A Practical Extension to Microfacet Theory for the Modeling of Varying Iridescence

Laurent Belcour, Pascal Barla

A Practical Extension to Microfacet Theory for the Modeling of Varying Iridescence

LAURENT BELCOUR, Unity Technologies

PASCAL BARLA, Inria

Fig. 1. Material appearance such as that of leather is usually reproduced with microfacet models in computer graphics. A more realistic result is achieved by adding a thin-film coating that produces iridescent colors [Akin 2014]. We replace the classic Fresnel reflectance term with a new Airy reflectance term that accounts for iridescence due to thin-film interference. Our main contribution consists in an analytical integration of the high-frequency spectral oscillations exhibited by Airy reflectance, which is essential for practical rendering in RGB. For the leather material on the CHAIR model, we used a thin film of index $\eta_2 = 1.3$ and thickness $d = 290$nm, over a rough dielectric base material ($\alpha = 0.2$, $\eta_3 = 1$). When the scene is rotated, goniochromatic effects such as subtle purple colors may be observed at grazing angles.

In this work, we introduce an extension to microfacet theory for the rendering of iridescent effects caused by thin-films of varying thickness (such as oil, grease, alcohols, etc) on top of an arbitrarily rough base layer. Our material model is the first to produce a consistent appearance between tristimulus (e.g., RGB) and spectral rendering engines by analytically pre-integrating its spectral response. The proposed extension works with any microfacet-based model: not only on reflection over dielectrics or conductors, but also on transmission through dielectrics. We adapt its evaluation to work in multi-scale rendering contexts, and we expose parameters enabling artistic control over iridescent appearance. The overhead compared to using the classic Fresnel reflectance or transmittance terms remains reasonable enough for practical uses in production.

CCS Concepts: • **Computing methodologies → Reflectance modeling**;

Additional Key Words and Phrases: Iridescence, SV-BRDF model, spectral aliasing, thin-film interference

1 INTRODUCTION

A surface is called iridescent when its color changes when viewed or lit from different directions. Such *goniochromatic* effects are due to interference between light waves that are scattered in a wavelength-dependent way, hence yielding rich color variations. Iridescent appearance is common in nature, as with birds, insects, snakes, and even some fruits; but it also occurs in man-made products such as with oil leaks, window defects, soap bubbles or car paints. Some iridescence effects are more subtle and may even go unnoticed to the untrained eye: these include traces of grease or alcohol (e.g., finger traces on kitchen appliance) or finishes to protect base materials (e.g., leathers, metals). Yet such subtle details are essential to reproduce the look and feel of real-world materials in computer graphics imagery [Akin 2014].

Two causes of goniochromism are to be distinguished: diffraction produced by light reflection on microscopic structures at a scale similar to the visible wavelengths; interferences produced by light interaction with films of thickness close to the visible wavelengths.

In this paper we focus on iridescence due to thin-films of *varying thickness*. In practice, thickness may be varied directly by artists during editing sessions, or across a surface to reproduce the appearance of traces or irregular finishes for instance. Formally, we

model iridescence using a dielectric thin-film laid on top of a base material (dielectric or conductor) of arbitrary roughness (Section 3). The main issue with such a configuration is that reflected radiance produces oscillations in the spectral domain, which require a dense sampling of wavelengths to avoid *spectral aliasing*. This is illustrated in Figure 2: the color fringes due to a dielectric thin-film on top of a metallic base layer are well-reproduced only if the number of spectral samples is large enough. However, such a numerical integration is impractical: for instance, when the film thickness varies spatially, integration should be performed at each and every point of the surface. Precomputing reflectance colors for all possible variations of material parameters would not be a viable solution either: as explained later on, this would require high-dimensional lookup tables with very high resolutions.

Our main contribution is an *analytical spectral integration* formula for reflectance due to thin-film iridescence (Section 4). It works as an extension of micro-facet theory, hence providing a modular solution to the use of thin-films on arbitrarily rough base materials. The resulting iridescent material model has many practical advantages (Section 5): it is fast to evaluate in RGB space while providing reflectance very close to brute-force spectral rendering; it gives interactive feedback for the exploration of physical parameters; and it is easily adapted to multi-scale rendering.

Our approach relies on Airy summation, which correctly models the reflectance due to a thin-film, including multiple scattering, polarization and phase changes for both conductor and dielectric base materials. Even though this equation has been known for decades in the physics literature, its adaptation to the constraints and demands of computer graphics requires a different approach. We show that our graphics-oriented model approximates the physical ground truth with unprecedented accuracy across different types of rendering engines (Section 6), which makes it particularly adapted to the faithful previsualization of high-quality renderings.

2 RELATED WORK

Reflectance models in computer graphics are often based on the micro-facet theory [Cook and Torrance 1982] where color is solely due to the wavelength-dependent Fresnel reflectance term, which itself depends on the (potentially complex) refractive index of the material. However, color may also emerge from the micro-structuration of matter. Such *structural colors* require to study light as a wave propagation phenomenon.

Wave Optics. Modeling the propagation of light waves is a difficult endeavor in the general case; however, many specific models for particular structurations of matter are readily found in the optics literature, such as gratings, or thin-films for instance [Hecht 2001]. In particular, iridescence effects due to multiple scattering inside a thin film have been characterized in closed form using Airy summation [Yeh 2005]. The formula is presented in detail in Section 3, where it is shown that it is highly sensitive to phase shifts between scattered light waves. Phase is also modified on reflection depending on polarization, with formula differing between dielectrics and conductors as detailed by Born and Wolf [1999] (we reproduce and

illustrate their formula in Supplemental Material for completeness). Such polarization-dependent phase shifts may be ignored when interference effects can be safely neglected, as in geometric optics; however, in the case of iridescence (due to either diffraction or thin-films), they often have profound effects on reflectance colors, and hence cannot be avoided.

Diffraction Models. Early diffraction models in computer graphics [He et al. 1991; Stam 1999] have relied on the Beckmann-Kirchhoff theory of scattering from statistically-defined rough surfaces [Beckmann and Spizzichino 1963], which assumes height variations of the micro-surface to be small enough such that interferences do not average out as in geometric optics models. Further work has considered more complex micro-structures such as the ones arising in biological patterns [Dhillon et al. 2014], or mutual dependencies of (wave-based) reflectance between neighboring surface patches [Cuypers et al. 2012]. In this paper, we altogether avoid diffraction by assuming that micro-surface variations are much larger than visible wavelengths, as in geometric optics models.

Thin-film Models. Goniochromatic effects in thin-films are due to phase shifts between different paths in a layered structure. One of the earliest treatments of thin-film interference in computer graphics is due to Smits and Meyer [1992]. Their model is the closest to our goal as it permits to reproduce iridescent effects due to thin films of varying thickness. However, their method is limited in several ways: it is defined only for a perfectly smooth surface, it does not account for polarization, it does not consider inter-reflections inside the thin-film, nor does it work with conductors as a base material. Subsequent work has not addressed these limitations, but rather consider different types of layered structures. For instance, Icart and Arquès [1999] combined diffraction and thin-films, expanding on the Beckmann-Kirchhoff theory. Granier and Heidrich [2003] propose to model a thin-film with its interface not parallel to the base surface. Sun [2006] models the stacking of multiple thin-films to simulate natural patterns found in animals and insects, such as on the Morpho butterfly wings. In this paper, we pursue the work of Smits and Meyer, and focus on a single thin-film that may vary in thickness over an arbitrary micro-surface.

Layered Materials. Our goal may seem related to layered material models such as those of Ershov et al. [2001], Weidlich and Wilkie [2007], Jakob et al. [2014], or Ergun et al. [2016]. However, in all these models, the separation between layers is assumed to be large enough such that a geometric optics approximation may be taken, hence forbidding interference due to layers. Nevertheless, car paint models [Ergun et al. 2016; Ershov et al. 2001] do incorporate goniochromatic effects by embedding iridescent flakes inside layers. Iridescent effects may then be safely precomputed since all flakes are assumed to have the same reflectance independently of their location on the surface. This workaround is not possible when modeling variations in iridescence as we do.

Spectral Aliasing. Any spectral reflectance model will at some point require spectral integration over tristimulus sensitivity functions to yield an RGB color, which is a resource-demanding process. An alternative is to pick three representative wavelengths, one for each of the R, G, and B channels. This approach has been taken by

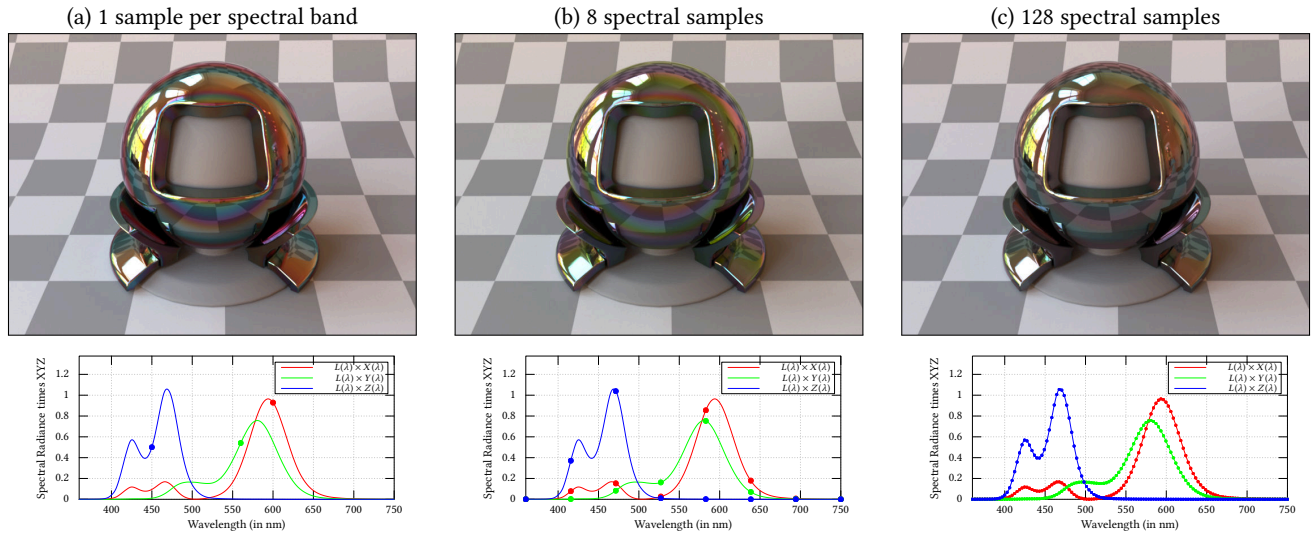(a) 1 sample per spectral band (b) 8 spectral samples (c) 128 spectral samples

Fig. 2. The addition of a thin film layer requires a dense spectral sampling to render colors properly. We show this in the top row: (a) with only one sample per spectral band (as done in RGB rendering), colors appear unnaturally over-saturated; (b) 8 spectral samples produce different color patterns; (c) a reference using 128 samples per spectral band shows the correct colors. These color artefacts are due to *spectral aliasing*, as shown in the bottom row where spectral sampling is visualized for one surface point. The colored curves correspond to reflected radiance multiplied by either the CIE X, Y, or Z sensitivity curve, in red, green and blue respectively. They show oscillations that obviously cannot be captured with a single sample (a), or even 8 samples (b), hence yielding aliasing issues in these two cases. The use of 128 samples correctly captures these oscillations but is impractical when thin-film thickness, and hence reflectance, varies.

Granier and Heidrich [2003], or for the real-time rendering of gem stones by Guy and Soler [2004]. However in the case of thin-film iridescence, this leads to severe aliasing issues in the spectral domain as shown in Figure 2. The correct number of spectral samples is not fixed either: it will depend on the thin-film properties since the number of oscillations increases with film thickness, requiring an ever-increasing number of samples. An appealing solution has been suggested by Smits and Meyer [1992]: pre-integrate spectral reflectance in a two-dimensional RGB texture indexed by film parameters and incidence angle. However, with the extensions we introduce in this paper, this solution is not viable anymore: it would require a *nine*-dimensional texture in the general case of an RGB conductor base material and should be stored in a high resolution lookup table due to the high-frequency color fringes produced by thin-film interference (see Figure 9). In contrast, we perform pre-integration of spectral reflectance using an *analytic* method that requires no material-dependent lookup table. Using pre-integrated reflectance may still produce over-saturated colors compared to the ground truth in the presence of inter-reflections. However, we consider this to be a general limitation of tri-stimulus rendering engines, for which a common solution is to reproject RGB colors inside the reproducible gamut [Meng et al. 2015].

3 SPECTRAL BRDF MODEL

We begin by presenting the assumptions we make on the micro-structuration of an iridescent material, and formally describe the model as an extension of microfacet theory. We then highlight the importance of phase shifts and state the main problem addressed in this paper: the spectral integration of reflectance due to thin-films. Even though we focus on reflectance throughout the exposition, our model also applies to transmittance as detailed in the Appendix and as demonstrated in Section 6 and in the supplemental video.

Assumptions. The structure of our model is illustrated in Figure 3. The base material consists of a micro-facet surface with a complex index of refraction $\eta_3 + i\kappa_3$. A single dielectric thin-film layer of thickness d (in nanometers) and real index η_2 is applied on top of the base layer. We assume d to be constant on a single micro-facet, which amounts to consider that spatial variations of d are of low frequency, a reasonable assumption in practice. The exterior medium has an index $\eta_1 = 1$ for air unless specified otherwise. In the following, we will describe the Bidirectional Reflectance Distribution Function (BRDF) of such a configuration.

Model. In micro-facet theory, a BRDF is defined as:

$$\rho(\boldsymbol{\omega}_o, \boldsymbol{\omega}_i; \lambda) = \frac{D(\mathbf{h})G(\boldsymbol{\omega}_o, \boldsymbol{\omega}_i)F(\mathbf{h} \cdot \boldsymbol{\omega}_i; \lambda)}{4(\boldsymbol{\omega}_o \cdot \mathbf{n})(\boldsymbol{\omega}_i \cdot \mathbf{n})}, \qquad (1)$$


Fig. 3. Our reflectance model is composed of a micro-facet surface of complex index $\eta_3 + i\kappa_3$ with a thin dielectric film of index η_2 on top of each micro-facet. The small thickness d of the film requires to treat light interactions at the surface using wave optics.

where $\boldsymbol{\omega}_o$, $\boldsymbol{\omega}_i$, and $\mathbf{n}$ are the outgoing, incoming and surface normal vectors, λ is the wavelength, D is an arbitrary micro-facet distribution evaluated at the halfway vector $\mathbf{h} = \frac{\boldsymbol{\omega}_o + \boldsymbol{\omega}_i}{\|\boldsymbol{\omega}_o + \boldsymbol{\omega}_i\|}$, G is the associated geometric (masking/shadowing) term, and F is the Fresnel reflectance term evaluated at $\mathbf{h} \cdot \boldsymbol{\omega}_i = \cos \theta_1$, with θ_1 the difference angle (see Figure 3). The Fresnel reflectance is the only term that depends on wavelength; in the general case, it also depends on polarization. However, assuming natural (i.e., randomly polarized) illumination, $F = \frac{1}{2}(F^\perp + F^\parallel)$, with $F^\perp$ (resp. $F^\parallel$) the Fresnel reflectance for light waves polarized perpendicularly (resp. parallely) to the plane of incidence containing $\boldsymbol{\omega}_i$ and $\mathbf{n}$.

Our extension consists in replacing the classic Fresnel reflectance term, F, by a more complex term, R accounting for all inter-reflections inside the thin-film layer, including constructive and destructive interference effects. This requires to consider the wave nature of light since interference effects are due to phase differences between light waves.

Airy Reflectance. For a given wavelength and polarization, reflectance is defined as the ratio $R = \frac{A_o}{A_i}$ with $A_{o,i}$ the powers of outgoing and incoming light waves, which are related to the wave amplitudes $a_{o,i}$ by $A_{o,i} \propto |a_{o,i}|^2$. We will follow the convention of using lowercase symbols for amplitudes, and uppercase symbols for powers. We thus write $R = \frac{|a_o|^2}{|a_i|^2} = |\mathbf{r}|^2$ where $\mathbf{r}$ is a *complex* reflection coefficient. The thin-film reflection coefficient $\mathbf{r}$ for an arbitrary polarization is obtained by adding the contributions of all reflected rays (see Section 4.2 in Yeh [2005]), as illustrated in Figure 4:

$$
\begin{aligned}
\mathbf{r} &= \mathbf{r}_{12} + t_{12}\mathbf{r}_{23}t_{21}e^{i\Delta\phi} + t_{12}\mathbf{r}_{23}\mathbf{r}_{21}\mathbf{r}_{23}t_{21}e^{2i\Delta\phi} + \dots \\
&= \mathbf{r}_{12} + \sum_{k=1}^{+\infty} t_{12}\mathbf{r}_{23}[\mathbf{r}_{21}\mathbf{r}_{23}]^{k-1}t_{21}e^{ik\Delta\phi} \qquad (2) \\
&= \mathbf{r}_{12} + \frac{t_{12}\mathbf{r}_{23}t_{21}e^{i\Delta\phi}}{1 - \mathbf{r}_{21}\mathbf{r}_{23}e^{i\Delta\phi}}, \qquad (3)
\end{aligned}
$$

where $\mathbf{r}_{ab} = r_{ab}e^{i\phi_{ab}}$ (resp. t_{ab}) is a complex reflection (resp. real transmission) coefficient when going from medium a to medium b, $\Delta\phi$ is the phase shift due to the optical path difference (OPD) between the primary and secondary light paths, and k is the number of inter-reflections, or order. Equation 3 is due to Sir George Biddell Airy and known as *Airy summation* in optics. Assuming a randomly polarized illumination as before, we write $R = \frac{1}{2}(|\mathbf{r}^\perp|^2 + |\mathbf{r}^\parallel|^2)$. We will denote this reflectance by the term "Airy reflectance" to distinguish it from the Fresnel reflectance term commonly used in microfacet models. It takes into account the phase shifts coming not only from OPD, but also due to reflection.

Phase Shifts. Fresnel equations show that there is no phase shift on transmission, which is why we only consider *real* transmission coefficients. However, phase shifts do occur on reflection (and depend on polarization), which is why we consider *complex* reflection coefficients. The phase shift $\Delta\phi = 2\pi\nu\mathcal{D}$ linearly depends on $\nu = 1/\lambda$, as well as on the first-order optical path difference $\mathcal{D} = 2\eta_2 d \cos\theta_2$, with $\cos\theta_2 = \sqrt{1 - \frac{\eta_1^2}{\eta_2^2}(1 - \cos^2\theta_1)}$ according to Snell's law. The


Fig. 4. The complex reflectance coefficient $\mathbf{r}$ from Equation 2 is obtained by summing the reflectance coefficients of light paths of all orders (here we show orders $k = \{0, 1, 2\}$), taking into account their interference due to phase shifts. For instance, at order $k = 1$, this corresponds to the phase shift between the path going from A to D, and the path going from A to B to C.

OPD at an arbitrary order k is simply $\mathcal{D}(k) = k\mathcal{D}$. For the sake of completeness, formula for transmission and reflection coefficients are reproduced in detail and illustrated in our supplemental material, along with a derivation of the OPD at order k. The phase shifts due to both reflection and OPD play an important role in iridescent appearance, since they impact the corresponding Airy reflectance R.

Spectral Integration. Rendering engines use a small discrete set of spectral bands, most commonly three for RGB rendering, even though some offline renderers may use around a dozen or even randomly selected spectral bands for better color fidelity. Each spectral band j has a corresponding sensitivity function s_j; for instance, we have three such functions s_R, s_G, and s_B for RGB rendering engines. We may now write the reflected radiance equation [Kajiya 1986] with an explicit mention of spectral bands using j subscripts:

$$
L_j^\uparrow(\boldsymbol{\omega}_o) = \int \frac{s_j(\lambda)}{\|s_j\|} \int_\Omega \rho(\boldsymbol{\omega}_o, \boldsymbol{\omega}_i; \lambda)L^\downarrow(\boldsymbol{\omega}_i; \lambda)(\boldsymbol{\omega}_i \cdot \mathbf{n})d\boldsymbol{\omega}_i \, d\lambda,
$$

with $L_j^\uparrow$ the reflected radiance integrated over the jth spectral band, $L^\downarrow$ the spectrally-dependent incoming radiance and Ω the upper hemisphere of directions. The sensitivity function s_j is normalized to express $L_j^\uparrow$ in units of radiance (i.e., $W\,sr^{-1}\,m^{-2}$). Each spectral band is treated independently of others and integration is performed over the support of each sensitivity function.

In most rendering engines, light sources are pre-integrated with respect to spectral bands. The BRDF is similarly pre-integrated:

$$
\rho_j(\boldsymbol{\omega}_o, \boldsymbol{\omega}_i) = \int \rho(\boldsymbol{\omega}_o, \boldsymbol{\omega}_i; \lambda)\frac{s_j(\lambda)}{\|s_j\|}d\lambda,
$$

where we use the normalized sensitivity function as before to express ρ_j in BRDF units (i.e., sr^{-1}). At the same time, it ensures energy conservation, provided ρ is itself energy-conserving. The reflected radiance equation for the jth spectral band then becomes:

$$
L_j^\uparrow(\boldsymbol{\omega}_o) \approx \int_\Omega \rho_j(\boldsymbol{\omega}_o, \boldsymbol{\omega}_i)L_j^\downarrow(\boldsymbol{\omega}_i)(\boldsymbol{\omega}_i \cdot \mathbf{n}) \, d\boldsymbol{\omega}_i,
$$

which assumes material and lighting are not spectrally correlated.

(a) Spectral integration problem   (b) Scaled sensitivity functions   (c) moduli in Fourier   (d) Integration in Fourier
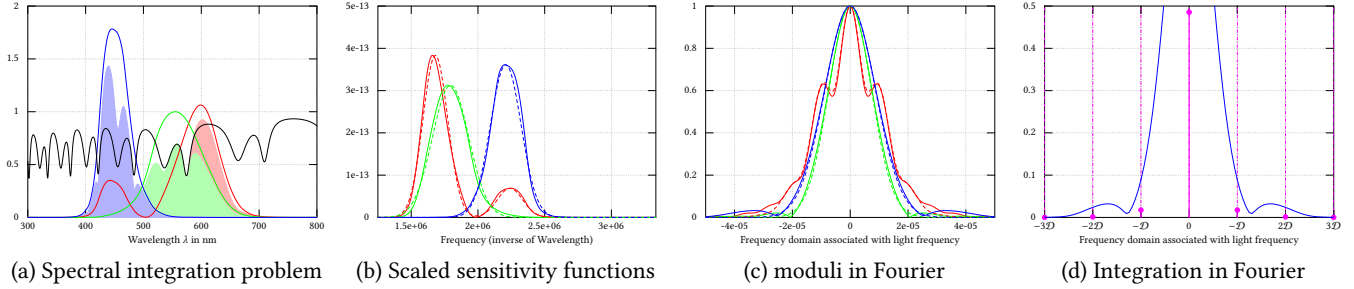
Fig. 5. The spectral integration problem is illustrated in (a) for the CIE XYZ color space: the reflectances $\{R_X, R_Y, R_Z\}$ (light colored areas) are obtained by integrating the product of each sensitivity curve (s_X in red, s_Y in green, s_Z in blue) with the spectral Airy reflectance R (in black). In our approach, the sensitivity functions are first re-expressed in terms of $\nu = 1/\lambda$ in (b), yielding $\{S_X, S_Y, S_Z\}$. They are then transformed in Fourier space: the moduli $\{|\hat{S}_X|, |\hat{S}_Y|, |\hat{S}_Z|\}$ are shown in (c). Since the Fourier transform $\hat{R}$ of the Airy reflectance term is composed of diracs, the integral in Fourier becomes analytical as shown in (d) for the evaluation of R_Z. The dashed curves in (b) and (c) show Gaussian fits, which provide reasonable approximations in practice.

The dependency of micro-facet BRDF models on wavelength only occurs in the reflectance term R; hence the BRDF pre-integration may be directly carried out to the this term:

$$R_j(\mathbf{h} \cdot \boldsymbol{\omega}_i) = \int R(\mathbf{h} \cdot \boldsymbol{\omega}_i; \lambda) \frac{s_j(\lambda)}{\|s_j\|} d\lambda. \qquad (4)$$

This spectral integral is illustrated in Figure 5(a) where we use the sensitivity functions of the CIE XYZ space.

The main issue with iridescent materials is that with variations of either the thin-film or base layer properties, R_j will have to be recomputed from scratch. This prevents its use in interactive editing scenarios where the artist freely modifies material parameters. It also forbids spatial variations as they would require a costly spectral integration at each and every surface point. The central problem is thus to provide a fast and accurate evaluation of Equation 4.

4 ANALYTIC SPECTRAL INTEGRATION

Our solution is to perform integration in the Fourier domain: we first derive an explicit formula for the spectral Airy reflectance; we then use a fast analytical spectral integration in Fourier space. The different steps of our method are illustrated in Figure 5. We validate our approach against a ground truth obtained by numerical spectral integration at the end of this section.

Spectral Airy reflectance. We will rely on Equation 2, which defines the complex reflection coefficient of a thin-film as an infinite sum. It is easily reformulated as a sum of complex numbers of the form $\mathbf{r} = \sum_{k=0}^{+\infty} c_k e^{i\phi_k}$, where for $k \geq 1$:

$$\begin{aligned} c_k &= t_{12} r_{23} [r_{21} r_{23}]^{k-1} t_{21}, \\ \phi_k &= k(\Delta\phi + \phi_{23} + \phi_{21}) - \phi_{21}, \end{aligned}$$

and we write $c_0 = -r_{21}$ and $\phi_0 = \phi_{21}$ for later convenience. Recall that r_{ab} and ϕ_{ab} denote the modulus and phase of $\mathbf{r}_{ab}$.

Expressing the spectral Airy reflectance using the c_k and ϕ_k yields:

$$\begin{aligned} |\mathbf{r}|^2 &= \left| \sum_{k=0}^{+\infty} c_k \cos\phi_k \right|^2 + \left| \sum_{k=0}^{+\infty} c_k \sin\phi_k \right|^2 \\ &= \sum_{k=0}^{+\infty} c_k^2 + 2 \sum_{k=0}^{+\infty} \sum_{l<k} c_k c_l \cos(\phi_k - \phi_l), \end{aligned}$$

where we have used a multinomial expansion for each term of the first line, then grouped terms using the trigonometric identity $\cos\phi_k \cos\phi_l + \sin\phi_k \sin\phi_l = \cos(\phi_k - \phi_l)$.

If we now write $\phi_k - \phi_l = m\Phi$ with $m = k - l$ and $\Phi = \Delta\phi + \phi_{23} + \phi_{21}$, then the spectral Airy reflectance becomes:

$$|\mathbf{r}|^2 = \sum_{k=0}^{+\infty} c_k^2 + 2 \sum_{m=1}^{+\infty} \sum_{k=0}^{+\infty} c_k c_{k+m} \cos(m\Phi),$$

More succinctly, using $C_0 = \sum_{k=0}^{+\infty} c_k^2$ and $C_m = \sum_{k=0}^{+\infty} c_k c_{k+m}$:

$$\boxed{R = C_0 + 2 \sum_{m=1}^{+\infty} C_m \cos(m\Phi).} \qquad (5)$$

The C_m terms are illustrated in Figure 6 for $m \in \{0, 1\}$, where m represents the offset in orders between pairs of light paths. They are derived in closed form in the Appendix, yielding:

$$C_0 = R_{12} + R_\star; \quad C_m = \left(\sqrt{R_{23} R_{21}}\right)^m \left(R_\star - \sqrt{T_{12} T_{21}}\right), \qquad (6)$$

with R_{ab} and T_{ab} denoting Fresnel reflectances and transmittances between media a and b, and where $R_\star = \frac{T_{12} T_{21} R_{23}}{1 - R_{23} R_{21}}$ encapsulates all inter-reflections inside the thin-film layer.

We also rewrite Equation 5 in closed-form in the Appendix:

$$R = R_{12} + R_\star + \frac{2 \left(R_\updownarrow \cos\Phi - R_\updownarrow^2\right)}{1 - 2 R_\updownarrow \cos\Phi + R_\updownarrow^2} \left(R_\star - \sqrt{T_{12} T_{21}}\right),$$

where we have used $R_\updownarrow = \sqrt{R_{23} R_{21}}$. However, its interpretation in Fourier is not straightforward; hence we will only use it for computing ground-truth reflectances in the following.

$$C_0 = c_0 c_0 + c_1 c_1 + \ldots \qquad\qquad C_1 = c_0 c_1 + c_1 c_2 + \ldots$$

Fig. 6. We illustrate the C_m terms of Equation 5. Left: C_0 is obtained by summing interferences for all pairs of light paths of the same order. Right: C_1 is obtained by summing interferences for all pairs of light paths with an offset of $m = 1$ in orders.

Spectral Integration in Fourier. In order to simplify the spectral integration of Equation 4, we are now going to assume that all Fresnel amplitude and phase coefficients are constant per spectral band (we will evaluate the pertinence of this approximation later). A direct consequence is that the spectral dependence in R now only occurs in the phase shift $\Delta\phi = 2\pi\nu\mathcal{D}$, which is linear in $\nu = 1/\lambda$, with $\mathcal{D}$ the optical path difference. We must now re-express Equation 4 in terms of ν using a change of variables, yielding:

$$R_j = \int R(\nu) \frac{s_j\left(\frac{1}{\nu}\right)}{\|s_j\|\nu^2}\, d\nu. \tag{7}$$

Writing $S_j(\nu) = \dfrac{s_j\left(\frac{1}{\nu}\right)}{\|s_j\|\nu^2}$ and using Parseval's theorem yields:

$$R_j = \int \hat{R}(\mu)\, \hat{S}_j(\mu)^{\star}\, d\mu, \tag{8}$$

where $\hat{R}(\mu)$ and $\hat{S}_j(\mu)$ are the Fourier transforms of $R(\nu)$ and $S_j(\nu)$ respectively, and $\star$ denotes the complex conjugate. The scaled sensitivity functions S_j and the moduli of their Fourier transform $|\hat{S}_j|$ are shown in Figures 5(b) and 5(c) respectively, for each of the CIE X, Y, and Z bands.

The (unitary, ordinary frequency) Fourier transform of R follows from Equation 5 by using Euler's formula to separate the spectrally-independent phase shift $\phi_2 = \phi_{21} + \phi_{23}$ from $\Delta\phi$, yielding:

$$\hat{R}(\mu) = C_0 \delta(\mu) + \sum_{m=1}^{+\infty} C_m \left[e^{im\phi_2}\delta(\mu - m\mathcal{D}) + e^{-im\phi_2}\delta(\mu + m\mathcal{D}) \right], \tag{9}$$

where δ is the dirac function. As shown in Figure 5(d), $|\hat{R}|$ is a distribution of diracs each separated by $\mathcal{D}$, with the amplitudes of the DC term and the mth harmonic given by C_0 and C_m respectively. Therefore Equation 8 may now be evaluated analytically.

Since S_j is a real function, the real part of $\hat{S}_j$ is symmetric, and its imaginary part is anti-symmetric; hence $\hat{S}_j(-\mu) = \hat{S}_j(\mu)^{\star}$. Plugging this formula and Equation 9 inside Equation 8 yields:

$$R_j = C_0 + \sum_{m=1}^{+\infty} C_m \left[e^{im\phi_2}\hat{S}_j(-m\mathcal{D}) + e^{-im\phi_2}\hat{S}_j(m\mathcal{D}) \right],$$

where we use the fact that $\hat{S}_j(0) = 1$ by construction.

If we explicitly write $\hat{S}_j(\pm\mu) = \mathfrak{R}_j(\mu) \pm i\mathfrak{I}_j(\mu)$ with $\mathfrak{R}_j$ and $\mathfrak{I}_j$ the real and imaginary parts of $\hat{S}_j$ respectively, then after a few straightforward simplifications, we obtain:

$$R_j = C_0 + 2 \sum_{m=1}^{+\infty} C_m \begin{bmatrix} \cos(m\phi_2) \\ \sin(m\phi_2) \end{bmatrix}^T \begin{bmatrix} \mathfrak{R}_j(m\mathcal{D}) \\ \mathfrak{I}_j(m\mathcal{D}) \end{bmatrix}. \tag{10}$$

It should be noted that since C_0 and C_m are defined in terms of Fresnel reflectances and transmittances (see Equation 6), they depend on indices of refraction that we assumed constant per spectral band. For a given band, a natural choice is to take the refractive index for which the sensitivity function is maximum (i.e., its mode). Strictly speaking, the C_0 and C_m terms should be subscripted by the spectral band index j; but we prefer to avoid this notation for clarity.

Equation 10 is our main result. If we choose a maximum value for m, it provides a closed-form approximation to Airy's reflectance for a given spectral band. Note that with a dielectric base, $\phi_2 = 0$ or π meaning that only the real part of $\hat{S}_j$ needs to be evaluated. In this case, we should also consider transmission through both layers. The simplest approach is then to use energy conservation and define an Airy transmittance term by $T_j = 1 - R_j$, which may then be plugged in any microfacet-based BTDF model. For completeness, we also re-derive Airy transmittance from the corresponding Airy summation formula in the Appendix.

Validation. We validated our pre-integration strategy using a pair of datasets: one with constant indices for both layers, the other with spectrally-varying indices. We used an integration step of $1nm$ to generate the ground truth and performed comparisons in CIE XYZ color space. For each band, the reference indices for the C_m terms in Equation 10 are those corresponding to the peak of each sensitivity curve (i.e., at $\lambda_X = 600nm$, $\lambda_Y = 560nm$ and $\lambda_Z = 450nm$).

The constant-index dataset permits to specifically compare various approximations of the Airy reflectance term. We use a dielectric thin-film of thickness comparable to visible wavelengths over a smooth dielectric base (i.e., D is a dirac distribution in Equation 1). Results are shown as curves parametrized by θ_1 plotted in the CIE xy chromaticity space, since we are mostly concerned by iridescent color effects. Note that when $\theta_1 \rightarrow \pi/2$, reflectances for all bands tend to 1 and the curves in chromaticity space tend toward the equi-luminant point $E = (\frac{1}{3}, \frac{1}{3})$. In Figure 7(a), we show that the naive approach that uses one wavelength sample per sensitivity curve produces results far from the ground truth, even creating over-saturated colors that run outside of the CIE RGB gamut. If instead we use a large number of wavelength samples as with the ground truth but then resort to the simplified model of Smits and Meyer [1992] as shown in Figure 7(b), we obtain less saturated colors, but the curve still exhibits a significant disparity compared to the reference curve. We attribute these differences to the limitations of their model: no multiple scattering, and no account of base conductors or polarization. In Figure 7(c) we show our result using the analytical spectral integration of Equation 10. It shows that truncating m at 3 or even 1 yields curves nearly indistinguishable from the reference curve. We make the same comparison for each of the X, Y and Z spectral

(a) Naive approach   (b) Smits and Meyer   (c) Our approach   (d) Ours for X, Y, Z bands
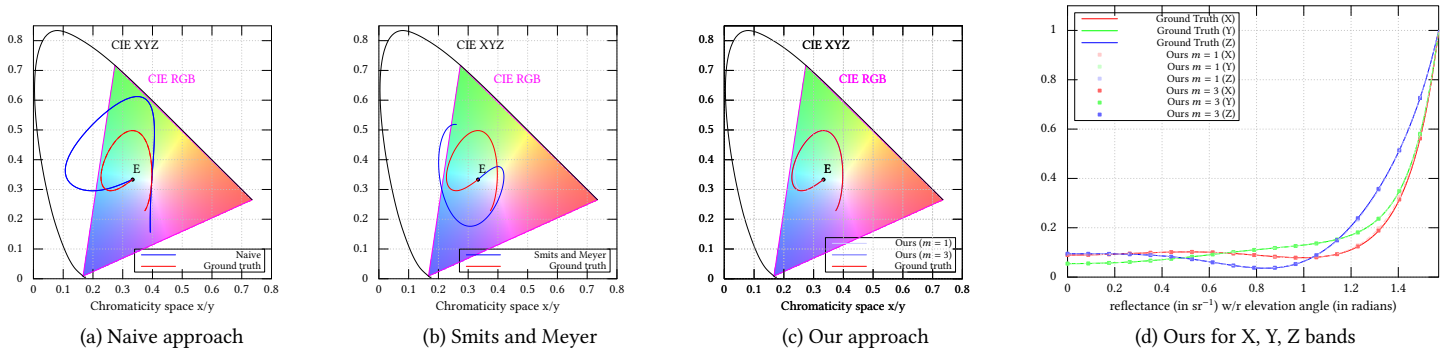
Fig. 7. We compare different approximations to a ground truth reflectance for a dielectric film ($\eta_2 = 1.5$) of thickness $d = 525nm$ over a dieletric base ($\eta_3 = 1.09$). In (a-c), reflectance curves parametrized by θ_1 are plotted in CIE xy chromaticity space, with the ground truth in red and approximate solutions in blue. The naive approach in (a) uses the exact Airy reflectance, but only one wavelength sample per spectral band, which yields poor results. The approach of Smits and Meyer [Smits and Meyer 1992] in (b) makes several simplifications to the reflectance model (e.g., $\kappa_3 = 0$); even with a dense spectral sampling identical to the ground truth, it still shows significant disparities. Our approach is shown in (c) for different truncations of the offset m. Computationally, it is only slightly more complex than the naive approach; yet it produces results nearly indistinguishable from the ground truth, as is also shown for each of the X, Y, and Z spectral bands in (d).

bands separately in Figure 7(d).

Detailed comparisons using the varying-index dataset are provided in our supplemental material. For the sake of brevity, we only provide visual comparisons on rendered spheres in Figure 8. The figure shows three examples of a dielectric thin-film (buthanol) applied over a dielectric or conducting base, with both layers having refractive indices that vary with wavelength. This permits to evaluate the impact of assuming that Fresnel phase and amplitude coefficients are constant per spectral band. We observe that our approximation becomes slightly less accurate for a conducting base layer such as copper, whose index of refraction varies non-monotonically with wavelength. Of course, these differences would likely be reduced if more spectral bands were available.

5 PRACTICAL CONSIDERATIONS

Having described and validated our BRDF model, we now discuss practical issues one must consider when incorporating any material model in modern rendering engines: how to evaluate and prefilter

(a) Glass base   (b) Mercury base   (c) Copper base

Fig. 8. Comparisons between the ground-truth Airy reflectance and our approximation, on spheres rendered in the Ufizzi environment lighting. A thin-film of buthanol of thickness $d = 666nm$ is laid on top of a base layer made of either (a) glass, (b) mercurial, or (c) copper. Each image is split in two, with our result on the left and the ground truth on the right.

it efficiently, especially for multi-scale rendering; which parameters should be brought to artists to control iridescent appearance.

BRDF Evaluation. The most direct method for evaluating the pre-integrated Airy reflectance (or transmittance) term is to tabulate $\hat{S}_j$, the Fourier transform of the scaled sensitivity curves. When working with a tristimulus rendering engine, the real and imaginary parts of $\hat{S}_j$ may be stored in each row of a $N \times 2$ three-channel texture, where N is the resolution in the Fourier dimension. Evaluating Equation 10 would then normally require two texture fetches to get the real and imaginary parts; however, a single fetch at texture location $(m\mathcal{D}, \gamma)$ using bilinear interpolation is enough since:

$$\begin{bmatrix} \cos(m\phi_2) \\ \sin(m\phi_2) \end{bmatrix}^T \begin{bmatrix} \mathfrak{R}_j(m\mathcal{D}) \\ \mathfrak{I}_j(m\mathcal{D}) \end{bmatrix} = \beta \, \mathrm{lerp}\left(\mathfrak{R}_j(m\mathcal{D}), \mathfrak{I}_j(m\mathcal{D}), \gamma \right),$$

with $\beta = \cos(m\phi_2) + \sin(m\phi_2)$ and $\gamma = \sin(m\phi_2)/\beta$.

Another method consists in approximating scaled sensitivity functions (and their Fourier transforms) using Gaussians. This is shown in Figure 5(b-c), where we have used two Gaussians for the X band, and one Gaussian for each of the Y and Z bands. This approximation results in an appearance close to the ground-truth BRDF when dealing with dielectric base layers. However, it tends to produce slightly over-saturated colors for conducting base layers, or at grazing angles. We attribute these differences to the subtle oscillations in $\hat{S}_j$ (see Figure 5(c)) that are not captured by the Gaussian fits. The practical advantage of this approximation is that it requires only a few input parameters for the Gaussians, as demonstrated by the GLSL implementation provided in our supplemental material.

For importance sampling, we follow the approach of existing techniques that uniquely rely on the microfacet distribution D of Equation 1 and ignore the reflectance term R.

Fig. 9. We show 2D slices of the Airy reflectance term for a base material of index $\eta_3 = 1.1$, with either $\kappa_3 = 0$ (top row) or $\kappa_3 = 1.5$ (bottom row). In each slice, the horizontal axis maps to θ_1, while the vertical axis maps to $\eta_2 \in [1..2]$. Airy reflectance is highly sensitive to variations of $\mathcal{D}_{\text{inc}}$ the OPD at normal incidence.

Multi-scale Rendering. Our model is linearly dependent on the thickness d of the thin-film, which makes it adapted to the prefiltering of spatial thickness variations. Formally this requires to integrate our BRDF model against a distribution P of thickness values. We make two simplifying assumptions: P is modeled as a Gaussian distribution and is not correlated with the microfacet distribution D (i.e., for every normal the associated distribution of thicknesses is the same: $P(d|\mathbf{h}) = P(d) \; \forall \mathbf{h}$). As a result of the latter, integration only takes place in the reflectance term, yielding:

$$\bar{R}_j(\mathbf{h} \cdot \boldsymbol{\omega}_i) = \int R_j(\mathbf{h} \cdot \boldsymbol{\omega}_i) P(d) \, \mathrm{d}d.$$

Using Equation 10 and moving outside of the integral the terms that are constant with respect to d, we obtain:

$$\bar{R}_j = C_0 + 2 \sum_{m=1}^{+\infty} C_m \begin{bmatrix} \cos(m\phi_2) \\ \sin(m\phi_2) \end{bmatrix}^T \int \begin{bmatrix} \mathfrak{R}_j(m\mathcal{D}) \\ \mathfrak{I}_j(m\mathcal{D}) \end{bmatrix} P(d) \, \mathrm{d}d.$$

If we now explicitly write $m\mathcal{D}$ in term of d and perform a change of variable $t \leftarrow m\mathcal{D}$ in the integral, we get:

$$\int \begin{bmatrix} \mathfrak{R}_j(m\mathcal{D}) \\ \mathfrak{I}_j(m\mathcal{D}) \end{bmatrix} P(d) \, \mathrm{d}d = \int \begin{bmatrix} \mathfrak{R}_j(t) \\ \mathfrak{I}_j(t) \end{bmatrix} \frac{P\left(\frac{t}{\tau}\right)}{\tau} \, \mathrm{d}t,$$

with $\tau = m2\eta_2 \cos \theta_2$. Specifically, $P'(t) = P\left(\frac{t}{\tau}\right)/\tau$ is a Gaussian distribution that is shifted and scaled with respect to P, since $\mathrm{E}[P'] = \tau \, \mathrm{E}[P]$ and $\mathrm{Var}[P'] = \tau^2 \, \mathrm{Var}[P]$.

In practice, we first prefilter the real and imaginary parts of $\hat{S}_j$ with Gaussian kernels of increasing variance. The result is stored in a dedicated mip-map, with the scale dimension being indexed by the variance parameter of the Gaussian. Then, at run time, depending on τ (which varies with m and local material parameters), the mip-map is fetched at a different location $\mathrm{E}[P']$ and scale $\mathrm{Var}[P']$.

Fig. 10. The Airy reflectance term is directly visualized on spheres with three values of η_2 (one per column); other material parameters are held fixed to $\mathcal{D}_{\text{inc}} = 1600$nm, $\eta_3 = 1.1$, and $\kappa_3 = 1.5$. More color fringes are obtained with smaller values of η_2.

Fig. 11. The Airy reflectance term at normal incidence is visualized for $\mathcal{D}_{\text{inc}} \in [0..6400]$nm with three values of η_2 (one per row); other material parameters are held fixed to $\eta_3 = 1.1$ and $\kappa_3 = 1.5$. Color saturation and intensity increase with larger values of η_2.

Parametric Control. In our model, the thin-film layer is controlled by a pair of physical parameters, d and η_2. Since d only appears in the optical path difference (OPD), we rather provide a direct control over the OPD at normal incidence, denoted by $\mathcal{D}_{\text{inc}} = 2\eta_2 d$; we then use $\mathcal{D} = \mathcal{D}_{\text{inc}} \cos \theta_2$. The parameter space of the Airy reflectance term is visualized using $(\cos \theta_1, \eta_2)$-slices at various values of $\mathcal{D}_{\text{inc}}$ in Figure 9, for both dielectric and conducting base layers. First observe that even with an achromatic base layer, it exhibits high-frequency oscillations that would require a very high-resolution lookup table if it were to be precomputed. This would be even more problematic with a colored base layer, since the look up table would also increase in dimensionality, up to nine dimensions for a conductor in RGB (two for the thin-film, six for the base, and one for $\cos \theta_1$). Second, note that when $\mathcal{D}_{\text{inc}}$ is made large enough, iridescent effects begin to vanish. This is because for large values of the OPD, the diracs of $\hat{R}$ (see Figure 5(d)) will be distant enough so that only the DC term will significantly contribute to the spectral integral in Equation 9. We thus define a maximal OPD $\mathcal{D}_{\text{max}}$ such that $\forall \mu \geq \mathcal{D}_{\text{max}}, |\hat{S}_j(\mu)| \leq \epsilon$: beyond $\mathcal{D}_{\text{max}}$, iridescence is considered negligible. We use $\epsilon = 0.05$ in our implementation, which yields $\mathcal{D}_{\text{max}} \approx 25$ microns for the CIE XYZ color space. Even though $\mathcal{D}_{\text{max}}$ is a valid criterion only at normal incidence, we use it at all incidence angles. This is justified by the observation that C_0 dominates the C_m terms at all incidence angles, the latter eventually vanishing at $\theta_1 = \frac{\pi}{2}$ (see Figure 22).

Such a higher bound on the OPD permits to control iridescence as a whole, since the latter only occurs when $0 < \mathcal{D}_{\text{inc}} < \mathcal{D}_{\text{max}}$. When $\mathcal{D}_{\text{inc}} = \mathcal{D}_{\text{max}}$, $R \approx C_0$, which is consistent with the reflectance of

thick layers (e.g., see [Jakob et al. 2014]). When $\mathcal{D}_{\text{inc}} = 0$, R should become equal to R_{13} since the thin-film then effectively vanishes. Unfortunately, this configuration is not physically-valid since it yields a pair of superimposed interfaces with different pairs of indices on each side. A physically-realistic treatment would require to model the case of extremely thin layers (below a few nm) with quantum optics, which is clearly out of the scope of this paper. Our alternative solution is to force $\eta_2 \to \eta_1$ when $\mathcal{D}_{\text{inc}} \to 0$, hence ensuring that we end up with a single effective interface.

The thin-film index η_2 provides a more subtle control over appearance, with two main effects. First, assuming a constant $\mathcal{D}_{\text{inc}}$, η_2 controls the number of color fringes swept by when θ_1 varies from 0 to $\pi/2$. This is shown in Figure 10 using spheres rendered in a white furnace environment. Second, assuming a constant incidence angle θ_1, variations in η_2 modify the intensity and saturation of color fringes. This is shown in Figure 11 where a fronto-parallel surface is rendered with $\mathcal{D}_{\text{inc}}$ varying linearly from left to right.

Finally, the complex index $\eta_3 + i\kappa_3$ of the base layer may be directly given for each spectral band, or it may be obtained from more intuitive color input [Gulbrandsen 2014].

6 RESULTS

We have implemented our approach in GLSL shaders for Disney's BRDFExplorer [Disney 2011] and Gratin [Vergne and Barla 2015], using the GGX distribution [Walter et al. 2007] for the former, and the Ward distribution [Ward 1992] for the latter. For global illumination results, we have created a plugin for Mitsuba [Jakob 2010], using the GGX distribution again. Note that our method is completely orthogonal to the choice of microfacet distribution. We obtained slightly less saturated colors when using Mitsuba due to the choice of color space in the API (sRGB instead of CIE RGB). As demonstrated in our supplemental video, our GLSL implementation does not suffer from this limitation since we fully control the spectral conversion. The BRDFExplorer shader and the Mitsuba plugin are provided in our supplemental code.

Direct Lighting. We demonstrate our GLSL shader in BRDFExplorer through an editing session in our supplemental video. It shows that our model reproduces the ground truth very accurately, while permitting interactive manipulation. We highlight the artifacts produced by the naive approach with one wavelength sample per color channel: it yields wrong colors and iridescent effects do not vanish when $\mathcal{D}_{\text{inc}} \to \mathcal{D}_{\text{max}}$ as they should.

Our approach is particularly useful when the thickness of the thin-film is varied spatially, which is shown in Figure 12 and in our supplemental video for two types of materials, on a SPEED SHAPE model. We add a red diffuse base to the dielectric material on the left column using a simple Lambertian model ρ_d. We also add a clear-coat layer on top of both materials using the approach of Weidlich and Wilkie [2007] (we set $\eta_1 = 1.1$) to reproduce the appearance of car paint. Thin-film variations are controlled with a texture that remaps $\mathcal{D}_{\text{inc}}$ between 0 and a maximum OPD. When the maximum OPD is increased, the iridescent effects sweep through a series of


Fig. 12. The SPEEDSHAPE model is rendered with a dielectric base in the left column ($\alpha = 0.1$, $\eta_2 = 1.25$, $\eta_3 = 1.72$, $\rho_d = \{0.191, 0.015, 0.015\}$) and a rough conductor base in the right column ($\alpha = 0.18$, $\eta_2 = 1.25$, $\eta_3 = 1.1$, $\kappa_3 = 1.63$). The thin-film of spatially-varying thickness is revealed by remapping $\mathcal{D}_{\text{inc}}$ from 0 to 100nm (first row), to 200nm (second row), and to 400nm (third row), controlled by the texture shown on top.

color fringes that provide a realistic look to the SPEED SHAPE.

The texture map used to control the thin-film may be chosen for aesthetic purposes, as shown in Figure 13. On the left side, we re-use an ambient occlusion map to introduce subtle color variations in cavities of the ROBOT model. On the right side, we replace parts of the texture with regular patterns to convey a futuristic look.

Spatial variations of iridescence also bring realism to the rendering of transparent objects. This is shown in Figure 14, where we have created a thick transparent slab using a displacement map on

Fig. 13. The ROBOT model is made of a metallic base ($\alpha = 0.07$, $\eta_3 = 1.87$, $\kappa_3 = 1.182$) to which we add a thin-film of index $\eta_2 = 1.42$ and varying thickness ($\mathcal{D}_{\text{inc}} \in [0..1125]$nm). On the left, we use an ambient occlusion texture to guide thickness variations. On the right we add regular patterns to give a futuristic look. The image on top is rendered without iridescence.
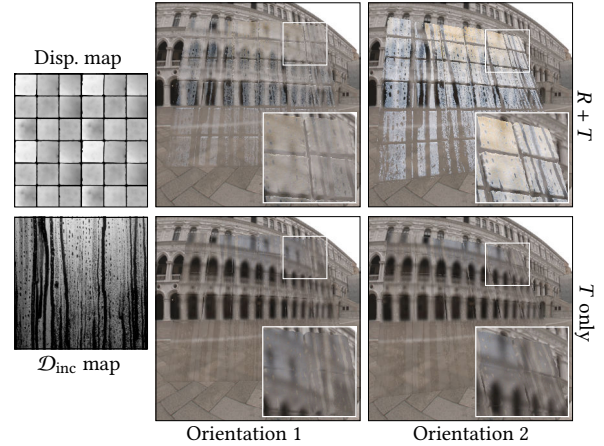

Fig. 14. Our model works with transparent objects, here a transparent slab obtained by applying a displacement map to a plane, with a rough dielectric base ($\alpha = 0.2$, $\eta_3 = 1.45$). We control variations of the thin-film ($\eta_2 = 2.0$) using a texture that remaps $\mathcal{D}_{\text{inc}}$ to the $[0..415]$nm range, mimicking a condensation effect that is more or less revealed depending on the slab orientation. We also show the transmitted radiance in isolation in the bottom row.

a plane, then added thin-film variations using a texture controlling $\mathcal{D}_{\text{inc}}$ as before. We not only visualize the combination of reflection and transmission, but also transmission in isolation, which reveals subtle color variations on close inspection (see insets).

Aliasing artifacts may occur when using highly detailed thin-film variations. As shown in Figure 15, our prefiltering solution reduces these artifacts efficiently, with visual results close to the reference.

Our model runs in real time when the base layer is perfectly smooth: the distribution D in Equation 1 then becomes a dirac, hence R becomes a function of $\mathbf{n} \cdot \omega_i$. A common approximation for real-time rendering of rough materials is to prefilter the environment lighting [Kautz et al. 2000], which is then evaluated once in the specular direction. We combine this technique with an evaluation of R in the specular direction as well in Figure 16 (our supplemental video shows this scene captured in real time). It gives satisfying visual results when the material is smooth or even moderately rough, but produces over-saturated colors with rougher materials.

Global Illumination. We illustrate the full range of appearances achievable by our model in supplemental material for conductor and dielectric materials.

We focus here on material parameters adapted to specific cases, such as in Figure 1 where the thin-film is applied to a leather CHAIR model. This example is inspired by Akin [2014], where thin-film iridescence is produced by a brute-force approach. Iridescent colors indeed push the realism of the material: when rotating the object, colors exhibit subtle changes in hue. Similarly subtle, yet important color effects may be obtained by applying our model to the metallic


Fig. 15. A highly-detailed texture is mapped to a sphere to produce spatial variations of the film thickness, here applied on a dielectric base layer ($\alpha = 0.01$, $\eta_2 = 1.68$, $\eta_3 = 2.0$, $\mathcal{D}_{\text{inc}} \in [0..1470]$ nm). When the sphere is moved away from the camera (middle column), aliasing due to the Airy reflectance term becomes visible, unless we use our dedicated anti-aliasing (AA) technique. As shown in the zoomed insets, our method gives results close to the reference zoomed image.

BOAR model, as seen in Figure 17.

Special effect pigments for car paints are usually less subtle. We reproduce such an example in Figure 18, adding thin-film variations on the doors to give a custom touch to the BEETLE model. As with the SPEED SHAPE, we add a clear-coat layer to reproduce a car paint appearance.

Our approach may be applied to transparent objects, as shown in the classic example of a SOAP BUBBLE in Figure 19. In this case, both reflection and transmission at each interface (front and back) are responsible for the rich color patterns observed in the result.
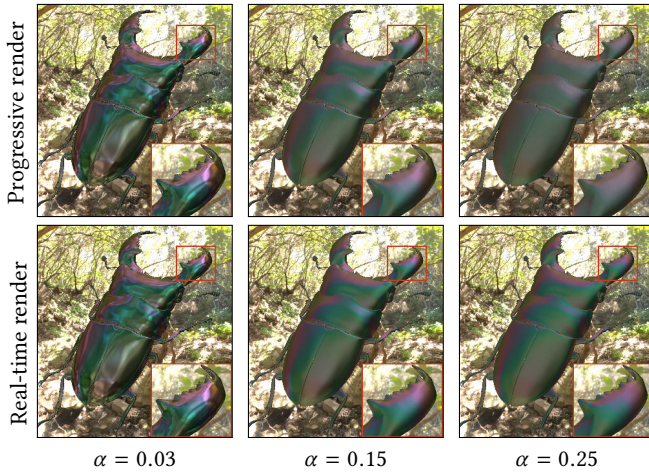
Fig. 16. The STAG BEETLE model is rendered with a dielectric base ($\eta_2 = 1.21$, $\eta_3 = 2.0$, $\mathcal{D}_{inc} = 1740$nm) using progressive rendering in the top row, and a real-time approximation in the bottom row (see text). When the roughness α of the base layer increases, the real-time approximation is less and less accurate; in particular, color fringes become over-saturated.

Comparisons. Compared to the method of Smits and Meyer [1992], our approach approximates the ground truth much more accurately, even when evaluated in RGB. This is not only demonstrated in Figure 7, but also shown in Figure 20, where we have used our Mitsuba plugin on the MAT PREVIEW scene. There are still slight differences in color saturation between our method and the ground truth, which are due to the choice of RGB color space in Mitsuba as previously explained.

We compare our analytical micro-facet model to the more complex micro-flake model of Ergun et al. [Ergun et al. 2016] in supplemental material. From a physical point of view, this amounts to consider that micro-facets act as uncorrelated iridescent flakes, a condition that is satisfied when Smith's geometric term is used for G in Equation 1. A couple example comparisons are given in Figure 21, where the appearance obtained with their model is imitated by adjusting the parameters of our model by hand. We achieve similar color fringes overall, even though their results tend to yield slightly more saturated colors, which we attribute to the effect of multi-layered thin-film interferences. A clear advantage of our approach is its efficiency: the material parameters can be modified interactively with our GLSL implementation. In comparison, their approach requires a costly preprocess that takes several seconds and requires around 10mb of storage per material configuration. However, their model is not designed for the same usage: it permits to accurately predict the appearance of specific car paints, while ours only imitate them.

Interactive implementation. As seen in the supplemental video, our model is nearly as efficient as the naive model (i.e., using one wavelength sample per color channel) in our BRDFExplorer implementation. More specifically, to render a sphere at 800spp in 1024×1024 resolution *without* the thin-film layer takes 1.1sec on a NVidia GeForce GTX555. The naive model takes approximately 1.5


Fig. 17. We reproduce the subtle iridescent appearance of oxide layers on top metals with the BOAR model. Iridescence adds a touch of green and purple (right inset) to a base 'white' material (left inset). The base layer index is here given by $\eta_3 = 1.5$ and $\kappa_3 = 3$, and the thin film is $d = 600$nm thick and has an index of $\eta_2 = 1.33$.


Fig. 18. With the BEETLE model we tried to reproduce the look of special effet car paints (see our supplemental video for an example). We show that even for a configuration where the variation of colors is moderate and dense sampling is probably not required, the naive RGB model still produces inacurrate colors. Here we used a conducting base layer ($\eta_3 = 1.2$, $\kappa_3 = 0.5$) with a thin-film of $d = 505$nm and index $\eta_2 = 1.39$. We also added a clear coat of $\eta_1 = 1.2$ to further match the car paint appearance. The door-side sticker is created by varying thickness using a texture.


Fig. 19. With the SOAP BUBBLE model, we show how iridescence behaves in transmission & reflections with multiple scattering. In this example, the multiple transmissions and reflections (TRT, $TRRT$, etc) add to the purely reflected signal to produce a convincing look. Here we use a film of $d = 400$nm and $\eta_2 = 1.7$, and a base layer of air (i.e., $\eta_1 = \eta_3 = 1.0$).

Fig. 20. We use the MAT PREVIEW scene to compare our approach both with the ground truth, and with the model of Smits and Meyer, which here exhibits incorrect green colors at grazing angles. We use a dielectric film ($\eta_2 = 1.33$, $d = 550$nm) on top of a conducting base ($\eta_3 = 1.9$, $\kappa_3 = 1.5$).


Fig. 21. We have reproduced by hand the results of Ergun et al. (their Figure 7) by adjusting the parameters of our model. Here we show two out of the eight materials presented in supplemental material: our approach permits to reproduce very similar color fringes.

seconds, while ours takes 1.6, 1.9 and 2.2 seconds when we truncate m to 1, 2, and 3 respectively in Equation 10 (we truncate m to 3 in the video). The overhead of our approach compared to the naive model (i.e., the simplest implementation of thin-film iridescence) is thus very reasonable. In comparison, using the ground truth is prohibitively slow: it takes 70 minutes using 94 wavelength samples per color band (each separated by 5nm), which is likely due to the poor performance of long loops in GLSL.

Offline implementation. We report rendering times for the different offline scenes in the following table:

	No irid RGB	Naive RGB	Spectral	**Ours (RGB)**
MAT PREVIEW	2.0m	2.4m	11.8m	2.6m
SOAP BUBBLE	33s	55s	8.1m	57s
CHAIR	43s	50s	3.2m	53s
BEETLE	3.2m	4.0m	16.6m	4.3m
BOAR	3.1m	4.0m	23.4m	4.5m

In the first column, we report rendering times without interferences (by replacing the Airy reflectance term R by the Fresnel reflectance term F). We then provide rendering times when using the naive RGB model with Airy reflectance, a spectral rendering (with 128 samples per spectral band), and our analytical model (cutting off m at 2, in blue). Since Airy reflectance evaluates multiple Fresnel reflectances, it requires more time: between 20% and 30% for the naive RGB model. Using our method adds less than 10% to

these timings, but achieves results very close to the ground truth spectral model, which is much more expensive in comparison.

7 DISCUSSION

We have introduced an extension to microfacet theory that models thin-film interference effects based on Airy summation, in practice replacing the Fresnel term by a new Airy term. Our main contribution resides in an analytical spectral integration taking advantage of the simple forms of Airy reflectance and transmittance in Fourier (a series of Diracs). As a result, the extended micro-facet model may be edited with immediate feedback, used with surface-varying parameters with negligible overhead, and rendered at multiple scales without producing aliasing artifacts.

Limitations. The assumptions we made in our approach forbid some specific micro-structures. For instance, we assume the thin-film to be flat and parallel to the underlying microfacet. If this is not the case, reflectance colors will likely appear desaturated and might change in tint due to the average of different optical path differences. While the desaturation could be handled with our multi-scale approach, we cannot account for the potential change in tint. Parallelism with the microfacet will eventually be untenable for films thicker than a few microns, in which case iridescence starts to become irrelevant and other phenomena such as absorption and scattering begin to matter. An interesting challenge would be to model the transition from thick to thin films, as is the case of a liquid drying on a rough surface for instance.

On the practical side, our model permits interactive editing in progressive rendering for arbitrary material parameters. For real-time rendering, the simple approximation presented in Section 6 remains limited to material of small roughness. A more promising solution would be to pre-integrate the Airy term for each possible view direction, in essence yielding a color-varying directional albedo. Such an integration will of course depend on the choice of microfacet distribution.

The analytical spectral integration of Equation 10 assumes that spectral bands are fixed throughout rendering. Some high-end rendering engines rather perform spectral sampling from scratch at each light-surface interaction to better account for complex lighting effects [Wilkie et al. 2014]. The pre-integrated Airy term could still be used as a control variate [Hammersley and Morton 1956] in such cases, and hence serve as a guide to spectral sampling. This would actually be beneficial in the case of iridescence to get rid of spectral aliasing; in particular, using 4 randomly shifted spectral samples as in the original hero wavelength sampling technique is clearly not sufficient to avoid spectral aliasing due to thin-film iridescence.

Future Work. Our model reproduces the iridescent effects occurring with thin-films, which are due to the interference of parallel waves. Interferences of non-parallel light waves are rather due to diffraction, which not only modifies the colors but also the angular spread of scattering lobes. A challenge for future work is to model the effect of diffraction in tristimulus engines while retaining physical plausibility, analytic formulation and parametric control.

Finally, even though the design of material acquisition devices has greatly improved in recent years, measuring goniochromatic materials remains a challenge. Not only do they require a dense sampling in angular and spatial dimensions, but also in the spectral dimension owing to the complex iridescent colors they produce. With a sufficiently rich database of measured goniochromatic materials, it would become possible to investigate the fitting of data using our microfacet-based extension. Another interesting direction of future work would be to study whether our parametric model could be used to guide the material acquisition process directly in RGB.

ACKNOWLEDGEMENTS

REFERENCES

Attila Akin. 2014. Pushing the limits of realism of materials. http://blog.maxwellrender.com/tips/pushing-the-limits-of-realism-of-materials. (2014).

P. Beckmann and A. Spizzichino. 1963. *The scattering of electromagnetic waves from rough surfaces.* Pergamon Press; [distributed in the Western Hemisphere by Macmillan, New York]. https://books.google.fr/books?id=QBEIAQAAIAAJ

Max Born and Emil Wolf. 1999. *Principles of Optics* (7th edition ed.). Cambridge University Press.

R. L. Cook and K. E. Torrance. 1982. A Reflectance Model for Computer Graphics. *ACM Trans. Graph.* 1, 1 (Jan. 1982), 7–24. DOI:https://doi.org/10.1145/357290.357293

Tom Cuypers, Tom Haber, Philippe Bekaert, Se Baek Oh, and Ramesh Raskar. 2012. Reflectance Model for Diffraction. *ACM Trans. Graph.* 31, 5, Article 122 (Sept. 2012), 11 pages. DOI:https://doi.org/10.1145/2231816.2231820

D. S. Dhillon, J. Teyssier, M. Single, I. Gaponenko, M. C. Milinkovitch, and M. Zwicker. 2014. Interactive Diffraction from Biological Nanostructures. *Computer Graphics Forum* 33, 8 (2014), 177–188. DOI:https://doi.org/10.1111/cgf.12425

Disney. 2011. BRDF Explorer. https://www.disneyanimation.com/technology/brdf.html. (2011).

Serkan Ergun, Sermet Önel, and Aydin Ozturk. 2016. A General Micro-flake Model for Predicting the Appearance of Car Paint. In *Eurographics Symposium on Rendering.*

Sergey Ershov, Konstantin Kolchin, and Karol Myszkowski. 2001. Rendering Pearlescent Appearance Based On Paint-Composition Modelling. *Computer Graphics Forum* 20, 3 (sep 2001), 227–238. DOI:https://doi.org/10.1111/1467-8659.00515

Xavier Granier and Wolfgang Heidrich. 2003. A simple layered RGB BRDF model. *Graphical Models* 65, 4 (2003), 171–184. DOI:https://doi.org/10.1016/S1524-0703(03)00042-0

Ole Gulbrandsen. 2014. Artist Friendly Metallic Fresnel. *Journal of Computer Graphics Techniques* 3, 4 (2014), 64–72. http://jcgt.org/published/0003/04/03/

Stephane Guy and Cyril Soler. 2004. Graphics Gems Revisited: Fast and Physically-based Rendering of Gemstones. *ACM Transactions on Graphics* 23, 3 (Aug. 2004), 231–238.

JM Hammersley and KW Morton. 1956. A new Monte Carlo technique: antithetic variates. In *Mathematical proceedings of the Cambridge philosophical society*, Vol. 52. 449–475.

Xiao D. He, Kenneth E. Torrance, François X. Sillion, and Donald P. Greenberg. 1991. A Comprehensive Physical Model for Light Reflection. In *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '91)*. ACM, New York, NY, USA, 175–186. DOI:https://doi.org/10.1145/122718.122738

Eugene Hecht. 2001. *Optics* (4th edition ed.). Addison-Wesley.

Isabelle Icart and Didier Arqués. 1999. An Illumination Model for a System of Isotropic Substrate- Isotropic Thin Film with Identical Rough Boundaries. In *Eurographics Workshop on Rendering*, Dani Lischinski and Greg Ward Larson (Eds.). The Eurographics Association. DOI:https://doi.org/10.2312/EGWR/EGWR99/261-272

Wenzel Jakob. 2010. Mitsuba renderer. (2010). http://www.mitsuba-renderer.org.

Wenzel Jakob, Eugene D'Eon, Otto Jakob, and Steve Marschner. 2014. A comprehensive framework for rendering layered materials. *ACM Transactions on Graphics* 33, 4 (jul 2014), 1–14. DOI:https://doi.org/10.1145/2601097.2601139

James T. Kajiya. 1986. The Rendering Equation. In *ACM SIGGRAPH*, Vol. 20. 143–150.

Jan Kautz, Pere-Pau Vázquez, Wolfgang Heidrich, and Hans-Peter Seidel. 2000. Unified Approach to Prefiltered Environment Maps. In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000*. Springer-Verlag, London, UK, UK, 185–196. http://dl.acm.org/citation.cfm?id=647652.732274

Johannes Meng, Florian Simon, Johannes Hanika, and Carsten Dachsbacher. 2015. Physically Meaningful Rendering using Tristimulus Colours. *Computer Graphics Forum* 34, 4 (2015), 31–40.

Brian E. Smits and Gary W. Meyer. 1992. Newton's Colors: Simulating Interference Phenomena in Realistic Image Synthesis. In *Photorealism in Computer Graphics*. Springer, 185–194. DOI:https://doi.org/10.1007/978-3-662-09287-3_13

Jos Stam. 1999. Diffraction Shaders. In *ACM SIGGRAPH*.

Yinlong Sun. 2006. Rendering Biological Iridescences with RGB-based Renderers. *ACM Transactions on Graphics* 25, 1 (Jan. 2006), 100–129.

Romain Vergne and Pascal Barla. 2015. Designing Gratin, A GPU-Tailored Node-Based System. *Journal of Computer Graphics Techniques (JCGT)* 4, 4 (19 November 2015), 54–71. http://jcgt.org/published/0004/04/03/

Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. 2007. Microfacet Models for Refraction Through Rough Surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques (EGSR'07)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 195–206. DOI:https://doi.org/10.2312/EGWR/EGSR07/195-206

Gregory J. Ward. 1992. Measuring and Modeling Anisotropic Reflection. *SIGGRAPH Comput. Graph.* 26, 2 (July 1992), 265–272. DOI:https://doi.org/10.1145/142920.134078

Andrea Weidlich and Alexander Wilkie. 2007. Arbitrarily Layered Micro-facet Surfaces. In *ACM GRAPHITE*. 171–178.

A. Wilkie, S. Nawaz, M. Droske, A. Weidlich, and J. Hanika. 2014. Hero Wavelength Spectral Sampling. *Computer Graphics Forum* 33, 4 (jul 2014), 123–131.

Pochi Yeh. 2005. *Optical Waves in Layered Media*. Wiley.

APPENDIX

Derivation of C_m and R in Closed Form. C_0 is given by:

$$C_0 = c_0^2 + \sum_{k=1}^{\infty} c_k^2 = R_{12} + \frac{T_{12}T_{21}R_{23}}{1 - R_{23}R_{21}} = R_{12} + R_\star,$$

with $R_{ab} = |r_{ab}|^2$ and $T_{ab} = \frac{\eta_b \cos \theta_b}{\eta_a \cos \theta_a}|t_{ab}|^2$ denoting Fresnel reflectances and transmittances respectively. The derivation is similar to that of Equation 3, and we have used $T_{12}T_{21} = |t_{12}|^2|t_{21}|^2$ to express C_0 uniquely in terms of reflectances and transmittances.

For C_m, we must first separate the 0th order terms from the others:

$$C_m = c_m c_0 + \sum_{k=1}^{+\infty} c_k c_{k+m}$$
$$= c_m c_0 + \left(\sqrt{R_{23}R_{21}}\right)^m \sum_{k=1}^{+\infty} c_k^2,$$

with $c_m c_0 = -\sqrt{T_{12}T_{21}}\left(\sqrt{R_{23}R_{21}}\right)^m$ and $\sum_{k=1}^{+\infty} c_k^2 = R_\star$. We may now factorize by $\left(\sqrt{R_{23}R_{21}}\right)^m$ to obtain C_m in closed form:

$$C_m = \left(\sqrt{R_{23}R_{21}}\right)^m \left(R_\star - \sqrt{T_{12}T_{21}}\right).$$

The C_m terms for $m \in \{0..3\}$ are shown as functions of the incidence angle θ_1 in Figure 22.

To derive R in closed form, we first explicitly write Equation 5 in terms of Fresnel reflectances/transmittances and phase shifts as:

$$R = R_{12} + R_\star + 2(R_\star - \sqrt{T_{12}T_{21}}) \sum_{m=1}^{+\infty} (R_{\updownarrow})^m \cos(m\Phi),$$

where we have used $R_\updownarrow = \sqrt{R_{21}R_{23}}$ as a short notation. Now, applying Euler's formula $\cos(m\Phi) = \frac{1}{2}(e^{im\Phi} + e^{-im\Phi})$, we obtain

$$\sum_{m=1}^{+\infty} (R_\updownarrow)^m \cos(m\Phi) = \frac{P_+ + P_-}{2},$$

with $P_\pm = \sum_{m=1}^{+\infty} (R_\updownarrow e^{\pm i\Phi})^m = \frac{R_\updownarrow e^{\pm i\Phi}}{1 - R_\updownarrow e^{\pm i\Phi}}$. The sum $P_+ + P_-$ may be further simplified to yield:

$$R = R_{12} + R_\star + \frac{2\left(R_\updownarrow \cos\Phi - R_\updownarrow^2\right)}{1 - 2R_\updownarrow \cos\Phi + R_\updownarrow^2}\left(R_\star - \sqrt{T_{12}T_{21}}\right).$$


Fig. 22. The C_m terms are visualized for both polarizations (S for perpendicular, P for parallel), as a function of the incidence angle θ_1, for a material defined by $\eta_2 = 1.5$, $\eta_3 = 1.2$ and $\kappa_3 = 0.5$. All terms but C_0 vanish at $\theta_1 = \pi/2$.

The Case of Transmission. Assuming a dielectric base layer of real index η_3, the spectral Airy transmittance $T = \frac{\eta_3 \cos\theta_3}{\eta_1 \cos\theta_1}|\mathbf{t}|^2$ is computed using the following summation formula:

$$\mathbf{t} = e^{i\delta\phi} \sum_{k=0}^{\infty} t_{12}\left[\mathbf{r}_{23}\mathbf{r}_{21}\right]^k t_{23} e^{ik\Delta\phi}$$

where $\delta\phi$ is the phase shift between the primary reflected ray and all transmitted rays (it will get simplified later on).

As explained in our supplemental material, $\Delta\phi$ has the same expression in the reflection and transmission cases. However, the c_k and ϕ_k terms take on different forms:

$$\begin{aligned} c_k &= t_{12}\left[r_{23}r_{21}\right]^k t_{23} \\ \phi_k &= \delta\phi + k\left(\Delta\phi + \phi_{23} + \phi_{21}\right). \end{aligned}$$

Following the same steps as in Section 4, we obtain:

$$T = C_0 + 2\sum_{m=1}^{\infty} C_m \cos(m\Phi),$$

where $\Phi = \Delta\phi + \phi_{23} + \phi_{21}$ as before, since $\delta\phi$ vanishes. However, the C_m terms are now given (for all $m \geq 0$) by:

$$C_m = \left(\sqrt{R_{23}R_{21}}\right)^m T_\star \quad \text{with } T_\star = \frac{T_{12}T_{23}}{1 - R_{23}R_{21}},$$

where we have used $T_{12}T_{23} = \frac{\eta_3 \cos\theta_3}{\eta_1 \cos\theta_1}|t_{12}|^2|t_{23}|^2$.

Since the spectral Airy reflectance and transmittance have the same form, it follows that the analytic spectral integration of Equation 10 is also valid for the case of transmission. The resulting Airy

transmittance term may thus be used in place of the classic Fresnel transmittance term in any microfacet-based BTDF model.

The spectral Airy transmittance is also given in closed form by:

$$T = T_\star\left(1 + \frac{2\left(R_\updownarrow \cos\Phi - R_\updownarrow^2\right)}{1 - 2R_\updownarrow \cos\Phi + R_\updownarrow^2}\right),$$

following a derivation similar to the one used for the Airy reflectance.