CS284A Final Project: Rendering Iridescence

Jason Li (CS184)
University of California - Berkeley
Berkeley, California
jjli@berkeley.edu

Katherine Song (CS284A)
University of California - Berkeley
Berkeley, California
kwsong@berkeley.edu

Alex Yang (CS184)
University of California - Berkeley
Berkeley, California
yangalex@berkeley.edu

Figure 1: Examples of iridescence in nature.

ABSTRACT

In this project, we render objects that exhibit iridescence from thin-film interference. We derive a first-order approximation of complex reflectance from a thin-film coating on various objects and apply this to multiple BRDF material models. From this approach, we generate various images that demonstrate that our model allows us to render a variety of surfaces that have dramatically iridescent properties. We also model the effect of film thickness variation across the surface of an object.

CCS CONCEPTS

• **Computing methodologies → Reflectance modeling**.

KEYWORDS

graphics, iridescence, thin-film interference

1 INTRODUCTION

Iridescence is an optical phenomenon of certain materials in which nanoscale structures at/near the surface of the material (e.g. one or more layers of thin films or diffraction gratings) create interference among waves of light that hit the surface of the material. This causes the color of each point at the surface to change, sometimes dramatically, across different viewing angles. As a result, objects appear"rainbowed" and dynamic as they move through space (or as the viewer's position changes). Accounting for iridescence is needed in computer graphics to render surfaces such as soap bubbles, butterfly wings, some minerals, and some fish scales. Some real-life examples of these are shown in Figure 1.

In this project, we aim to render iridescence that arises specifically due to thin-film interference. One challenge of rendering iridescence is that we must treat light as a wave (i.e. more than just a simple vector with magnitude and direction) to track phase (color) shifts that occur when a wave of light bounces between the surfaces of a thin film before finally exiting to travel to the viewer.

2 RELATED WORK

Several other efforts have been made to render iridescence. For example, Hirayama *et al.* modeled and rendered iridescence that arises from multiple stacked films, which is important for accurately rendering complicated structures with subtle effects, such as butterfly wings [Hirayama et al. 2000]. In 2017, Belcour and Barla demonstrated a real-time method of rendering iridescent specular highlights on leather by extending a microfacet BRDF model [Belcour and Barla 2017]. These and most other published works emphasize the issue that because the phase that accumulates from traveling through a thin film is wavelength-dependent, achieving iridescent effects in computer graphics calls for spectral integration or sampling, which can become computationally expensive and require clever techniques to be able to render scenes in reasonable amounts of time[Gondek et al. 1994; Sun 2006].

While these approaches are extremely valuable when it is necessary to reproduce coloration as accurately as possible, it is not always possible to know all of the material properties needed as inputs for these models. Additionally, there exist some applications for which iridescent effects are desired, but it is not necessary for the effects to perfect mirror the real-life effects that would be achieved with exact parameter inputs. For such applications, the overhead associated with many existing models is too large to make implementation worthwhile.

Team 39, April 28, 2020, Berkeley, CA

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/none

Here, we show that we can achieve iridescent effects with a first-order approximation of the complex reflectance that results from a single thin film coating on a material even without spectral integration or complex photon mapping, providing a much simpler approach that may be suited for many applications.

3 APPROACH

3.1 Material modeling

We start with the codebase from Assignments 3-1 and 3-2 of this course to incorporate our new iridescent BRDF. Modeling iridescence requires that we account for two main issues. First, since each incoming light ray must be treated as a wave with both magnitude and phase, we need to derive an expression for complex reflectance and transmission terms that properly track changes in phase. Secondly, because each incoming ray is split into components that may directly reflect off of one surface or may bounce between surfaces several times before exiting to the viewer, we must sum across multiple wave components to find the total reflectance instead of simply tracing a single ray. Furthermore, since the total phase change of a wave as it travels between the two boundaries of the thin film is wavelength-dependent, we must perform this procedure for every color channel. Ideally, we would at least sample across multiple wavelengths within a color channel, but we found that even just applying this approach to 1 wavelength per channel was sufficient to demonstrating the effects we intended, as will be shown in subsequent sections.

Inspired by parts of existing approaches in literature, we first extended a microfacet material model to include thin-film interference effects [Bacterius 2013; Belcour and Barla 2017]. We start with our equation for a microfacet BRDF:

$$f(\omega_i, \omega_o) = \frac{F(h \cdot \omega_i) * G(\omega_i, \omega_o) * D(h)}{4 * (n \cdot \omega_i) * (n \cdot \omega_o)} \tag{1}$$

We leave the shadow-masking term, $G(\omega_i, \omega_o)$, and the distribution of normals, $D(h)$, unmodified and focus on adapting the Fresnel term $F(h \cdot \omega_i)$.

We need to replace the Fresnel reflectance F with a complex term that is a summation of all the waves exiting into the external medium that are:

- reflected directly off the the air-film interface
- transmitted through the thin film, reflected at the film-material interface, transmitted through the thin film and then finally transmitted back through the air-film interface
- transmitted through the thin film, reflected at the film-material interface, reflected at the film-air interface, etc. for 1 or more times before being transmitted through the air-film interface back to the viewer

A schematic of the path of such waves is shown in Figure 2.

We do this first by noting the Fresnel equations for computing the reflected and transmitted amplitudes of the components of a wave traveling from medium i to medium j:

$$r_{s,i|j} = \frac{n_i cos(\theta_i) - n_j cos(\theta_j)}{n_i cos(\theta_i) + n_j cos(\theta_j)} \tag{2}$$

$$t_{s,i|j} = \frac{2n_i cos(\theta_i)}{n_i cos(\theta_i) + n_j cos(\theta_j)} \tag{3}$$

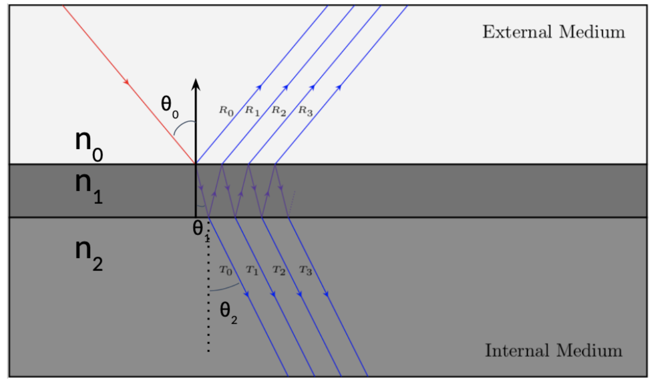

Figure 2: Thin-film interference [Bacterius 2013].

$$r_{p,i|j} = \frac{n_j cos(\theta_i) - n_i cos(\theta_j)}{n_i cos(\theta_j) + n_j cos(\theta_i)} \tag{4}$$

$$t_{p,i|j} = \frac{2n_i cos(\theta_i)}{n_i cos(\theta_j) + n_j cos(\theta_i)} \tag{5}$$

We assume a 1:1 mix of s and p polarized light.

We can use these terms to compute the amplitude of each transmitted component. We can see that T_0 in Figure 2 will have an amplitude $t_{0|1}t_{0|1}$. Tracing T_1, we see that it has an amplitude $t_{0|1}r_{1|2}r_{1|0}t_{1|2} = t_{0|1}t_{1|2}r_{1|0}r_{1|2}$. More generally, the kth transmitted wave will have an amplitude:

$$A_k = t_{0|1}t_{1|2}r_{1|0}^k r_{1|2}^k \tag{6}$$

To calculate phase changes, we must account for the phase change that occurs due to optical path length and those due to reflections at boundaries. Mathematically, the phase of the kth transmitted wave, ψ_n, is:

$$\psi_k = k[\frac{2\pi}{\lambda}(2n_1 \delta cos\theta_1) + \Delta)] \tag{7}$$

where λ is the wavelength of the wave, δ is the thickness of the film, and Δ is the sum of 180-degree phase changes that occur when the wave travels through a medium and reflects upon a medium that has a lower index of refraction (i.e. $\Delta_{i|j} = \pi$ for $n_i > n_j$). The $\frac{2\pi}{\lambda}(2n_1\delta cos\theta_1)$ term reflects the phase change that accumulates from optical path length.

We then can express our estimate for our total transmitted intensity, T, as an infinite sum of waves with the above amplitudes and phases, with a normalization term to account for conservation of energy:

$$\begin{aligned} T &= \frac{n_2 cos\theta_2}{n_0 cos\theta_0} |\sum_{k=1}^{\infty} A_k e^{ik\psi}|^2 \\ &= \frac{n_2 cos\theta_2}{n_0 cos\theta_0} |\sum_{k=1}^{\infty} t_{0|1}t_{1|2}r_{1|0}^k r_{1|2}^k e^{ik\psi}|^2 \\ &= \frac{n_2 cos\theta_2}{n_0 cos\theta_0} |t_{0|1}t_{1|2}\sum_{k=1}^{\infty}(r_{1|0}r_{1|2}e^{i\psi})^k|^2 \end{aligned} \tag{8}$$

The above expression is a geometric sum that can be evaluated as:

$$T = \frac{n_2 cos\theta_2}{n_0 cos\theta_0} \left| \frac{t_{0|1} t_{1|2}}{1 - (r_{1|0} r_{1|2}) e^{i\psi}} \right|^2 \tag{9}$$

Finally, we find our reflectance term R by noting

$$T + R = 1 \tag{10}$$

This complex reflectance term accounts for the fact that rays carry phase information in addition to magnitude and direction, which enables us to model the constructive or destructive interference of the rays that travel back towards the viewer. It is a first-order approximation that does not account for higher-order subtleties, but as we will see, it is sufficient to demonstrating iridescent effects. To more accurately reflection and transmission coefficients, we can refer to [Born et al. 1999] and the Supplemental Materials of [Belcour and Barla 2017] for a more detailed dependence on the film and material paramenters.

This approach may similarly be applied to add thin-film coatings to other materials beyond conductive microfacets. For instance, when modeling glassy materials, the Fresnel reflectance is conventionally used to determine the ratio of reflected to refracted light. We can replace the simple Fresnel reflectance by our modified thin-film interference reflection to render iridescence on such materials.

3.2 Modeling film thickness variation

Thin films on objects in real life exhibit small variations in thickness that also contribute to irregular iridescent effects. To do this, we incorporate Perlin noise to vary the thickness of our modeled thin films when sampling our BRDF [Perlin 2002]. The film thickness at a sampled point is adjusted to be

$$thickness = \frac{noise + scaling}{scaling} \tag{11}$$

We empirically find that $scaling = 40$ yielded the best aesthetic results (i.e. variations in film thickness are 20 times smaller than the thickness of the film).

3.3 Applying our model

To test our model, we created scenes by adding thin-film information (index of refraction and film thickness) to microfacet `.dae` scenes from Assignment 3-2. We additionally modified the Collada parser to handle such files.

4 RESULTS

4.1 Microfacet materials

A rendering of a microfacet `CBdragon.dae` with a 550nm thin film soap coating (n=1.33) is shown in Figure 3. The walls of the Cornell box are all colored white to avoid the confounding effects of reflection from the conventional red and blue walls falsely coloring the dragon. We indeed do see dramatic variations in color that would be expected from an iridescent film.

An animation demonstrating the changes in color as the viewer moves with respect to the dragon is provided on our project website.

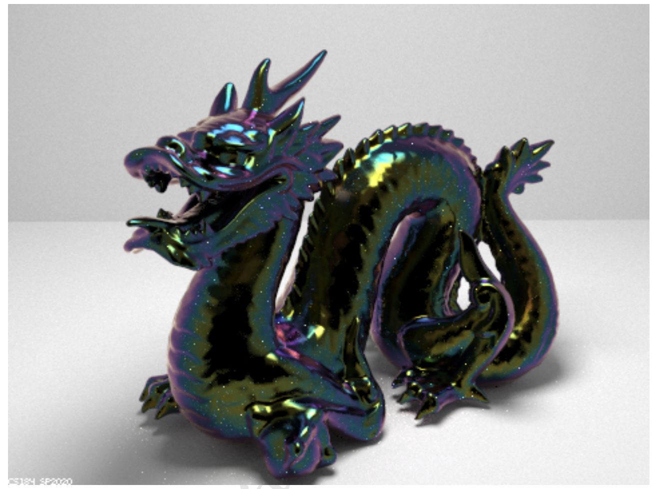

Figure 3: CBdragon scene modeled with thin-film iridescence

4.2 Other materials

We can also use our calculated reflectance to model thin films on materials other than microfacet conductors. By using our calculated reflectance as a probability of whether an incoming wave will reflect or refract, we can also model thin films on transparent materials such as glass or air. Figure 4 shows the results of applying this approach.

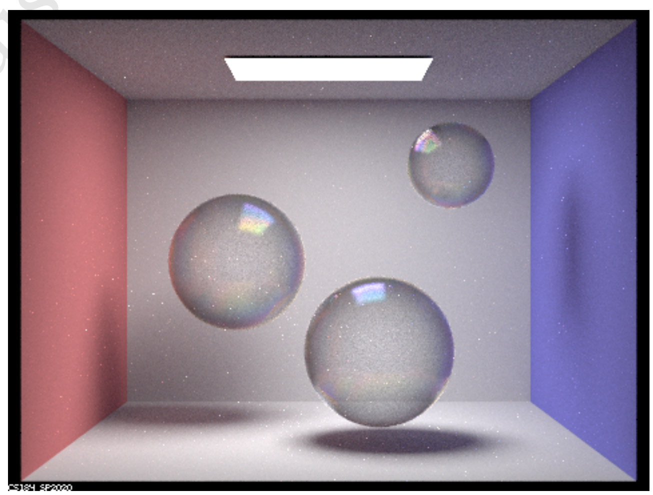

Figure 4: Rendering of a soapy thin film on a transparent material

4.3 Film thickness variation

As the thickness of the thin film changes, different iridescent effects can be achieved. This is illustrated in Figure 5, which shows CBdragon with 3 different film thicknesses that are constant over the surface of an object.

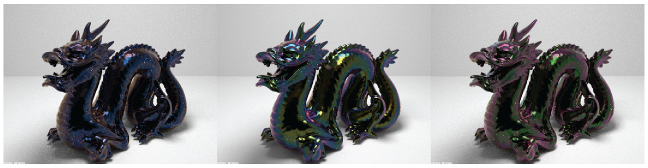

Figure 5: CBdragon rendered with a 250nm soap layer (left), a 550nm soap layer (middle), and a 750nm soap layer (right)

In Figure 5, the thickness of the film on the dragon is constant across the dragon's entire body. Of course, in reality, the thickness of a thin-film is rarely constant due to sagging from gravity, irregularities on the surface of the material that cause disruptions in surface tension, etc. This leads to second-order variations in color across the object. Figure 6 shows the effect of adding Perlin noise to model such variations in thin-film thickness.

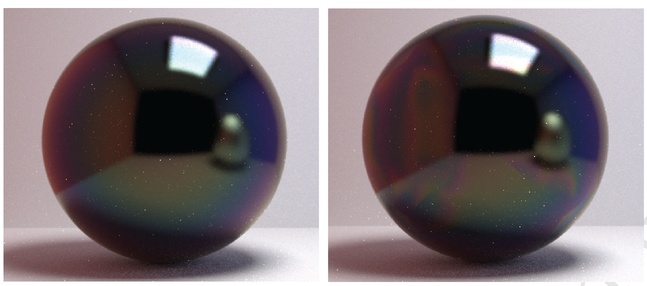

Figure 6: Left: before incorporating variations in thin-film thickness. Right: after incorporating variations in thickness based on Perlin noise.

5 EVALUATION

As can be seen in results shown above, our approach successfully captures the variation in color across an object that thin-film interference effects would cause. However, because we only use 1 wavelength sample per color channel, the actual colors of our object are not entirely physically accurate and tend to be more saturated than one might expect. A more advanced approach that integrates across multiple wavelengths within a channel is explored in literature and can achieve colorations that are less saturated and more realistic [Belcour and Barla 2017]. However, we believe that our approach is sufficient for applications that do not require exact color matching.

6 CONCLUSION

We demonstrate a simple iridescence model that we integrate into material BRDFs in a path tracing program. Although we do not implement spectral shading, we are able to render various objects and materials with thin-film coatings that show variations in color across viewing angle, as we expect. In the future, our approach can be easily ported to a shader program for fast, real-time rendering as well.

7 SUPPLEMENTAL MATERIAL

7.1 Video

`https://youtu.be/n6I63s2jexM`

7.2 Website

`https://kwsong.github.io/cs184_final/final.html`

7.3 Team Contributions

Jason: Thin-film thickness variation modeling via Perlin noise, applying model to transparent materials, debugging
Katherine: Thin-film reflectance modeling and implementation, image generation, documentation
Alex: Scene generation, custom COLLADA file handling, preliminary reflectance modeling

REFERENCES

Bacterius. 2013. Thin Film Interference for Computer Graphics. https://gamedev.net/tutorials/_/technical/graphics-programming-and-theory/thin-film-interference-for-computer-graphics-r2962/. Accessed: 2020-05-12.

Laurent Belcour and Pascal Barla. 2017. A Practical Extension to Microfacet Theory for the Modeling of Varying Iridescence. *ACM Trans. Graph.* 36, 4, Article 65 (July 2017), 14 pages. https://doi.org/10.1145/3072959.3073620

Max Born, Emil Wolf, A. B. Bhatia, P. C. Clemmow, D. Gabor, A. R. Stokes, A. M. Taylor, P. A. Wayman, and W. L. Wilcock. 1999. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light* (7 ed.). Cambridge University Press. https://doi.org/10.1017/CBO9781139644181

Jay S. Gondek, Gary W. Meyer, and Jonathan G. Newman. 1994. Wavelength Dependent Reflectance Functions. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '94)*. Association for Computing Machinery, New York, NY, USA, 213–220. https://doi.org/10.1145/192161.192202

H. Hirayama, Y. Yamaji, K. Kaneda, H. Yamashita, and Y. Monden. 2000. Rendering iridescent colors appearing on natural objects. In *Proceedings the Eighth Pacific Conference on Computer Graphics and Applications*. 15–433.

Ken Perlin. 2002. Improving Noise. *ACM Trans. Graph.* 21, 3 (July 2002), 681–682. https://doi.org/10.1145/566654.566636

Yinlong Sun. 2006. Rendering Biological Iridescences with RGB-Based Renderers. *ACM Trans. Graph.* 25, 1 (Jan. 2006), 100–129. https://doi.org/10.1145/1122501.1122506