeled. Tensor approximation errors and rendering results are shown for different clustering parameters in Fig. 9 and Fig. 10, respectively. In our implementation, we use 32 eigenvectors and a lighting dimension of 64. Lighting dimensions below 64 have little impact on performance, which is affected more significantly by the number of eigenvectors. High-frequency lighting from one-bounce interreflections are demonstrated for caustics in Section 5.

Two-bounce Interreflection For the 5D two-bounce PTTs T_x^{I2} , we obtain a substantial reduction in data size by using an even coarser ($n_L = 6 \times 2 \times 2$) environment map. In addition to this simplification of the source radiance, we employ CPCA in a manner similar to that for one-bounce interreflections.

Since the transferred incident radiance of two-bounce interreflections is treated as having a uniform distribution, high-frequency two-bounce lighting cannot be generated in our current implementation. Although indirect lighting of more than two interreflection bounces can contribute to scene appearance, these visual effects are often subtle. In principle, our framework could be extended to include indirect lighting with more interreflections and greater precision, at a cost of greater computation.




Figure 12: Indirect lighting in an indoor scene. Top row: direct lighting only; plus one- and two-bounce interreflections. Bottom row: after different BRDF adjustments for the floor and sofa.

Precomputation The precomputation of PTTs is substantially accelerated by employing interpolation techniques and recursive computation. For one-bounce PTTs, we trace rays over the hemisphere of each vertex, and when a ray intersects a surface point, the point’s incident radiance from direct lighting is used for evaluating its one-bounce contribution. To facilitate processing, the visibility of the intersection point is barycentrically interpolated from the vertices of its mesh triangle. In the same manner that one-bounce PTTs utilize direct lighting, two-bounce PTTs use one-bounce PTTs to expedite precomputation.

Rendering Performance In run-time rendering, an approximate breakdown of computation for the examples in this paper when lighting, viewpoint, and BRDFs are all dynamic is as follows: direct lighting 25%, one-bounce 70%, two-bounce and specular lobe 5%. Performance can be roughly considered as linearly proportional to scene complexity. For direct lighting, performance is linearly dependent on the number of vertices. For one-bounce interreflections, performance depends on the number of tensor approximation clusters and the reduced tensor dimensions in the clusters. For two-bounce interreflections, performance is a function of the number of eigenvectors and the number of clusters.

The rendering and precomputation performance is also dependent on the number of adjustable BRDFs in the scene. For one-bounce interreflections, there exists a linear relationship, while for two bounces the performance is linear to the squared number of adjustable BRDFs.

5 Results

We implemented our system on a workstation with a dual 3.20 GHz Xeon CPU, 4.0 Gbytes of RAM, and a Nvidia GeForce 7800 GT graphics card. The computation of the transferred incident radiance is performed on the CPU, and the final shading is done on the GPU. In the following, we present scenes rendered with changing BRDFs, lighting conditions, and view directions.

Fig. 11 demonstrates the visual quality of PTT-based relighting in comparison to direct illumination and ground truth ray tracing. With one- and two-bounce interreflections, a close approximation of the ray tracing result is obtained, including the color bleeding of the surface onto the bird’s chest.

The indoor scene of Fig. 12 contains twelve different materials,




Figure 13: Airplane rendered using Blinn-Phong BRDFs with constant diffuse/specular weights. Top row: change in specular exponent of the base. Bottom row: change in specular exponent of the airplane.

some of them with mapped textures. Renderings with and without indirect lighting are shown. With identical source radiance but different BRDFs for the floor and sofa, the resulting change in specular reflections and indirect illumination, as exhibited in the color bleeding effects, dramatically alters the appearance of the scene.

The PCA basis of BRDFs in PTTs, together with the BRDF tensor approximation scheme, offers sufficient accuracy to generate specular caustics, as shown in the top row of Fig. 13. Changes in interreflection effects are also shown in the bottom row. When the airplane is rendered with a relatively diffuse material, the color bleeding onto the table surface exhibits the red color of the plane. By increasing the shininess of the airplane, greater specular reflection from the plane transforms the interreflection colors on the surface towards that of the white light source.

The wristwatch scene shown in Fig. 14 highlights the effects of specular lobe separation, and demonstrates the utility of our method for material design and BRDF editing. The four watches made of different metallic materials exhibit differences in interreflection.

In Fig. 15, we show a chess scene in which the pieces of each side are collectively treated as an adjustable BRDF region. With illumination from the far side of the board, there exist dark shadows on the pieces when viewed with only direct lighting. With the inclusion of one-bounce and two-bounce interreflections from PTTs, the shadowed sides of the pieces become illuminated with indirect lighting reflected from other pieces and the chessboard, bringing clear visibility within these shadowed areas.

The Venus sequence in Fig. 1 demonstrates the application of our technique to BRDF animation. We animate the material of the sculpture by navigating among the measured BRDFs of Fig. 7 along a curve in the Fresnel-roughness space. Corresponding indirect illumination effects can be seen on the surrounding wall. Please notice not only the changes in reflected colors on the wall, but also the differences in the distribution of interreflected light that result from Venus’s reflectance properties.



Figure 14: BRDF editing with indirect illumination effects. The wristwatches are composed of different metallic materials.

Table 1 lists memory costs and performance statistics for each of these examples. The precomputation time depends primarily on the number of vertices, while the compression time is additionally related to the number of light paths. Since the ‘Venus’ scene contains just a single region with dynamic material, it requires relatively little compression time. The ‘Chess’ and ‘Watch’ scenes consist mostly of regions without texture, so to save memory we compress only a single color channel for light paths that contain only diffuse reflections or non-dynamic BRDFs.

6 Conclusion

The nonlinear relationship of scene appearance and BRDFs makes handling of interreflections highly non-trivial, and to address this problem we proposed the precomputed transfer tensor framework. This method enables interactive relighting generalized to include dynamic BRDFs. With PTTs, a factorized tensor representation of BRDFs, and separate processing of specular lobes, our technique not only generates compelling transformations in the appearance of objects, but also realistically produces the resulting indirect lighting effects in the surrounding scene.

Relighting with dynamic BRDFs, viewpoint, and illumination involves a tremendously large sampling space. In this paper, we showed that this relighting task can be accomplished with manageable memory and computation costs that are in line with related techniques:

- ◊ Performance: Rendering speed is 1.5-5.9 fps in our examples when BRDF, lighting, and viewpoint are *all* dynamic. In BRDF adjustment applications such as scene editing, it is often the case that only one of these factors is being adjusted at a time. As shown in Table 1, considerably higher performance is attainable in such cases. We note that our rendering algorithm is well-suited to parallel implementation on multiple CPUs, which would substantially elevate run-time performance.
- ◊ Precomputation: With a parallel implementation, precomputation time could also be significantly accelerated. Parallel processing has been exploited in other PRT methods such as the relighting

Scene	BRDF Space	Adjust. BRDFs	Vertices	Clusters	Eigen-vectors	Storage (Mbytes)	Precomp (hrs)	Compress (hrs)	V (fps)	V&B (fps)	V,B&L (fps)
Tweety (Fig. 11)	CT	2	13.0K	13	32	148	2.6	4.0	121.58	28.73	5.82
Airplane (Fig. 13)	BP	2	13.1K	13	32	144	2.5	6.4	139.50	30.79	5.91
Indoor (Fig. 12)	BP	2	35.7K	35	32	311	4.2	15.1	74.52	12.57	3.40
Venus (Fig. 1)	CT	1	39.7K	40	16	389	5.0	7.3	57.09	10.94	2.62
Chess (Fig. 15)	CT	2	53.9K	54	16	539	13.7	10.5	48.30	8.23	1.67
Watch (Fig. 14)	CT	2	62.1K	62	16	603	15.7	14.3	43.06	7.23	1.53

Table 1: Performance on exhibited examples. For the BRDF space, ‘BP’ refers to Blinn-Phong, and ‘CT’ to Cook-Torrance. ‘V’ refers to performance when only the viewing direction is dynamic; ‘V&B’ is for when only viewpoint and BRDF are changing; and ‘V,B&L’ refers to view, BRDF, and lighting all being dynamic. The lighting basis is set to 64 for all examples.




Figure 15: Chess scene. Top: with direct lighting only; Bottom: with inter-relations from PTTs.

technique in [Kristensen et al. 2005], which precomputes a scene with 43K triangles in four hours on a cluster of 32 2.8GHz PCs. Our precomputation time would be even lower on such a cluster.

- Memory: Given the generality of our technique in handling adjustable BRDFs, lighting, and viewpoint with high frequency effects, the memory size is understandably large, but is comparable to that of other all-frequency PRT relighting techniques, including [Ng et al. 2003; Ng et al. 2004; Wang et al. 2004].

In future work, we plan to examine ways to further decrease the memory and precomputation time needed for PTTs. For applications in which full generality in relighting is not needed, such as scenarios where the viewpoint or lighting conditions are fixed or limited, considerable reductions in PTT size and precomputation may be possible. We also intend to expand the functionality of our relighting technique. One interesting direction is to efficiently handle dynamic geometry in addition to BRDFs, such that BRDF animation can be applied to object morphing. We also plan to extend our method to additionally allow changes in surface mesostructure, by incorporating BTF data into our framework.

Acknowledgements

We would like to thank the reviewers for their valuable comments, and Kit Thambiratnam for his help in video production. Xin Sun and Jiaoying Shi were partially supported by the 973 Program of China (No.2002CB312105).

References

- ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. 2000. A microfacet-based BRDF generator. In *SIGGRAPH '00*, 65–74.
- AUPPERLE, L., AND HANRAHAN, P. 1993. A hierarchical illumination algorithm for surfaces with glossy reflection. In *SIGGRAPH '93*, 155–162.
- BEN-ARTZI, A., OVERBECK, R., AND RAMAMOORTHI, R. 2006. Real-time brdf editing in complex lighting. *ACM Trans. Graph.* 25, 3, 945–954.
- BLINN, J. F. 1977. Models of light reflection for computer synthesized pictures. In *SIGGRAPH '77*, 192–198.
- COOK, R. L., AND TORRANCE, K. E. 1981. A reflectance model for computer graphics. In *SIGGRAPH '81*, 307–316.
- DANA, K. J., VAN GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. 1999. Reflectance and texture of real-world surfaces. *ACM Trans. Graph.* 18, 1, 1–34.
- DEMPSTER, A. P., LAIRD, N. M., AND RUBIN, D. B. 1977. Maximum-likelihood from incomplete data via the em algorithm. *J. Royal Statistical Society* 39, 1, 1–38.
- DORSEY, J., ARVO, J., AND GREENBERG, D. 1995. Interactive design of complex time dependent lighting. *IEEE Comp. Graph. and Appl.* 15, 2, 26–36.
- FORSYTH, D. A., AND ZISSERMAN, A. P. 1991. Reflections on shading. *IEEE Trans. Pattern Analysis and Machine Intelligence* 13, 7, 671–679.
- GERSHBEIN, R., AND HANRAHAN, P. 2000. A fast relighting engine for interactive cinematic lighting design. In *SIGGRAPH '00*, 353–358.
- GREEN, P., KAUTZ, J., MATUSIK, W., AND DURAND, F. 2006. View-dependent precomputed light transport using nonlinear gaussian function approximations. In *Symp. 13D*.
- HAŠAN, M., PELLACINI, F., AND BALA, K. 2006. Direct-to-indirect transfer for cinematic relighting. *ACM Trans. Graph.* 25, 3, 1089–1097.
- KAJIYA, J. T. 1985. Anisotropic reflection models. In *SIGGRAPH '85*, 15–21.
- KAJIYA, J. T. 1986. The rendering equation. In *SIGGRAPH '86*, 143–150.
- KAUTZ, J., AND MCCOOL, M. D. 1999. Interactive rendering with arbitrary brdfs using separable approximations. In *Eurographics Rendering Workshop*, 281–292.
- KRISTENSEN, A. W., AKENINE-MÜLLER, T., AND JENSEN, H. W. 2005. Precomputed local radiance transfer for real-time

- lighting design. *ACM Trans. Graph.* 24, 3, 1208–1215.
- LATHAUWER, L. D., MOOR, B. D., AND VANDEWALLE, J. 2000. On the best rank-1 and rank- (r_1, r_2, \dots, r_n) approximation of higher-order tensors. *SIAM J. Matrix Anal. Appl.* 21, 4, 1324–1342.
- LIU, X., SLOAN, P., SHUM, H.-Y., AND SNYDER, J. 2004. All-frequency precomputed radiance transfer for glossy objects. In *Eurographics Symp. Rendering*, 337–344.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. In *SIGGRAPH '03*, 759–769.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. Efficient isotropic brdf measurement. In *Eurographics Symp. Rendering*, 241–247.
- MCCOOL, M. D., ANG, J., AND AHMAD, A. 2001. Homomorphic factorization of brdfs for high-performance rendering. In *SIGGRAPH '01*, 171–178.
- MEYER, M., AND ANDERSON, J. 2006. Statistical acceleration for animated global illumination. *ACM Trans. Graph.* 25, 3, 1075–1080.
- NG, R., RAMAMOORTHY, R., AND HANRAHAN, P. 2003. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Trans. Graph.* 22, 3, 376–381.
- NG, R., RAMAMOORTHY, R., AND HANRAHAN, P. 2004. Triple product wavelet integrals for all-frequency relighting. *ACM Trans. Graph.* 23, 3, 477–487.
- NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of brdf models. In *Eurographics Symp. Rendering*, 117–226.
- OREN, M., AND NAYAR, S. K. 1994. Generalization of lambert's reflectance model. In *SIGGRAPH '94*, 239–246.
- PELLACINI, F., VIDIMCE, K., LEFOHN, A., MOHR, A., LEONE, M., AND WARREN, J. 2005. Lpics: a hybrid hardware-accelerated relighting engine for computer cinematography. *ACM Trans. Graph.* 24, 3, 464–470.
- SAITO, T., AND TAKAHASHI, T. 1990. Comprehensible rendering of 3-d shapes. In *SIGGRAPH '90*, 197–206.
- SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *SIGGRAPH '02*, 527–536.
- SLOAN, P.-P., HALL, J., HART, J., AND SNYDER, J. 2003. Clustered principal components for precomputed radiance transfer. *ACM Trans. Graph.* 22, 3, 382–391.
- TABELLION, E., AND LAMORLETTE, A. 2004. An approximate global illumination system for computer generated films. *ACM Trans. Graph.* 23, 3, 469–476.
- TSAI, Y.-T., AND SHIH, Z.-C. 2006. All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. *ACM Trans. Graph.* 25, 3, 967–976.
- VASILESCU, M. A. O., AND TERZOPOULOS, D. 2004. Tensor-textures: multilinear image-based rendering. In *SIGGRAPH '04*, 336–342.
- WANG, R., TRAN, J., AND LUEBKE, D. 2004. All-frequency relighting of non-diffuse objects using separable brdf approximation. In *Eurographics Symp. Rendering*, 345–354.
- WANG, H., WU, Q., SHI, L., YU, Y., AND AHUJA, N. 2005. Out-of-core tensor approximation of high dimensional visual data. In *SIGGRAPH '05*, 527–535.
- WANG, R., TRAN, J., AND LUEBKE, D. 2006. All-frequency relighting of glossy objects. *ACM Trans. Graph.* 25, 2, 293–318.
- WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *SIGGRAPH '92*, 265–272.