SIGGRAPH Asia 2009 Course

Exploring the Potential of Layered BRDF Models

Andrea Weidlich

Alexander Wilkie

September 9, 2009

Abstract

This course serves as a guide on the considerable potential of layered surface models. The key advantage of using such layered BRDFs over traditional, more general shading language constructs is that the end result is automatically highly physically plausible.

In particular, we demonstrate on a simple layered surface model that combines several traditional BRDF components how a surprisingly large number of interesting and important surface types can be efficiently represented by using the same, not particularly complex, BRDF code. We also show how handy such an approach is for the eventual end user, whose main concern is the ease with which one can describe object appearance based only on a few intuitive parameters.

We first discuss layered surface models in computer graphics and the constraints of modelling object appearance in a physically plausible fashion. We then demonstrate the techniques that can be used to efficiently evaluate layered BRDF models, give examples of the surface types that can be described in this way. We also go beyond plain surface models, and showcase how a texture-based combination of layered surface components can be used to describe highly complex object appearance attributes, while implicitly remaining physically plausible.

Course Overview

1 minute: Welcome and Introduction

Andrea Weidlich and Alexander Wilkie

Overview of the course, and motivation for attending it.

29 minutes: Layered Surfaces in Computer Graphics

Alexander Wilkie

This part of the course outlines the main differences to traditional, shader-language based techniques to describe object appearance. Several different state-of-the-art layered surface models will be compared. We will also put layered BRDFs in a broader, comparative context with analytic methods or measured BRDFs of more complex materials and discuss when is it appropriate to use them.

35 minutes: Combining Individual BRDFs into Layered Models

Andrea Weidlich

The technical core of the course: while the idea of layering BRDFs is extremely simple, a few things are necessary to do this properly and efficiently in a physically-based stochastic renderer. We will further present a layered surface model that is comparably fast as well as easy to implement. Although the model was originally presented in the context of a stochastic rendering environment, we will show how it can be modelled with a modern shader language.

15 minutes: Classifying Materials - Using Layered BRDFs to Describe Object Appearance

Andrea Weidlich

In this part we will classify the appearance of different materials into several groups according to their reflectance properties. This is necessary since we later use this groups as a guideline for modelling.

25 minutes: Modelling with Layered Surfaces

Alexander Wilkie

This part of the course showcases the power of layered BRDFs in a practical setting. For some types of surface where BRDF measurements are available, we also discuss the performance of the layered model relative to these measurements, and also compared to other, simpler combined BRDFs. For each type, we discuss the applications, and the limits of the layered approach.

About the Lecturers

Andrea Weidlich

Vienna University of Technology, Austria weidlich@cg.tuwien.ac.at

Andrea Weidlich is currently an assistant professor at Vienna University of Technology, Austria, Europe, from where she graduated in 2005 with a M.Sc. in computer science and a M.Sc. in computer science management. She has defend her Ph.D. in March 2009. Her current research interests include predictive rendering with a special focus on gemstone prototyping and appearance modeling.

Alexander Wilkie

Charles University in Prague, Czech Republic alexander@wilkie.at

Alexander Wilkie is currently a senior lecturer at Charles University in Prague. He achieved his habilitation in Applied Computer Science from Vienna University of Technology in June 2008, and spent the time from 2000 to 2008 as an assistant professor at that university. He obtained both his master's degree (1996) and Ph.D. (2001) in computer science there. His main research interests are predictive rendering, colour science, and appearance modelling. He has authored numerous peer-reviewed papers about computer graphics, and has extensive teaching experience in the areas of photorealistic rendering and colour science, which he obtained both at Vienna University of Technology, and now at Charles University.

Contents

1	Intro	oduction									5
	1.1	Motivation						•	•	•	. 5
	1.2	Overview of the Course Material	•		•			•	•	•	. 6
2	Laye	ered Surfaces in Computer Graphics									7
	2.1	BRDF Models									. 7
		2.1.1 Requirements for Analytical BRDFs									. 7
		2.1.2 Traditional Analytical Reflectance Models						•			. 8
		2.1.3 Multi-Layer Reflectance Models			•						. 9
		2.1.4 Differences to Existing Techniques	•		•			•			. 10
3	Com	nbining Individual BRDFs into Layered Models									11
	3.1	Arbitrarily Layered Micro-Facet Surfaces									. 11
		3.1.1 Simplification of the Problem									. 11
		3.1.2 Overview of the Model									. 12
	3.2	Using the Approach in a Monte Carlo Image Renderer						•			. 15
		3.2.1 Path Propagation vs. BRDF Evaluation						•			. 15
		3.2.2 Global Reflection Model						•			. 15
		3.2.3 Local Reflection Model						•			. 16
		3.2.4 Sampling PDFs						•			. 16
	3.3	RenderMan SL Implementation						•			. 17
		3.3.1 Basic Components						•			. 17
		3.3.2 Absorption						•			. 19
		3.3.3 Internal Reflection			•			•			. 20
		3.3.4 Final Shader						•			. 21
	3.4	Simplifications for Real-Time Rendering						•			. 21
		3.4.1 Precomputations and Approximations						•			. 21
		3.4.2 Performance						•			. 22
	3.5	Overall Performance in a Rendering System	•		•	•••	• •	•	•	•	. 22
4	Clas	ssifying Materials - Using Layered BRDFs to Describe Object	t A	pp	ear	an	ce				24
	4.1	Material Appearance and Visual Clues			•			•			. 24
		4.1.1 Metallic vs. Non-Metallic Materials						•	•		. 24
		4.1.2 Surface Roughness						•			. 25
	4.2	Classifying Materials According to the Reflection	•		•			•	•	•	. 26

5	Mod	lelling v	vith Layered Surfaces	29
	5.1	Traditi	onal Materials	29
		5.1.1	Glossy Paint – Ceramics and Acrylic Lacquer	29
		5.1.2	Frosted Paint – Sugar, Spraypaint and Latex	30
		5.1.3	Metallic Foil – Brass and Christmas Orbs	30
		5.1.4	Metallic Paint – Automotive Paint	31
		5.1.5	Frosted Metal – Metallic Gold and Silver, Blue Metallic Paint	31
		5.1.6	Patina – Silver Paint	31
		5.1.7	Multi-Layer – House Paint, Fabric	31
	5.2	Specia	l Materials	32
		5.2.1	Heterogeneous Surfaces	32
		5.2.2	Fluorescent Surfaces	34

Chapter 1

Introduction

This course material describes the technical details necessary to implement physically plausible layered BRDF models in a modern rendering system, and also gives indications how a simplified form of such models could be included in a modern shader language like Renderman and for real-time rendering applications. It also highlights numerous practical usage scenarios for such analytical reflectance models.

1.1 Motivation

Modern computer graphics have reached astonishing levels of visual realism and artistic expression. As our technological abilities have progressed, several distinct sub-fields, each with its unique technical challenges and goals, have emerged within the discipline; examples are non-photorealistic rendering, scientific visualisation, production techniques geared for the entertainment industry, or physically based rendering. The latter is a comparatively small but highly interesting area, and is not, as one might surmise from the name, only useful for predictive rendering applications. Techniques from this domain can also be utilised in many contexts for believable rendering purposes, mostly to bring selective realism to scenes that are being assembled under artistic control.

Conventional, believable rendering techniques, such as the workflows used in RenderMan or Maya, have reached a level of capability that is apparently quite sufficient even for highly demanding projects. By contrast, physically based rendering has some way to go before it can be considered equally mature and robustly useful for all application domains that it might be considered for. Significant progress has been made over the past years, but some sub-problems, such as efficient and intuitive appearance modelling under the constraints of such rendering systems, or the still considerable performance penalties associated with unbiased global illumination techniques, are not entirely solved yet.

For the purposes of believable rendering, the normal solution to describing the appearance of an object that cannot be directly modelled using a single BRDF is the use of a shader language. While this approach can deliver good-looking results, it is the time-honored mainstay of most production graphics and not always easily to achive.

One solution is to resort to using BRDF measurement data for such "problematic" surfaces instead.

While this is of course a viable solution for some applications, this has the drawback of restricting the creativity of the end user: only previously measured surfaces can be applied to objects. Especially for usage scenarios where physically based techniques are used to augment scenes that are being modelled for artistic purposes, a higher degree of control over object appearance is desirable.

For some of the cases where no simple analytical model exists (such as metallic car paint, or glazed ceramics), layered combinations of several BRDFs could be used instead; this is a fairly obvious approach for all surfaces that have an intrinsic layering. But so far, such layered surface models have not been in particularly widespread use, mainly because the derivation of a compound BRDF is rather difficult for a truly general arrangement of layers, and exact simulations of light transport in such surfaces are quite costly.

The goal of this course is to comprehensively explain one of the recent advances in the area of physically plausible appearance modelling. The technique in question is a high quality approximative solution to the problem of computing aggregate BRDFs for layered surfaces. The course briefly reviews the state of the art and the technological challenges in this sub-field of BRDF construction, to demonstrate how the simple and efficient approximative combination of layered BRDFs that was presented in [37] works in practice, and to showcase the significant potential of this approach by discussing practical usage scenarios for it. The excellent visual quality of the achievable results, coupled with the ease of use of this method, and the fact that it is still physically plausible, make it a potentially valuable asset for all forms of physically-based rendering that do not require total radiometric accuracy. The two big advantages of this approach are that it is physically based (if not entirely physically *correct* in the narrow sense of the word because of the approximations that are introduced into the evaluation) and therefore produces feasible looking results, and intuitive to use at comparable low computational costs.

The idea to submit such a course occurred to us because since the time the paper was published, we have received several requests from engineers in the graphics industry to clarify minor points of the proposed layering technique. To us, this suggests that a) the method is of industrial relevance (the requests were all of the form "*we are trying to implement this in our system, could you help us with some details*?"), and that b) it apparently could do with some further explanations that simply did not fit into the original 8 pages of the conference paper. These space restrictions also prevented that we fully expanded on the usage possibilites of this kind of technique, which are quite numerous, and which we feel might be of interest to the larger SIGGRAPH community.

1.2 Overview of the Course Material

As a means of preparing the ground for the subsequent treatment of the layering, chapter 2 first discusses the general technical issues that have to be solved if one wants to implement a physically plausible reflectance models. In this chapter, we also briefly review those BRDF models that are known to be reasonably physically accurate. Chapter 3 discusses the actual technique that enables us to define layered surfaces in a fairly free, but still physically plausible, way. Chapter 4 Chapter 5 finally showcases the numerous surface types for which such a layered model can be used, and discusses the potential issues with them. We also compare some of the results with measured BRDFs and go beyond plain, homogenous surfaces and explore the use of such layered BRDFs for the characterisation of the appearance of complex, textured objects.

Chapter 2

Layered Surfaces in Computer Graphics

In this chapter, we briefly review known layered BRDF models and discuss the constraints that physically plausible BRDF models have to fulfil (energy conservation, reciprocity, etc.). While the salient points are of course also covered in textbooks on global illumination renderers, the discussion of the actual layering algorithm references these concepts quite heavily, so a re-cap of the relevant theory is given here along the following lines.

2.1 BRDF Models

Computer Graphics uses two distinct approaches to deliver realistic surface reflectance behaviour. The first one is the measurement of real surfaces, and the subsequent use of the gathered reflectance data in the rendering process. The second one is the derivation and use of analytical models of varying realism. In this course we focus on the analytical approach.

During the last decades a large number of analytical reflectance models have been developed for computer graphics use. These models can roughly be divided into two groups:

- 1. *empirical* and often fundamentally physically implausible models, which deliver reasonably good-looking results at moderate computational cost, and
- 2. those where comparatively expensive, *physically plausible* computations of light interacting with matter are used for highly convincing depictions of surfaces.

Especially the more sophisticated specimens of the second category offer a great potential for creating very convincing renderings, but have to be handled with some care. Also, there is only a limited number of BRDF models that are entirely physically plausible, and not all types of common surfaces can be described with them.

2.1.1 Requirements for Analytical BRDFs

BRDFs have to two important physical constraints, namely the law of conservation of energy, and the Helmholtz reciprocity rule. The first one requires that the radiant power reflected in all directions must be less than (or, in certain rare cases, equal) to the radiant power of the incident light. The second

condition ensures the symmetry of V and V'. Mathematical representations of BRDFs ought to fulfil the following criteria:

- 1. use reasonable amounts of storage
- 2. faithfully capture the key features of the reflection characteristics
- 3. permit fast and easy sampling by Monte Carlo methods

Apart from the limit cases of perfectly diffuse surfaces and perfect mirrors, arbitrary reflection properties of a surface are only tractable through MC rendering. The possibility of casting of rays according to distribution function is crucial.

2.1.2 Traditional Analytical Reflectance Models

Models based on a micro-facet approach are normally used to simulate rough surfaces. They assume the surface to consist of a large number of very small statistically distributed micro-facets, which are oriented according to some given probability distribution function, and which can be either isotropic or anisotropic.

Torrance-Sparrow The Torrance-Sparrow model [33] was originally proposed for use in physics, long before computer graphics had reached a point where such a sophisticated technique could be used in a meaningful way. For the purposes of this model, one considers a surface to be a collection of a large number of tiny, symmetric V cavities with two opposing facets. These facets are assumed to be perfect mirrors, the reflectance of which is governed by the Fresnel terms. The model takes mutual masking and shadowing between the facets into account, and their normals are considered to be normally distributed with zero mean. Within the constraints of the model, the specular reflectance term which results from this approach is physically accurate, and is given by

$$f_r = \frac{FDG}{4 \cdot (N \cdot L)(N \cdot V)} \tag{2.1}$$

- *D* is the distribution function of the micro-facets.
- *G* is the geometric attenuation term that influences self-shadowing when the incident light is blocked, and self-masking when the reflected ray is blocked.
- *F* is the Fresnel term for each micro-facet which describes the amount of light that is refracted and reflected.

When considering the Torrance-Sparrow model by itself, one has to keep in mind that it only provides a physically accurate formula for the reflectance off a rough version of the material used for the microfacets. Which considerably limits its immediate usefulness in graphics, since in practice this means that nothing except rough metals and transparent dielectrics can be directly modelled.

Cook-Torrance Micro-facet theory was brought to computer graphics proper in 1982 by Cook and Torrance[6], who introduced a somewhat simplified and refined version of the original Torrance-Sparrow model that had been adapted to graphics use. Here, the total reflected radiance consists of separate specular and a diffuse parts. The specular component represents light waves reflected by

only one micro-facet, and the diffuse component is due to light waves which are reflected by several micro-facets, or which are scattered internally. The overall bi-directional reflectance is the sum of these specular and diffuse components.

Oren-Nayar The model of Oren and Nayar[29] is insofar unique, as it is up to now the only technique to use Lambertian micro-facets - as opposed to the perfect mirror facets of all other approaches. The resulting BRDF is mainly diffuse, with a noticeable retro-reflective component for higher roughness values.

Ward The isotropic version of the Ward model[36] is a simpler version of the Cook-Torrance BRDF, which omits the geometric attenuation and Fresnel coefficients. Ward also introduced an anisotropic version of this model which is – in contrast to the isotropic version – energy conserving, and analytically invertible.

Ashikhmin Like the Ward BRDF, this anisotropic BRDF [4] has diffuse and specular components, although the diffuse term is not Lambertian in this case. Ashikhmin et al. [3] also proposed a second, micro-facet based surface. In this model, the appearance of a given surface is matched by deriving a suitable micro-facet distribution. The main drawback of this otherwise powerful approach is that it is not well suited for surfaces with an appearance that is not dominated by a specular highlight.

2.1.3 Multi-Layer Reflectance Models

Although layered surface models offer a great potential for creating very convincing renderings, they do not appear to have received an extensive amount of attention in computer graphics up to now.

Classical layered surface models are those of Kubelka and Munk [21] and Hanrahan and Krueger [13]. Both these models have no closed mathematical form, and are therefore rarely used in practice – at least for image synthesis purposes.

Neumann and Neumann [27] were amongst the first to propose layered surface models. They discuss two models, one that consists of a single perfectly smooth, transparent layer over an arbitrary surface, and one with an arbitrary number of layers. Both these models include absorption, but not internal reflection. Since they do not give a sampling PDF for their compound BRDFs, and neither a closed expression for the entire BRDF nor an algorithm to compute it, their work has to be considered somewhat incomplete.

Kelemen and Szirmay–Kalos [20] used the Cook-Torrance model in conjunction with layered surfaces, albeit in a simplified form. Their model lacks the ability to simulate absorption and internal reflections, and relies on a simplified variant of the Cook-Torrance model for its specular component. The diffuse component is described by a Lambertian term, and the combination of the two is dependent on the incident angle. They also provide an efficient scheme for sampling the BRDF in a stochastic renderer.

Another relevant surface model is that of Schlick [31]. Although it is not an actual layered model, it should be mentioned here since it can be used to mimic the appearance of layered surfaces fairly well. The appearance of multi-layered models was also indirectly simulated by Lafortune et al. [22]. A

wide number of different surfaces can be reproduced by their technique, but it has two disadvantages: there is a high discrepancy to real surface behaviour near grazing angles, and it is a purely empirical approach.

Wilkie et al. [39] use a multi-layer reflectance model to describe the reflectance properties of diffuse fluorescent surfaces. The combination of a transparent, rough dielectric layer over a normal diffuse surface yields a behaviour that closely corresponds to real diffuse fluorescent objects, such as cardboard.

2.1.4 Differences to Existing Techniques

The considerable body of related work in this area mandates a discussion of the novelty factor of the method discussed in this tutorial:

Combination The multi-lobe approach as suggested by Cook and Torrance works well for some surfaces (see [28]), but cannot be used to describe retro-reflective surfaces, because a backscattering lobe cannot be expressed in a physically correct way by combining several forward scattering lobes. Also, their multi-lobe approach does not work for strongly dissimilar lobes, e.g. a combination of Oren-Nayar and a highly specular lobe. We instead opted to combine suitable BRDFs of arbitrary type – which can include the genuinely retro-reflective Oren-Nayar model if this property is needed – in a physically correct way, and as a consequence can produce a much wider range of BRDFs, than the combination of quite similar forward-scattering lobes could achieve on its own.

Patinas The presented layering technique can be seen as being complementary to, rather than an improvement on, the work of Dorsey and Hanrahan [7]. A large part of their work deals with determining where layers of patina should be placed on a model, and which layers are present in a given patina. Conceivably, one could use our technique to evaluate the resulting compound BRDFs in a GI renderer as an alternative to the method proposed by the original authors.

Weighting We do not use a fixed parameter to determine the ratio between diffuse and specular components of the reflected light, like e.g. the classical Cook-Torrance model does. Since in reality the specularity of a surface increases with decreasing incident angle, such a fixed ratio between these components is physically implausible. Like Ashikhmin and Shirley [4], we use the Fresnel coefficients as a weighting factor instead.

Physical Correctness In contrast to the empirical Lafortune model (which just combines Phong lobes to fit a given BRDF at given angles of incidence), our model combines its components in a physically based fashion through the simulation of multiple surface layers. Also, we go further than Kelemen and Szirmay–Kalos [20], since we use the original, non-simplified versions of the BRDFs we combine.

Chapter 3

Combining Individual BRDFs into Layered Models

In this section, the actual technique of combining layered, physically plausible BRDFs is discussed. It includes the actual source code of a RenderMan implementation. Also, a discussion on how such an approximative layering approach can be used in a real-time system can be found in this section.

3.1 Arbitrarily Layered Micro-Facet Surfaces

As noted earlier, the idea of using individual surfaces of somewhat limited applicability – such as perfect mirrors or Lambertian surfaces – as layered components of a more sophisticated BRDF, is immediately appealing due to its simplicity and usefulness.

While the concept of using layered surfaces is simple, actually using it in a renderer is not – at least if the unrestricted case is considered, because in this case the computation of the entire BRDF involve sub-surface scattering computations within the layers.

3.1.1 Simplification of the Problem

The key to using such surfaces without expensive sub-surface scattering computations is to perform four simplifications:

- 1. Any micro-facets are considered to be much larger in horizontal extent, than the layers are thick.
- 2. All rays that are generated by sampling of lower BRDF layers, are assumed to exit at the original point of incidence.
- 3. Refraction rays that are generated for the computation of the entire BRDF, are assumed to meet at a single point on the next layer interface.
- 4. All light scattering is due to reflection at the boundaries between layers; no scattering occurs within individual layers.

Figure 3.1 gives an overview of the geometrical simplification associated with this process of generating directional samples.

It is worth noting that none of these assumptions is entirely implausible: the first is consistent with the notion of only applying very thin layers atop a base substrate, while the second and third are reasonable simplifications under the circumstances – especially when one considers that the micro-facets involved are assumed to be statistically distributed entities in the first place.

In addition to these three simplifications, one additional restriction is imposed: namely that the material used for any (partially) transparent layers only attenuates light passing through it, and does not contribute any secondary scattering effects of its own. Since most clear and tinted varnishes (which are one of the main targets for this model) do not exhibit noticeable scattering, this is not a particularly hard restriction, though. One also assumes any varnish layers to be homogeneous. The three simplifications, taken together with the restriction to non-scattering varnishes, are what allows one to omit a full sub-surface scattering computation.

It has to be noted, though, that due to the simplifications that are performed, this approach does not constitute a general solution for surface layers of arbitrary thickness and is limited to surfaces that have layers thick enough to have the influence of absorption, and sufficiently thin enough that the simplifications hold true. However, given the high quality of the results obtained by the method, this is not an immediate concern, at least not for those layered surface types which are readily modelled by this approximative approach.



Figure 3.1: Computation geometry and simplification of the sub-surface scattering in a layered surface during BRDF sampling. We assume any micro-facets to be very large in relation to the layer thickness, which allows us to do the following: (1) a ray will always leave through the same micro-facet that it entered through, thereby eliminating the need to perform an intersection test with the nearby micro-facets. And (2), any exitant ray coming from a lower level will emanate from the original point of entry (yellow dot), regardless of how oblique the exitant angle is. The direction of these rays is computed according to the correct geometry (dashed line).

3.1.2 Overview of the Model

The basic idea behind the model is to simulate the physical process of light propagation inside a layered material as close as possible. The process can be summarised as follows.

1. Any light that hits an interface in the layer stack is partly reflected, and partly refracted. The actual amount of energy that will be reflected is determined by the BRDF $f_{r_1}(\theta_i, \theta_r)$ of the topmost surface; in our case the topmost surface is always a dielectric Torrance-Sparrow surface. Therefore the reflection intensity depends on the index of refraction (IOR) of the topmost surface and its roughness as well as the angle of incidence.

- 2. Only a part of the energy that encounters such a material is refracted into the material according to Snell's law. The amount is determined by the Fresnel transmission coefficient T_{12} for the air/material phase interface. A part of this will possibly be absorbed *a* by the varnish material, and the rest then interacts with the next, second surface in the stack. Here again computations are performed how much light is reflected from the surface $f_{r_2}(\theta_{i'}, \theta_{r'})$. Note that we have to use the refracted incoming and outgoing angle this time. If the second layer is another dielectric Torrance-Sparrow surface, this step is repeated until an opaque surface is hit.
- 3. All light that is reflected from lower layers is again attenuated by the varnish on its upward path, and possibly subjected to total internal reflection *t*. This means that any directional samples from lower layers have to be treated accordingly during the return from the recursion.



Figure 3.2: The recursive BRDF evaluation process used by the method; the numbers correspond to the steps described in the text.

Figure 3.2 illustrates the concept. Mathematically this also can be summarised as

$$f_r = f_{r_1}(\theta_i, \theta_r) + T_{12} \cdot f_{r_2}(\theta_{i'}, \theta_{r'}) \cdot a \cdot t$$
(3.1)

with

$$a = e^{-\alpha d \cdot \left(\frac{1}{\cos \theta_{t'}} + \frac{1}{\cos \theta_{r'}}\right).}$$
(3.2)

and

$$t = (1 - G) + T_{21} \cdot G \tag{3.3}$$

Note that the Fresnel reflection is already included in $f_{r_1}(\theta_i, \theta_r)$ in formula 3.1. Figure 3.1 shows a sketch of the simplified reflection and ray propagation geometry. In the following we will discuss the absorption and internal reflection terms in more detail.

Absorption Term a

Usually a part of light is absorbed inside a transparent material by travelling through it. The intensity of the absorption is defined by the Bouguer-Lambert-Beer law. According to that, the loss of intensity I of a light wave travelling through a material with a thickness of l is related to its initial intensity I_0 through

$$I = I_0 e^{-\alpha l} \tag{3.4}$$

where the constant α is the wavelength dependent absorption coefficient which defines the extent to which a material absorbs energy. The length of the path is determined by the thickness of the layer *d* as well as the cosines of the incident angle $\theta_{i'}$ and the outgoing angle $\theta_{r'}$ so that

$$l = d \cdot \left(\frac{1}{\cos \theta_{i'}} + \frac{1}{\cos \theta_{i'}}\right). \tag{3.5}$$

According to this law, internal absorbance, and thus the corresponding colour, increases with *l*. As a consequence the colour decreases and the saturation increases, but additionally it can change the hue. This phenomenon can be seen in action in figure 3.4, and figure 3.3 shows a comparison between the same surface being rendered with and without absorption taken into account.



Figure 3.3: A surface of type **b**) (as defined in figure 4.4) rendered with and without absorption, with otherwise identical colour, illumination and surface parameters. The three difference images are - from left to right - for the hue, chroma and lightness channels of these two images. As expected, there are practically no hue differences, but chroma and lightness change considerably, mainly at grazing angles.



Figure 3.4: Diffuse white spheres with a yellow varnish layer of varying thickness. Layer thickness values are 0.0, 0.2, 0.4, 0.6, 0.8, 1.2, 1.6, 3.0, 5.0, 7.5, 10.0 and 15.0. Note the progressive changes in colour, saturation and hue.

Internal Reflection Term t

On its way backwards the light is again attenuated by the Fresnel transmission coefficient T_{21} , but this for the material/air interface. However, since the index of refraction of the material is bigger than the refractive index of air, i.e. the light propagates from a denser medium into a less dense one, the light is possibly subject to total internal reflection (TIR). This means that if the angle of incidence is greater than the so-called critical angle, no light enters the second medium. If the index of refraction of a medium is very high, a light wave can easily become trapped inside a medium, and total internal reflection may occur several times before even a part of the light wave is able to leave the medium. The remaining part is reflected again and emerges somewhere else.

Ignoring these energy, i.e. culling the rays that are created on lower levels and effected by TIR, would lead to an energy loss. On the other side, an exact computation of the effect is problematic, since it would require an explicit simulation of the incident ray interacting with the micro-facet geometry multiple times. However, these computations would be rather expensive, so we use an approximation that works as follows.

We attenuate the light that is reflected on a lower level by T_{21} , but compensate the lost energy by adding a term that accounts for the energy that is not able to immediately leave the medium. This term we use for this purpose is the part of the energy that is blocked by the geometric attenuation factor *G* due to shadowing and masking, i.e. 1 - G. The idea behind this is that the light that is blocked by the geometric attenuation factor should actually enter the material for dielectric surfaces and return somewhere else.

Conceivably, one could also re-start the downward propagation for the ray that is subjected to TIR, although it remains to be seen whether the difference to the results obtained with the energy compensation term would be worth the additional computational effort since the overall influence of the term t is rather small, but increases with increasing index of refraction.

3.2 Using the Approach in a Monte Carlo Image Renderer

While images of perfectly diffuse surfaces and perfect mirrors can be rendered by a deterministic raytracer, surfaces with arbitrary reflection properties are basically only tractable through Monte Carlo rendering, e.g. bi-directional path tracing or photon tracing. For this purpose, we have to be able to evaluate the reflection model globally and locally.

3.2.1 Path Propagation vs. BRDF Evaluation

It is rarely explicitly mentioned in rendering literature that one has to be able to perform two distinct functions for each reflectance model one wishes to include in a stochastic renderer.

- The first concerns the ability to correctly continue an incoming path according to a chosen sampling PDF. This functionality is all that is needed for a primitive path tracer, and is rather easy to perform even for complicated multi-layer surfaces: for each layer a suitable propagation direction can be recursively calculated, and is weighted according to its sampling probability. One of these rays is then followed by random selection. This is what we call the *global model*.
- 2. However, one also has to be able to evaluate the *entire*, *combined* BRDF for arbitrary input and output directions for a given point. This is the *local model*, what one would normally associate with a BRDF, and used in shader-based languages or deterministic ray tracers (i.e. Whitted ray tracers) or for lightsource sampling when one wants to determine the influence of a lightsource. For arbitrary multi-layer surfaces the computation of this second piece of information is far from being trivial.

We also need an ability to efficiently sample a given BRDF, i.e. to derive a sampling PDF for it.

3.2.2 Global Reflection Model

The global reflection model is a straightforward implementation of section 3.1.2. The purpose of the global model is to find an outgoing direction and its intensity for a given incoming direction. That can be done as follows.

- 1. For incoming light is reflected on the top surface, a dielectric Torrance Sparrow surface. An outgoing ray is randomly chosen according to the microfacet distribution of the top surface and the reflection intensity is calculated. The ray is stored.
- 2. The incoming ray is refracted according to Snell's law and reflected on the next surface in the stack; an appropriate sampling direction is generated for the reflective component, its reflection intensity possibly attenuated by the absorption of the varnish. Again, the outgoing ray is stored. If there are any further layers, this step is repeated.
- 3. Every ray was created on a sub-level is again refracted according to Snell's law until it reaches the top layer. The reflection intensity is again attenuated by the varnish on its upward path, and possibly subjected to total internal reflection.

This process returns a list with at least one, but possibly also several additional rays.

3.2.3 Local Reflection Model

The computation of the entire BRDF also requires a recursive approach:

- 1. The BRDF of the topmost level f_{r_1} is evaluated for the two given, arbitrary incoming directions ω_i , and ω_o . This yields a reflection component, and, except at the lowest layer, two refraction directions.
- 2. Any energy that is refracted into the next level $R_{t_{12}}$ follows the two refraction directions associated with the initial incident directions, and is partly absorbed *a* by the medium.
- 3. These two refraction directions are assumed to meet at a single point on the next layer f_{r_2} , and the process is repeated from step 1 until an opaque layer without a refraction component is encountered.
- 4. On returning from the recursion, the individual BRDF components are attenuated by the Fresnel transmission coefficients $R_{t_{21}}$ for the level above them, and added to the total BRDF.

3.2.4 Sampling PDFs

Common to gathering expansion solvers of the rendering equation is the fact that at each recursion level, they attempt to evaluate the illumination integral through stochastic numerical integration. One of the standard techniques to accelerate the convergence of such a stochastic integration is to perform *importance sampling*, which requires that the integrand be randomly sampled using a probability density function that mimics the integrand as closely as possible.

In practice, this means that for any reflectance model one not only needs formulas for its BRDF values, but also an efficient sampling PDF for the BRDF. While formulas for the BRDF are usually given in literature, a sampling PDF is often omitted, which limits the immediate applicability of some published models.

Since our BRDF consists of many different functions, it is impractical to find a probability distribution that matches the whole BRDF. Instead, it is a common method to sample the component which most influences the appearance of rough surfaces, and use this distribution function to obtain the probability that a micro-facet is oriented in a specific orientation.

Again we have to distinguish between the two tasks of section 3.2.1. For the global model, the components of the surface can be sampled individually according to their respective PDFs like it would be done in a non-layered setting. For example, a Torrance-Sparrow-Lambertian stack would sample the top layer according to the microfacet distribution, then refract the incoming ray according to Snell'S law and sample the Lambertian layer with cosine sampling and the refracted incoming ray. Which ray is followed is selected randomly, but with respect to its intensity. As a consequence, importance sampling is still possible although a surface consists of BRDFs with different shapes.

Calculating the PDF for the entire BRDF, i.e. for the local model, is a little bit more complicated. What we want is a PDF that is as proportional to our BRDF as possible, and of course fulfils the two required properties that negative probabilities do not exist, and that the sum of the probabilities of all possible outcomes over a domain D is 1.

$$p(x) \ge 0$$
$$\int_D p(x) d\mu(x) = 1$$

To achieve this we again calculate the PDFs of each component and weight the probabilities p of the individual PDFs of the compound BRDF with constant values w so that

$$p = \sum w_i p_i \tag{3.6}$$

If both BRDFs have a similar shape (e.g. two Torrance-Sparrow surfaces of equal roughness), the weighting factors are equal. Otherwise we have to choose a bigger weight for the BRDF that dominates the appearance of the whole BRDF. The respective value depends on the reflection properties of the different layers. However, the sum of these weights has to be 1.

3.3 RenderMan SL Implementation

For application areas like computer games or film production rendering, global illumination solutions are often too expensive to compute. In these areas, shader language constructs, and purely local illumination models, are often used to describe the appearance of an object. To demonstrate the applicability of the layering approach to this type of graphics, we implemented the model in a shader-based language, namely the RenderMan SL. We chose this particular language because it is a very powerful and widely available SL that is used in production rendering, but that is still easy to implement. Moreover, several free RenderMan-compilant renderers exist, like e.g. 3Delight or Pixie.

3.3.1 Basic Components

To implement the model in RSL, we need two basic components that are later combined in the presented way. For this demonstration, we use an Oren-Nayar[29] and a Torrance-Sparrow component. The first one can simulate surfaces that range from perfectly diffuse to retro-reflective, the second one is the main component of the presented layered surface model. However, it should be noted that there are no restrictions; theoretically any BRDF could be used as a layering component.



Figure 3.5: Oren-Nayar model with $\sigma = 0.0, 0.2, 0.4, 0.6, 0.8$ and 1.0. With increasing σ the Oren-Nayar component becomes more and more retro-reflective.

Oren-Nayar As discussed in section 2.1.2, this BRDF is mainly diffuse, with a noticeable retroreflective component for higher roughness values. As it can be seen in figure 3.5, the retro-reflectivity can be controlled by the parameter sigma σ , which specifies the roughness of the surface. The larger sigma, the more retro-reflective is the material. The shader we use was originally implemented by Larry Gritz.

Torrance-Sparrow In literature, several different micro-facet distributions exist to predict the distribution of micro-facets. We decided to use the normalised Blinn distribution

$$D(\boldsymbol{\omega}_h) = \frac{m+2}{2\pi} (\boldsymbol{\omega}_h \cdot n)^m \tag{3.7}$$

where m is the root mean square of the slope, i.e. a measure of the roughness. The larger the average slope m of the micro-facets is, the more the reflection is spread out.

The geometric attenuation factor is even easier to compute.

```
float geom( normal Nn; vector H; vector L; vector V )
{
   float ndoth = Nn.H;
   float ndotv = Nn.V;
   float ndotl = Nn.L;
   float vdoth = V.H;
   float masking = 2 * ndoth * ndotv/vdoth;
   float shadowing = 2 * ndoth * ndotl/vdoth;
   return min( 1, min(masking, shadowing) );
}
```

Both functions are used in the Torrance-Sparrow shader together with a function that calculates the Frensel reflection coefficient.

Both components are now in a form that is useful for layering, as described in section 3.1.2.

3.3.2 Absorption

First we need to calculate the absorption term. For this purpose we need the incoming and outgoing angle as well as the layer thickness and the absorption coefficient. Note that the absorption coefficient specifies which amount of light is *absorbed* for which wavelength, i.e. which colours vanish. This means that we have to use the inverse colour, e.g. if we want to simulate a pink lacquer, alpha has to be green. All calculations have to be performed for each colour component since the effect is wavelength dependent. The shader code for the absorption can be found here:

```
extern vector I;
normal Nf, Nn;
float idotn;
Nn = normalize(N);
Nf = faceforward(Nn, I);
float r = 0, b = 0, g = 0;
vector Ln = normalize(L);
float ldotn = Ln.Nf;
float vdotn = V.Nf;
vdotn = clamp(vdotn, 0.001, 1.0);
ldotn = clamp(ldotn, 0.001, 1.0);
r = exp( -alpha[0] * d * (1.0/ vdotn + 1.0/ldotn) );
g = exp( -alpha[1] * d * (1.0/ vdotn + 1.0/ldotn) );
b = exp( -alpha[2] * d * (1.0/ vdotn + 1.0/ldotn) );
return (r,g,b);
```

Figure 3.6 shows the shader used with different parameters for the layer thickness.



Figure 3.6: A silver sphere that is coated by a blue varnish with increasing thickness. As in figure 3.4, a change in hue and satturation can be seen.

3.3.3 Internal Reflection

Calculating the internal reflection is straightforward. We only need the geometric attenuation factor and the Fresnel transmission coefficient.

Figure 3.7 shows two spheres rendered with and without taking the effect into account. The spheres in figure 3.7a and 3.7b have a much lower index of refraction than the spheres in figure 3.7d and 3.7e. What can be seen is that the effect is almost negligible for surfaces with a low indices of refraction since t takes values on between 0.95 and 1. However, for a high IOR, the effect becomes a little bit more prominent, albeit still in a rather subtle fashion. Figure 3.7c and 3.7f are visualisations of the respective values.



Figure 3.7: A silver coated with a blue lacquer with (left) and without (right) taking internal reflection into account. The white circles are visualisations of t.

3.3.4 Final Shader

The final shader combines all components. Care should be taken to not confuse the geometry normal with the half vector.

Figure 3.8 shows the individual components of the shader. Note that we added an ambient term since we rendered the scene without global illumination.



Figure 3.8: The final shader and its components. Note that we use an ambient term and an environment map to fake global illumination. These components are actually not part of the model, but greatly increase the visual appeal.

3.4 Simplifications for Real-Time Rendering

The model which was presented here can be implemented in a way that it can be rendered in realtime since recombining the layers does not introduce much overhead, assuming that some form of reflectance model has to be evaluated anyway.

3.4.1 Precomputations and Approximations

The only component that really allows precomputations are the Fresnel terms. For RGB, the most trivial solution is a function of 7 variables (1 angle, 3 real IOR components, 3 imaginary IOR components), which is unfortunately not possible. A better way would be to construct a 3D texture for each color channel, ending up with three 3D textures for RGB. Although this is better, for a reasonable resolution (128 for the angle and 32 x 32 for the (r,i) IOR's components using 16 bit floats) three 1MB textures would be generated. However, this would add additional overhead since three texture reads have to be performed instead of the usual one (also, the reads from 3D textures are not very fast, because these are not cached so easily). The third possibility would be to use one 1D texture parameterised just by the angle with a fixed IOR. Such a texture must generated for each material in the scene, but even by using a fine resolution of 1024 pixels for the angle, the textures would be only 6kB per texture/material (assuming RGB texture). With for e.g. 100 materials, the size of the texture atlas would be 800kB and just one texture read would be needed.

Moreover, the Fresnel coefficients can be approximated [23]. Comparisons with the full terms showed that the results are almost identical with the full Fresnel coefficients, even for metals. Another option is to use Blinn-Phong instead of Torrance-Sparrow as an approximation (and this is probably just what RT-graphics people would do, or use some other low-cost model, like e.g. the Strauss model).

3.4.2 Performance

The cost of the model is roughly comprised of twice the evaluation of any local reflectance model, the Lambert-Beer absorption term and the internal reflection term. So the overhead over a single layer model is one additional local reflection model evaluation, absorption term and internal reflection term. The final cost of course also depends on the used local reflection model for both layers - in the measurements and in the screen-shots Torrance-Sparrow was used, but as already said, one could also use the Blinn-Phong model or anything else, which would reduce the computational costs considerably.

For the real-time demo, only one layer was implemented, since three and more are not really affordable in terms of performance, and are also not necessary in most cases, e.g. for games. Screentshots can be seen in figure 3.9.

In concrete numbers, the cost of layered model fragment shader evaluation is 479 cycles, yielding a pixel throughput of 266 MPix/s, which means about 140 FPS at 1600 x 1200 resolution, assuming that the layered object covers the entire screen and nothing else is computed. The cost of Blinn-Phong is 83 cycles per fragment, producing a throughput of 1991 MPix/s, making about 1030 FPS on the mentioned resolution. This may seem like a big difference, but these are the results of the unoptimised layered model without using any precomputations.

With precomputations, the layered model's performance using the Fresnel term approximation is (using the same settings as before) 268 cycles per fragment, producing throughput of 614 MPix/s (2.3 times faster than the naive implementation), and the performance using precomputed Fresnel term is 195 cycles per fragment, producing 885 MPix/s (3.3 times faster than the naive implementation). So the cost is twice the cost of the original one-layered surface if the Fresnel coefficients are precomputed; however, remember that we are still speaking about using the Torrance-Sparrow model for the layers, so further performance increases are theoretically possible.

The performance was measured in FXComposer, using settings for the GeForce 8800 GTX, a mid-tohigh-end GPU on the market. The measurements are synthetic, so in real life the performance of the layered model would probably be even better. The number of texture reads is more or less fixed and these can cause the pipeline to stall, which lowers the preformance.

3.5 Overall Performance in a Rendering System

Since our model consists of a combination of various BRDF models which all have to be evaluated, its computation time mainly depends on the number of layers, and on the BRDFs of the individual surfaces. However, in the typical high-quality global illumination setting for which it is primarily intended, the evaluation of the reflectance models is an important, but not crucial aspect for the overall



Figure 3.9: Screenshots of the real-time implementation. The tori show (3.9a) a smooth clear layer on a diffuse base, (3.9b) a smooth tinted layer on a diffuse base, (3.9c) a torus with varying varnish thickness on a diffuse base layer and (3.9d) a torus with matte varnish and a bump-mapped surface.

performance. By and large, the inclusion of our model did not significantly increase the execution times of the plain and bi-directional path tracing computations we used to generate the result images.

There was particularly little relative increase, if run-time comparisons were made against a version of the same scene in which the layered surfaces were replaced by other, non-layered surfaces of similar specularity and albedo. This kept the amount of specular reflections and caustics – and hence the level of difficulty for the rendering algorithm – within comparable limits, although the scene appearance of course was different.

This fact was also confirmed in the real-time setting, where the overall overhead is not more than a few percent at most (depending of course on the application) if the layered model is not used on everything, just on the objects in the user's attention.

Chapter 4

Classifying Materials - Using Layered BRDFs to Describe Object Appearance

In this section we will describe how to use layered surfaces to simulate a variety of different surfaces. We will show that all these different surfaces can be simulated with the layered model presented in chapter 3, with only the layering and the parameters of the individual interfaces and varnish layers being changed. To simplify the modelling process, this section will contain a classification for layered surfaces according to the reflection properties of the individual layers.

4.1 Material Appearance and Visual Clues

Before we can start to model or put materials into certain groups, we first have to answer the question "What causes the appearance of a material"? Different materials like e.g. sand, velvet or metals have a distinct visual appearance. We are immediately able to distinguish between these materials alone from looking at them, and we immediately know how these materials would feel, if they are e.g. soft, wet or rough. This is possible because the human brain is possible to interpret the features of an object and concludes from past experiences what we have to expect since it assumes that objects with a similar appearance have similar properties.

These different visual clues are e.g. the blurriness of a highlight which informs us about the surface roughness or its intensity which provides information about the reflectivity of an object. Although the brain does not explicitly try to analyse the physical properties of a material, these clues are strongly related to the surface reflectance properties of an object, which brings us back to our question. Materials look different because they reflect light in a different way and the amount of light that is reflected in a certain direction causes the visual differences between these materials.

4.1.1 Metallic vs. Non-Metallic Materials

The different distribution and intensity of reflected energy is what leads to the different appearance between metallic and non-metallic materials. Event if they are rough, metallic materials reflect the environment much better than non-metallic materials, because metallic materials have a much higher



Figure 4.1: Metallic, coated metallic and non-metallic spheres. The respective spheres in 4.1b and 4.1c have the same colour and roughness, but still look different.

overall reflectivity. When considering the average reflectance over all incident angles, non-metallic surfaces reflect only 4 to 10 percent of the total incoming light. In contrast to that, metallic surfaces reflect, on average, about 80 to 90 percent of the total incoming energy.

Another difference between metallic and non-metallic materials is that for non-metallic materials the reflection intensity is not dependent on the wavelength, i.e. the index of refraction changes only slightly with the wavelength, and no energy is absorbed during reflection. This means that the reflected light has the same colour as the incoming light. However, the situation is quite different for metals; here the index of refraction strongly depends on the wavelength and often certain wavelengths are absorbed; consequently metallic materials change the colour of the incoming light.

Figure 4.1 illustrates this concept. Figure 4.1a shows three metallic spheres, made of gold, silver and copper. The highlights of the gold and copper spheres are coloured, because the gold sphere absorbs the blue part of the light, while the copper sphere reflects only reddish light and absorbs the rest. Note that this is not the case for the silver sphere, since silver, and white metals in general, reflect all light more or less equally, independent of its wavelength.

Figure 4.1b and 4.1c further demonstrate the difference between metallic and non-metallic materials. To match the appearance, the metallic spheres in figure 4.1b have been coated with a coloured lacquer so that they have the same colour as the non-metallic spheres in 4.1c, and the respective spheres have the same degree of roughness. Still, both images look quite different; although the non-metallic spheres do reflect the environment, the reflection is only weak, and the colours of the spheres are much more prominent. In contrast to that, reflections of the other objects can be clearly seen in the metallic spheres, especially in the purple sphere.

4.1.2 Surface Roughness

The surface highlight is the most important visual clue to estimate the surface roughness. One of the fundamental physical principles, the law of reflection, is that light is always reflected in the mirror direction, i.e. that the incoming angle equals the outgoing angle. This is, however, on a macroscopic level only true for very smooth materials. With increasing surface roughness from e.g. scratches or imperfections, light is also reflected not only in the ideal direction, but also offset from the ideal spec-



Figure 4.2: Several non-metallic (top) and metallic (bottom) real world materials with different degree of roughness. Note the broadening and vanishing of the highlight.

ular angle due to microscopic surface variations. This results in a broader, but less intense highlight, since the incoming energy is reflected in many different directions. The reflection becomes more and more diffuse until the same amount of energy is reflected in every direction over the hemisphere.

As it can be seen in figure 4.2, both metallic and non-metallic surfaces exist in various degrees of surface roughness. However, with increasing roughness the highlight is not focused anymore and starts to vanish.

4.2 Classifying Materials According to the Reflection

To simplify the modeling process, we will now propose a decision tree that allows us to put materials into groups. For our purpose, we will classify materials according to their reflectance behaviour instead of using every day language like e.g. gloss or luster. The full decision tree can be seen in figure 4.3.

The first decision we have to make is if we want to simulate a metallic or a non-metallic surface. If the environment is visible, i.e. the object is highly reflective, or if the highlight is coloured, the material we want to simulate is, with a very high probability, metallic.

For non-metallic materials, the next decision is pretty easy. If the material we want to simulate has a clear reflection of the lightsource, i.e. the highlight is sharp, it falls into the group of "smooth surfaces on top of a diffuse surface", it is a surface that specularly reflects light from the top, while the remaining part of the energy enters the material, is partly absorbed there and re-emerges in a diffuse way. We call this group "glossy paint" after its most obvious representative. Materials that exhibit such a behaviour are e.g. smooth plastics, smooth ceramics of lacquer. Otherwise, if the lightsource



Figure 4.3: Illustration of the decision tree. Green arrows indicate positive decisions, red negative. We can build six basic groups of materials on the basis of their reflection properties.

is blurred or not even visible, the material falls into the "frosted paint" group -a rough layer over diffuse surfaces. Such a materials are e.g. rough plastic, rubber, paper or wall paint.

Metals are a little bit more complicated to classify. Metallic surfaces that clearly reflect the environment belong to the group of "smooth layer on top of a smooth surface" which is called "metallic foils"; they are smooth metallic materials that are coated with a smooth transparent varnish layer. Classical examples are e.g. christmas ornaments that are made of glass and coated on their interior with a silver nitrate solution.

On the other side, if the environment is somewhat visible, but blurred, at least one rough component is involved. Surfaces that still have a sharp highlight, or a clearly visible, but weak reflection of the environment, are put into the group of "smooth layer on top of a rough surface" which we call "metallic paint". They consist of a rough metallic base material that is coated with a smooth varnish. A typical representative that falls into this group is metallic car paint, but also, as we will later see in section 5.2.1, all sorts of metallic inclusions belong to this group.

Finally, materials that are metallic but have only a weak and blurry highlight – or none at all – must have a rough top surface. The only question is now if this rough top surface coats a smooth or a rough metallic base material. The only indicator we have for this is the width of the highlight. Surfaces that fall into the group of "rough layer on top of a smooth surface", called the "frosted metal"-group, tend to have a smaller highlight than the materials that belong to the "patina"-group, the group of "rough layer on top of a rough surface". Also, completely rough surfaces should have a smaller and a broader highlight.

The final classification can be seen in figure 4.4, although it should be noted that this is probably not an exhaustive list of what can be achieved by this technique. Note that there are 8 instead of 6 types. One represents the further split into tinted and non-tinted varnish that can be performed for all types, and one stands for the group of materials that consists of more than two layers, i.e. multi-layered materials.



Figure 4.4: Examples of various surface types that can be generated by using our layered model in different configurations. In order to properly distinguish the various cases the icons do not exhibit the simplifying assumption shown in figure 3.1 which we use for all our actual BRDF computations. The micro-facets are much smaller than the layer thickness in this drawing.

Chapter 5

Modelling with Layered Surfaces

After discussing the general capabilities of layered surfaces, and after presenting a methodology for using such materials, we give practical examples of how certain materials can be modelled by this approach. The fact that the same BRDF layering code is used for all of them underlines that a large variety of different surfaces can be simulated using a single module, with only the layering and the parameters of the individual interfaces and varnish layers being changed. In the interest of reducing the complexity of this demonstration section, we only used the original Torrance-Sparrow model [33] and the Oren-Nayar model [29] as basic components f_{r_n} for these layered BRDFs. Additional component BRDFs would enhance the capabilities of this approach even further.

5.1 Traditional Materials

5.1.1 Glossy Paint – Ceramics and Acrylic Lacquer

A typical application for a diffuse surface covered with a smooth layer is the simulation of glossy paint, opaque ceramic glazing, as well as plastics; examples can be seen in figure 5.7.

To model these materials, a perfectly diffuse surface is used as base material. To limit the number of different BRDF components used in this section, we decided to use an Oren-Nayar surface with roughness of 0.0° instead of the more obvious Lambert surface; this caused a slight performance hit. If performance were a critical consideration, one would of course use a Lambert surface instead. The advantage of using an Oren-Nayar surface is that we could also simulate backscattering, which, however, is comparable rare for glossy objects.

The upper layer is a smooth Torrance Sparrow layer; the roughness is normally a value between 0.01° for very smooth surfaces and 3° or 4° for surfaces that are smooth, but have little imperfections like e.g. very small scratches which gives them a slightly rough appearance. The difference can be seen in figure 5.70, 5.7k and 5.7f. The blue sphere looks much smoother than the red one although the index of refraction is 1.7 for both spheres and both have a perfectly diffuse base. For the blue sphere we used a roughness of 0.01° while the red one has a roughness of 1° .

The clear varnish of such objects has normally an index of refraction between 1.35 for acrylic lacquers and 1.7 for ceramics. Plastics have a refractive index of about 1.5; ceramic glazing has a higher index

of refraction than the binder typically used for enamels – as a consequence, ceramic glazing has brighter specular reflections than paintwork. Choosing a higher index of refraction is often unrealistic for dielectric materials.

In addition, the lacquer layer of the sphere in figure 5.7k is tinted. For transparent lacquer layers, the thickness of the layer can be almost neglected since hardly any absorption occurs. In surfaces of this type, the colour usually increases in saturation towards grazing angles, where the colour of the sphere in figure 5.70 largely stays the same everywhere on the object regardless of curvature. Since the varnish layer has to be comparatively thin, is has to be strongly absorbing.

5.1.2 Frosted Paint – Sugar, Spraypaint and Latex

Garnet Red is a spray paint with a low gloss component. The material has bright specular highlights at grazing angles and at the specular angle. Viewed from certain angles the materials are almost ideally diffuse, though.

Figure 5.7a shows a red sphere that has been coated with a gloss-reducing finish; this is one of the examples from the Cornell BRDF database [1]. We combined a red Lambertian surface (Munsell BoC 7.5R 2/8) with a rather rough Torrance-Sparrow varnish. The average micro-facet slope is 12° and the refractive index of the varnish is 1.6. Here, the varnish layer is rather clear and thin – the thickness is 0.5.

The sugar sphere in figure 5.7h is covered by a layered surface that consists of a neutrally coloured Oren-Nayar base with a sigma of 0.26, covered by a very rough ($m = 26^{\circ}$) and completely transparent Torrence-Sparrow layer. This combination leads to a surface with very weak highlights, like a real sphere made of compacted sugar would exhibit.

The green latex sphere in figure 5.7l uses similar values, except that the Oren-Nayar sigma is now 0.34, that the Oren-Nayar surface is coloured green (NCS S4040-G50Y), and that the IOR of the transparent, rough top layer is even lower than on the sugar sphere (n = 1.3, the sugar sphere coating has n = 1.35).

5.1.3 Metallic Foil – Brass and Christmas Orbs

A very typical appearance results if a specular, colourless metal is covered by a smooth, tinted layer of varnish. The iconic object for such surfaces are christmas orbs, but polished brass surfaces such as figure 5.7g also belong to this category.

The brass material is a good example that the layering approach can also be used to successfully jerry-rig appearances, instead of always using it to exactly simulate the behaviour of real materials. Certain finishes of brass have frayed highlights that can not be simulated with a single Torrence-Sparrow surface. In this example, a fairly specular ($m = 4^\circ$), coloured metal surface is covered by an almost perfectly smooth ($m = 1^\circ$), tinted varnish layer with an IOR of 2.0. This gives a very good approximation of the desired appearance. Most of the colour of the finished surface comes from the varnish layer.

5.1.4 Metallic Paint – Automotive Paint

One very recognizable type of surface are the metallic paints used by the automotive industry. They consist of metal flakes suspended a binder, and which are covered by a transparent cover varnish. Optionally, the flakes can be suspended in a tinted binder, or the flakes themselves can be coloured. If one relinquishes the simulation of individual flakes (which, at least for settings that do not involve extremely close-up viewing, is usually permissible), this sort of paint lends itself to a layered modelling approach exceedingly well.

One uses a rough ($m = 12^{\circ}$), neutrally coloured metallic Torrance-Sparrow model as basis, and covers this with a tinted (Munsell 2.5B 5/10), smooth ($m = 1^{\circ}$) layer of varnish. Since this is a fairly good approximation of real metallic paint in terms of structure, it is not surprising that the sphere in figure 5.7b shows all the expected effects, such as darkening towards the edges, and a deep colouring that reveals metallic characteristics only for steep viewing angles.

The parameters for the paint shown in figure 5.7b correspond fairly closely to what one would expect to find in a real metallic paint. For the second car paint sample, the gold paint shown in figure 5.7d, and which is based on a MERL sample, this is quite as true. For the top layer, the IOR needed for a close correspondence with the measurements is 1.85, which is too high for normal varnish. However, such a high IOR could be seen as simulating a varnish layer with increased reflectivity, e.g. a layer with additional, small metal flakes in it.

5.1.5 Frosted Metal – Metallic Gold and Silver, Blue Metallic Paint

Gold metallic paint has a smooth metal shine, with fine aluminum powder and pigments to obtain the shine. This material has tiny bright spots caused by mirror like flakes. These spots are visible on grazing angles. The dominant overall reflectance property is rough specular (glossy). Ngan et al. [NDM05] show that there is no analytical reflectance model that can reproduce the material with an error under 0.02, which they demonstrated in there experimental results which is in fact rather good, but not very surprising since the Cook-Torrance model was developed to simulate this material type.

Figure 5.7i, blue metallic paint, is such a surface: a fairly smooth Torrance-Sparrow aluminium base $(m = 1^{\circ})$ covered by a rough $(m = 17^{\circ})$, tinted (NCS S2070-R80B) Torrence-Sparrow layer.

5.1.6 Patina – Silver Paint

Rough-over-rough material are arguably the rarest of the surface types discussed here. The silver paint (as the MERL sample in figure 5.7e is known by name) is such a surface: a colourless metal Torrance-Sparrow base ($m = 10^{\circ}$) covered by a clear ($m = 7^{\circ}$) Torrence-Sparrow layer of moderate IOR 1.7.

5.1.7 Multi-Layer – House Paint, Fabric

Some examples of the Cornell and MERL BRDF measurement databases do not lend themselves easily to simulation by a single layer set-up. In some cases, such as the light blue house paint, two layers atop a base substrate are needed to match the behaviour of the targeted material. In particular, matching the highlight shapes of a given object can require additional layers if the material in question exhibits features such as a strong fan-out of highlights.

In the case of the light blue house paint example shown in figure 5.7c, two transparent layers of equal IOR 1.3 are stacked atop a coloured (Munsell 5PB 4/10) Oren-Nayar base with sigma 0.34. The top layer has $m = 30^{\circ}$, while the lower has $m = 7^{\circ}$. The combination of the rough top layer over the intermediate, not so rough transparent layer creates the highlight shape one sees in the result.

5.2 Special Materials

5.2.1 Heterogeneous Surfaces

Many natural materials are coloured more or less irregularly, an effect that cannot be adequately described by a single layered surface model alone. In this section, we present two examples of how one can efficiently model more complex structures by combining layered surfaces with procedural textures [38]. In this context we use layered surfaces with different parameters, and use them as texture input to produce heterogeneous structures.



Figure 5.1: (a) Samples of various almost opaque, aventurescent minerals, Lavender Lepidolite (left stone), Raspberry Aventurine (top right) and a Green Aventurine specimen (middle). (b) Example renderings of various forms of Sunstone. The metallic glitter and orientation-dependent sheen have made this particular mineral a sought after gemstone. Note that only the specimen on the top right is opaque; the other two are translucent, volumetric objects with a glitter map.

As case study we describe the modeling of aventurescence in gemstones (see figure 5.1), a glittering effect that is caused by small, crystalline metallic inclusions in a parent mineral with a highly reflective surface of up to 1mm in size, but similar approach can be used to e.g. model flakes in a metallic lacquer. The obvious benefit of using layered surfaces in this context is that even though the shader-based texture patterns that govern the appearance of these objects are more or less ad-hoc creations, the overall appearance of these stones is still highly physically plausible, since the individual texture components are realistic.

As discussed in detail in [38], the core of achieving results such as those shown in figure 5.1 is the use of procedural texture layers, each of which contains surfaces that are described by layered BRDFs.

Figure 5.2 shows an overview of how various layered BRDFs can be integrated into a system of specialised Voronoi textures. In the following two sections, we briefly discuss the details of two of the gemstones shown in figure 5.1, namely the Layender Lepidolite and the Sunstone specimens.



Figure 5.2: Schematic illustration of the spatial arrangement of the Voronoi cell texture layers. The cells of textures that represent individual particle facets located on a deeper level are smaller (component (1) - the Surface Map), and the texture layer nearer to the top (layer (2) - the Glitter Map) includes fewer and larger cells that are assumed to stand for entire particles. Most of the cells in (2) are completely transparent, with only those responsible for the glitter remaining. Component (4) - the crystal facet bump map - provides sparkle facets for exactly those cells in layer (2) that are not transparent. *Note that the varying colour of the bump map facets in this drawing is not due to different types of surface being used, but due to shading effects of the now non-planar surface!* The combined result (5) mostly consists of polygonal facets of varying colour intensity, with some of them exhibiting additional sparkling effects due to their modified surface normal.

Lavender Lepidolite

The overall modelling of the Lavender Lepidolite in figure 5.1 can be inferred from figures 5.2 (for the basic idea behind the pattern) and 5.3 (for the variations in the pattern). The surface map of the dark tiles has 9 entries, but in contrast to the Aventurine glass, most of the surfaces have a Lambertian base of different colour (we used 6 different colours) and a smooth layer with index of refraction of 1.56. Only one entry has a metallic (in this case a silver) base with an average micro-facet slope of 1 degree. The layer is smooth and pinkish (similar to the colour of the tiles). The second surface map of the bright tiles is build similar, except that the colour of the Lambertian bases are brighter. Here we used only 5 different colours. Dark and bright tiles have the same size and are put in another surface map with bigger tiles.

The third surface map for the marble inclusions has only 3 surface entries which are even brighter than the bright tiles, but to disrupt the otherwise continuous line, entries from the bright and from the dark tiles are also used. Please note that the tiles of the marble surface map are much smaller than the dark or bright tiles.

The glitter map is built similar to that of the Aventurine glass with 4 entries. They have a silver base, but otherwise everything is the same. Only very few entries of the glitter map are filled.

Sunstone

Unlike the Lavender Lepidolite, the opaque Sunstone on the top right of figure 5.1 has no surface map, because it is so opaque that hardly any variation in colour can be seen. Therefore we used the same



Figure 5.3: (a) A close-up of the Lepidolite from figure 5.1, which reveals that – like in a real stone of this type – the overall appearance of the stone is caused by a complex pattern of polygonal particle facets, which are further influenced by colour changes driven by low-frequency, marble-like patterns within the material. To produce this effect, the simple arrangement of layered surfaces used for Aventurine Glass in [38] is replaced by a more complex combination of various surface maps outlined in (b). Each of these maps in turn consists of several layered surface elements, and a procedural marble texture function is used to switch between the individual maps.

Torrance-Sparrow surface with a Lambertian base and smooth varnish to fill all transparent tiles of the glitter map. Note that we used the reflection spectrum of a real Sunstone for the colour.

The glitter map has again four entries. The base material of the glitter is copper, the average microfacet slope is 3 degrees. The varnish has the colour of the Lambertian base and has an index of refraction of 1.58. The average micro-facet slope is 4 degrees. The relation between glitter and nonglitter tiles is as is the case of Aventurine glass. The glitter map has 3 entries, the base material is copper and the average micro-facet slope is 5 degrees. The tiles are tinted with the same colour the stone exhibits. The pink stone exhibits much more glitter, because many cells of the glitter map are filled. In contrast to that, the patterned Sunstone has less glitter, but the scattering function is no longer homogeneous. To produce the desired effect we combined three different Perlin functions with different parameters. The index of refraction is again 1.58 for both stones.

5.2.2 Fluorescent Surfaces

A less obvious application of layered surfaces is to use them to simulate reflection from fluorescent surfaces [39]. Unlike normal surfaces that just reflect more or less incident light depending on the viewing and illumination geometry, fluorescent surfaces have the unusual property that they can also change the wavelength of the reflected light through processes at the molecular level of the surface pigments. A remarkable property of fluorescent surfaces is that the reflection colour is directionally dependent. This effect can be seen in figure 5.5, where a green laser beam with different angle of incidence is reflected from an orange fluorescent day-glo cardboard sample, and can be easily replicated by anyone with access to a green laser pointer and a piece of day-glo cardboard. Note that while the bright dot where the laser hits the surface – and also parts of the reflection pattern – are a

bright orange caused by the frequency shift typical for this kind of fluorescent colorant, other parts of the reflection – essentially the specular component – are still the native green of the incident laser beam.



Figure 5.4: Concept of the rough varnish microfacet model used for the diffuse fluorescent BRDF needed to model the surface behaviour seen in figure 5.5. A rough layer of transparent Fresnel facets is assumed to cover a thin layer of transparent varnish on top of a fluorescent Lambertian base layer. Fluorescence only takes place for those rays that come in contact with the base layer.

As discussed in [39], this sort of bi-coloured reflection behaviour is pretty much intractable with normal, single-model BRDFs and normal combinations of BRDFs (such as the classical "Phong plus Lambert" model). However, a rough varnish type layering (the logic behind which can be seen in figure 5.4) can capture this sort of reflectance pattern quite well, and with minimal modelling effort. It is important to note that this usefulness of layered models for fluorescent BRDFs does not stop here: specular fluorescent paint, such as the paint found on ambulances and police cars (which also exhibits the bi-colored patterns, only in a much more focused way), is of course also easily modelled as a "glossy paint" surface, with the diffuse undercoat being fluorescent.



Figure 5.5: Reflection Patterns on a real diffuse day-glo orange cardboard sample: **a**) incident angle of less than 5 degrees; **b**) incident angle of about 20 degrees; **c**) incident angle of about 60 degrees. Note that while the bright dot where the laser hits the surface – and also parts of the reflection pattern – are a bright orange caused by the frequency shift typical for this kind of fluorescent colorant, other parts of the reflection – essentially the specular component – are still the native green of the incident laser beam. **a**) is a nice visualisation of the well-known effect that even seemingly very diffuse surfaces get more specular at grazing angles. In case **c**), almost no green specular component is evident; practically the entire reflection is orange and has been modified by the fluorescence effects in the colorant molecules.

Acknowledgements

We would like to thank Oskar Elek (oskee (at) centrum.cz) for doing the real-time implementation of the layered surface model and measuring its performance.



Figure 5.6: A rendering of a scene similar to the setting of real scene shown in figure 5.5 with similar incident angles of the laser beam: **a**) angle of incidence of 5 degrees; **b**) angle of incidence of 20 degrees; **c**) angle of incidence of 60 degrees. As described in the text, a layered surface model was used for the surface from which the laser beam is being reflected. The parameter m of the microfacet distribution was set to a value which corresponds to a mean slope angle of 17 degrees. Note that some of the visual differences to the images in figure 5.5 are due to the differences between the tone reproduction process used for the synthetic images, and the characteristics of the digital camera.



(a)





(d)





(g)







(j)







Bibliography

[1] Cornell BRDF measurement database. \newlinemeasurements/reflectance/.

http://www.graphics.cornell.edu/online/

- [2] MERL BRDF measurement database. http://www.merl.com/brdf/.
- [3] Michael Ashikhmin, Simon Premoze, and Peter Shirley. A microfacet-based brdf generator. In SIG-GRAPH, pages 65–74, 2000.
- [4] Michael Ashikhmin and Peter Shirley. An anisotropic Phong BRDF model. *Journal of Graphics Tools:* JGT, 5(2):25–32, 2000.
- [5] Mark Colbert, Sumanta Pattanaik, and Jaroslav Krivanek. Brdf-shop: Creating physically correct bidirectional reflectance distribution functions. *IEEE Comput. Graph. Appl.*, 26(1):30–36, 2006.
- [6] R. L. Cook and Kenneth E. Torrance. A reflectance model for computer graphics. ACM Trans. Graph., 1(1):7–24, 1982.
- [7] Julie Dorsey and Pat Hanrahan. Modeling and rendering of metallic patinas. In SIGGRAPH, pages 387– 396, 1996.
- [8] Sergey Ershov, Andrei Khodulev, and Konstantin Kolchin. Simulation of sparkles in metallic paints. In *Proceeding of Graphicon*, pages 121–128, Aug 1999.
- [9] Sergey Ershov, Konstantin Kolchin, and Karol Myszkowski. Rendering pearlescent appearance based on paint-composition modelling. *Comput. Graph. Forum*, 20(3), 2001.
- [10] Sergey Ershov, Roman Ďurikovič, Konstantin Kolchin, and Karol Myszkowski. Reverse engineering approach to appearance-based design of metallic and pearlescent paints. *The Visual Computer*, 20(8-9):586–600, 2004.
- [11] Xavier Granier and Wolfgang Heidrich. A simple layered rgb brdf model. *Graphical Models*, 65(4):171–184, 2003.
- [12] Chet S. Haase and Gary W. Meyer. Modeling pigmented materials for realistic image synthesis. ACM Trans. Graph., 11(4):305–335, 1992.
- [13] Pat Hanrahan and Wolfgang Krueger. Reflection from layered surfaces due to subsurface scattering. Computer Graphics, 27(Annual Conference Series):165–174, 1993.
- [14] Hideki Hirayama, Kazufumi Kaneda, Hideo Yamashita, and Yoshimi Monden. An accurate illumination model for objects coated with multilayer films. *Computers & Graphics*, 25(3):391–400, 2001.
- [15] Hideki Hirayama, Kazufumi Kaneda, Hideo Yamashita, Yoshiki Yamaji, and Yoshimi Monden. Visualization of optical phenomena caused by multilayer films with complex refractive indices. In *Pacific Conference on Computer Graphics and Applications*, pages 128–137, 1999.
- [16] Hideki Hirayama, Kazufumi Kaneda, Hideo Yamashita, Yoshiki Yamaji, and Yoshimi Monden. Visualization of optical phenomena caused by multilayer films based on wave optics. *The Visual Computer*, 17(2):106–120, 2001.

- [17] Hideki Hirayama, Yoshiki Yamaji, Kazufumi Kaneda, Hideo Yamashita, and Yoshimi Monden. Rendering iridescent colors appearing on natural objects. In *Pacific Conference on Computer Graphics and Applications*, pages 15–22, 2000.
- [18] Isabelle Icart and Didier Arquès. An illumination model for a system of isotropic substrate isotropic thin film with identical rough boundaries. In *Rendering Techniques*, pages 261–272, 1999.
- [19] Isabelle Icart and Didier Arquès. A physically-based brdf model for multilayer systems with uncorrelated rough boundaries. In *Rendering Techniques*, pages 353–364, 2000.
- [20] Csaba Kelemen and Laszlo Szirmay-Kalos. A microfacet based coupled specular-matte brdf model with importance sampling. In *Eurographics Short Presentations*, pages 25–34, 2001.
- [21] P. Kubelka and F. Munk. Ein Beitrag zur Optik der Farbanstriche. Z. tech. Physik, 12:593-601, 1931.
- [22] Eric P. Lafortune, Sing-Choong Foo, Kenneth E. Torrance, and Donald P. Greenberg. Non-linear approximation of reflectance functions. In SIGGRAPH, pages 117–126, 1997.
- [23] István Lazányi and László Szirmay-Kalos. Fresnel term approximations for metals. In WSCG (Short Papers), pages 77–80, 2005.
- [24] Andrew Lundberg, Lawrence B. Wolff, and Diego A. Socolinsky. New perspectives on geometric reflection theory from rough surfaces. In *ICCV*, pages 225–232, 2001.
- [25] Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. A data-driven reflectance model. ACM Transactions on Graphics, 22(3):759–769, July 2003.
- [26] Shree K. Nayar, Katsushi Ikeuchi, and Takeo Kanade. Surface reflection: Physical and geometrical perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.*, 13(7):611–634, 1991.
- [27] László Neumann and Attila Neumann. Photosimulation: Interreflection with arbitrary reflectance models and illuminations. *Comput. Graph. Forum*, 8(1):21–34, 1989.
- [28] Addy Ngan, Frédo Durand, and Wojciech Matusik. Experimental analysis of brdf models. In *Proceedings* of the Eurographics Symposium on Rendering, pages 117–226. Eurographics Association, 2005.
- [29] Michael Oren and Shree K. Nayar. Generalization of lambert's reflectance model. In *SIGGRAPH*, pages 239–246, 1994.
- [30] Martin Rump, Gero Müller, Ralf Sarlette, Dirk Koch, and Reinhard Klein. Photo-realistic rendering of metallic car paint from image-based measurements. *Computer Graphics Forum*, 27(2), April 2008.
- [31] Christophe Schlick. A Customizable Reflectance Model for Everyday Rendering. In *Fourth Eurographics Workshop on Rendering*, number Series EG 93 RW, pages 73–84, Paris, France, 1993.
- [32] Norihiro Tanaka, Shoji Tominaga, and Toshiyuki Kawai. Estimation of the torrance-sparrow reflection model from a single multi-band image. In *Proceedings of the International Conference on Pattern Recognition*, volume 3, pages 3600–3603, 2000.
- [33] Kenneth E. Torrance and Ephraim M. Sparrow. Theory for off-specular reflection from roughened surfaces. *J.Opt.Soc.Am*, 57(9):1105–1114, 1967.
- [34] Roman Durikovič. Explicit method of sparkling effect simulation. *Journal of Three Dimensional Images*, 16(4):96–100, Jan 2002.
- [35] Roman Ďurikovič and Tomas Ágošton. Prediction of optical properties of paints. *Central European Journal of Physics*, 5:416–427, September 2007.
- [36] Gregory J. Ward. Measuring and modeling anisotropic reflection. In SIGGRAPH, pages 265–272, 1992.
- [37] Andrea Weidlich and Alexander Wilkie. Arbitrarily layered micro-facet surfaces. In *GRAPHITE '07: Proceedings of the 5th international conference on Computer graphics and interactive techniques in Australia and Southeast Asia*, pages 171–178, 2007.

- [38] Andrea Weidlich and Alexander Wilkie. Modeling aventurescent gems with procedural textures. In *Proceedings of the Spring Conference on Computer Graphics (SCCG)*, pages 1–8. ACM, April 2008.
- [39] Alexander Wilkie, Andrea Weidlich, Caroline Larboulette, and Werner Purgathofer. A reflectance model for diffuse fluorescent surfaces. In Y. T. Lee, Siti Mariyam Hj. Shamsuddin, Diego Gutierrez, and Norhaida Mohd Suaib, editors, *GRAPHITE*, pages 321–331. ACM, 2006.

































































Oren-Nayar Shader Code (shader from Larry Gritz)
Eye =-normalize (1); theta_r = acos (Eye = M); sigma 2 = sigma = sigma, the photon = sigma, 2 = sigma = sig
A. Wedlich and A. Wilkie 37 SIGCRAPHASIA2009

















































































































