



# SIGGRAPH Asia 2011 Course Notes

## How to Write a Polarisation Ray Tracer

Alexander Wilkie

Andrea Weidlich

September 22, 2011

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Physics Background</b>	<b>4</b>
2.1	Light, a Transversal Electromagnetic Wave . . . . .	4
2.1.1	Oscillation of a Single Photon . . . . .	5
2.1.2	Micro- vs. Macroscopic Polarisation of Light . . . . .	6
2.2	Polarisation by Reflection and Transmission . . . . .	7
2.2.1	Perfectly Smooth Surfaces – The Fresnel Terms . . . . .	7
2.2.2	Reflection from Glossy Surfaces . . . . .	11
2.3	Emission Polarisation . . . . .	12
2.4	Polarisation by Scattering . . . . .	13
<b>3</b>	<b>Graphics Problems Where Polarisation Matters</b>	<b>15</b>
3.1	Computer Graphics vs. Physics . . . . .	15
3.2	Reflections from Specular Surfaces . . . . .	17
3.3	Polarised Area Light Sources . . . . .	17
3.4	Glowing Specular Surfaces . . . . .	17
3.5	Atmospheric Scattering . . . . .	19
3.6	Strain Visualisation – Stress-Induced Bi-refringence . . . . .	20
3.7	Organic Solutions . . . . .	21
<b>4</b>	<b>Mathematically Describing Polarised Light</b>	<b>22</b>
4.1	Stokes Vectors . . . . .	22
4.2	Mueller Matrices . . . . .	24
4.3	Common Mueller Matrices . . . . .	24
4.3.1	Plain Attenuation . . . . .	24
4.3.2	Depolariser . . . . .	24
4.3.3	The Ideal Linear Polarisation Filter . . . . .	25
4.3.4	The Mueller Matrix for Fresnel Reflectance . . . . .	25
4.4	Stokes Vector and Mueller Matrix Calculus . . . . .	26
4.4.1	Coordinate System Tracking – Reference Frames . . . . .	26
4.4.2	Operations on Stokes Vectors . . . . .	27

4.4.3	Light vs. Light Intensity . . . . .	28
4.4.4	Operations on Mueller Matrices . . . . .	29
4.4.5	Operations on Stokes Vectors and Mueller Matrices . . . . .	30
<b>5</b>	<b>Building and Debugging a Polarisation Ray Tracer</b>	<b>31</b>
5.1	Extending a Conventional Ray Tracer . . . . .	31
5.2	Visualisation Tools as Part of the Working Environment . . . . .	32

# Chapter 1

## Introduction

Polarisation Ray Tracing is a bit like spectral rendering: you might have heard about it, but the technical details – what exactly this is all about, how it is best done, what exactly the benefits are, where you ought to use it, and where you can safely omit it – are not that well known.

At least in principle, information about all this is available already – but it is scattered across many different sources, some of which are outside graphics literature, and consequently can be a bit hard to find at first. Also, a number of these resources, in particular publications in physics literature, were not written with graphics engineering in mind, and can prove hard to directly apply to the engineering problem of getting such a renderer to work.

This tutorial aims to provide a unified one-stop information resource on this topic, and should enable graphics engineers with a background in physically based rendering to properly assess whether the phenomenon has to be included in a given simulation, and what has to be done in order to properly allow for handling of the phenomenon in a ray-based renderer.

However, the contents of the tutorial are also potentially relevant for other areas of graphics beyond rendering proper: polarised light is currently being used in several other contexts, such as face scanning, highlight removal, or 3D projection technologies. Engineers and researchers in these other areas might benefit from the course insofar as the knowledge needed for an outright simulation of polarised light transport presented here is a superset of the information needed for their more specialised tasks, so this comprehensive overview would be a useful repository of knowledge for those other areas as well.



# Chapter 2

## Physics Background

The first module gives an introduction to what the polarisation state of light is, and how light can, under certain circumstances, become polarised in the first place.

### 2.1 Light, a Transversal Electromagnetic Wave

For the purposes of computer graphics, it is usually sufficient to describe light as an electromagnetic wave of a certain frequency – or, in the practice of physically-based rendering, a spectral mixture of frequencies<sup>1</sup>. Any such wave, or spectral mixture of waves, is assumed to travel linearly through space along the path of a discrete ray, and phenomena from the realm of wave optics are usually not considered.

This is of course a simplification of physical reality, with perhaps the most prominent omission being that the photons in a real beam of light actually oscillate in a fashion that is *transversal* relative to their direction of propagation. This property of electromagnetic waves has subtle and not so subtle repercussions on the way electromagnetic radiation interacts with surfaces and transmissive media. As we will see, the exact description of this phenomenon is not particularly complicated, but it does require more than just the notion of radiant intensity, which is the basis of practically all conventional representations of light in computer graphics.

A good introduction to this phenomenon from a physicist’s viewpoint is given by Shumaker [12] or Goldstein [5], and we refer the reader to these, and other related literature [2, 7] for an in-depth discussion that goes beyond the basic facts and the engineering aspects which we discuss in this tutorial.

---

<sup>1</sup>In this tutorial, we do not concern ourselves with the question whether one should conduct image synthesis in colour space, or with spectral representations of light. Or how a spectral mixture of frequencies is best stored and processed, if colour space rendering is deemed insufficient for the task at hand. Within Computer Graphics, this is the domain of Spectral Rendering research, and as such outside the scope of this work. We refer the reader to appropriate surveys on this topic, such as [3].

### 2.1.1 Oscillation of a Single Photon

As already mentioned, at the core of the whole issue is the fact that as an electromagnetic wave propagates through space, the tip of the wave vector oscillates in a transversal fashion. Since the propagation of light waves takes place in a three-dimensional setting, this of course potentially amounts to an oscillation in *two* transversal directions – the two directions that are orthogonal to the direction of propagation. This in turn means that the motion of this oscillation can be slightly counter-intuitive in some cases, because it does not necessarily involve the sort of *planar* oscillation one would normally associate with wave phenomena that can be described a single frequency measurement. Figure 2.1 shows two typical scenarios, only one of which corresponds to conventional notions of 2D sine waveforms propagating through space.




Figure 2.1: **Left:** a linearly polarised wave-train, which corresponds to case b) in figure 2.2. **Right:** an elliptically polarised wave-train, which in this case is a variant of scenario d) in figure 2.2 (the ellipse is axis-aligned, and not tilted, as in that example). Note that both images show a single wave-train, which conceptually corresponds to the oscillations of a single photon. Macroscopically, such behaviour is only observable if all photons in a ray of light exhibit exactly the same polarisation state. Also, note the fact that both waveforms have the same overall frequency, and therefore would appear as exactly the same colour to a human observer. Human eyes do not have a reliably useful capability to distinguish between different forms of polarisation. For Predictive Rendering purposes, the difference between them can be important, though: in some circumstances, the interaction of these two wave-trains with matter (e.g. when they are reflected by a specular surface) can be considerably different.

From a formal, mathematical viewpoint, the standard approach to understanding the phenomenon is to decompose the electric field of a wave that moves in the direction of  $z$  into two orthogonal components that are aligned with the  $x$  and the  $y$  axes, respectively. This is permissible since except at distances from the light source of a few wavelengths or less, the  $z$  component of the electric field will be negligible and the field will lie in the  $x - y$  plane. The  $x$  and  $y$  field components will be of the form

$$\begin{aligned} E_x &= V_x \cdot \cos(2\pi \cdot \nu \cdot t + \delta_x) \quad [\text{V} \cdot \text{m}^{-1}] \\ E_y &= V_y \cdot \cos(2\pi \cdot \nu \cdot t + \delta_y) \end{aligned} \tag{2.1}$$

where  $V_x$  and  $V_y$  are the amplitudes  $[\text{V} \cdot \text{m}^{-1}]$ ,  $\nu$  is the frequency  $[\text{Hz}]$ ,  $\delta_x$  and  $\delta_y$  are the phases  $[\text{rad}]$  of the electromagnetic wave-train, and  $t$  is the time  $[\text{s}]$ . The overall behaviour of the wave

vector can then be understood by considering the relationship between these  $x$  and  $y$  components of the field. Figure 2.1, and in particular figure 2.2, give examples of what this looks like: the most general case is elliptical polarisation, with circular and linear polarisation as important special cases.




Figure 2.2: Four typical scenarios of how a single wave-train of electromagnetic radiation can be polarised, and what the  $x$  and  $y$  components of the field look like in these cases. Scenario **b**) is what commonly is referred to as *linearly polarised* light, i.e. a state where the tip of the wave vector oscillates in a plane orthogonal to the direction of propagation. Figure 2.1a shows a 3D image of this. In such a state, the  $x$  and  $y$  components can be of different magnitudes, but are in phase. Vertically polarised light, as shown in **a**), is a trivial corner case of this scenario. However, the general case of polarisation is actually **d**): *elliptical* polarisation, where the  $x$  and  $y$  components are out of phase, and of differing magnitudes. In this scenario, the tip of the wave vector moves through space in a corkscrew-like motion – figure 2.1b gives a 3D view of this. Scenario **c**), so-called *circular* polarisation, is just a special case of this, where the  $x$  and  $y$  components are of equal magnitude, and exactly  $\pi/2$  out of phase.

## 2.1.2 Micro- vs. Macroscopic Polarisation of Light

For a practical understanding of this phenomenon in real environments, it is essential to always keep in mind that equation 2.1, as well as figures 2.1 and 2.2, describe the oscillations of a *single photon* – and that the behaviour of macroscopically observable light is the aggregate of huge numbers of such individual photons. A single photon is *always* polarised: its oscillation will exhibit one of the patterns shown in figure 2.2, or a variation thereof. By contrast, an entire light ray only exhibits properties of polarised light if *all* of the photons it consists of, or at least a significant percentage of them, are polarised in exactly the same fashion. This observation is highly important from an engineering perspective, since it shows us that a practical way of handling polarisation information in a rendering system has to be able to handle *partial* polarisation of a light ray – a property that does not make sense for an individual photon. But for

the huge numbers of photons that make up real light rays, such statistical concepts are needed. Light where only 10% of all photons are linearly polarised in a certain plane will only exhibit weak preferential filtering when it transits a linear polarisation filter; the 90% of its photons that are not linearly polarised will not be affected by the filter in a directionally-dependent way. They will, however, be subjected to the average filtering characteristics of the linear filter – which amounts to a more or less strong attenuation that is not directionally dependent. The data structures we use to describe light and reflectance, which we discuss in chapter 4, have to be able to take this sort of behaviour into account.

In the next sections, we give a very brief overview over some of the main mechanisms that can cause macroscopic light polarisation, i.e. similar polarisation of all, or at least a significant portion of, the many photons in a beam of light.

## 2.2 Polarisation by Reflection and Transmission

Apart from the corner case of perfectly diffuse surfaces (which act as a depolariser – see section 4.3.2), any light that interacts with a phase boundary will, in addition to being attenuated and deflected in some way, have its polarisation state altered and/or emerge from the interaction with a certain degree of polarisation. For perfectly smooth phase boundaries, the widely known Fresnel Terms give an exact solution to this problem, while for glossy surfaces, the situation is somewhat less clear. For such surfaces, various approximations have to be used instead.

### 2.2.1 Perfectly Smooth Surfaces – The Fresnel Terms

In their most extensive form, which we quote in equation 2.2, the Fresnel Terms consist of *two* pairs of equations, of which only the first is normally quoted in computer graphics literature. For the reflection geometry shown in figure 2.3, the first pair of equations determines the amounts of incident light intensity  $F_{\perp}, F_{\parallel}$  which are reflected for the  $x$  and  $y$  components of the incident wave-train. In the context of a reflection that is described by the Fresnel Terms, the  $(x, y)$  components of the wave-train are usually referred to as the *perpendicular* (denoted by  $\perp$ ) and *parallel* ( $\parallel$ ) components<sup>2</sup>. These qualifiers are to be understood in relation to the plane of incidence, which is shown in green in figure 2.3.

The first two formulas  $F_{\perp}, F_{\parallel}$  for reflected intensity are common graphics knowledge, and can be found in most computer graphics textbooks. The second pair for  $\delta_{\perp}, \delta_{\parallel}$ , which is less frequently mentioned in graphics texts [20], describes the *retardance* that the incident light is subjected to, i.e. the relative phase shift that the vertical and horizontal components of the wave-train undergo

---

<sup>2</sup>Sometimes, the subscripts for the two components are also given as  $s$  and  $p$ , respectively. These abbreviations derive from the German words *senkrecht* (perpendicular) and *parallel*. That particular custom is a carry-over from older optics literature, which was mostly in German.




Figure 2.3: Geometry of a ray-surface intersection with an optically smooth phase boundary between two substances, as described by the equation set (2.2). The plane of incidence is shown in green, and the planes in which the wave vectors of the incident and reflected rays  $I$  and  $R$  are traced is shown in red. The  $E$ -vectors for the transmitted ray  $T$  have been omitted for clarity. The  $(E_{\parallel}, E_{\perp})$  components in the two red planes shown here directly correspond to the  $(x, y)$  field components in figures 2.1 and 2.2.




Figure 2.4: **Left:** Fresnel reflectance  $F_{\parallel}$  (green),  $F_{\perp}$  (blue) and  $F_{\text{average}}$  (red, dashed), as well as the average degree of induced reflective polarisation (orange, dotted) for a typical dielectric substance with an index of refraction of 1.52. The cyan dashed line is at Brewster's angle. **Right:** Fresnel retardance  $\delta_{\parallel}$  (green) and  $\delta_{\perp}$  (blue) for this IOR, with Brewster's angle again shown as a cyan dashed line. Over the entire range of incident angles, such an interface causes linear polarisation, and only induces phase shifts between the incident wave-train components of either  $\pi$  (beneath Brewster's angle), or 0 (above).

during reflection:

$$\begin{aligned}
F_{\perp}(\theta, \eta) &= \frac{a^2 + b^2 - 2a \cos \theta + \cos^2 \theta}{a^2 + b^2 + 2a \cos \theta + \cos^2 \theta} \\
F_{\parallel}(\theta, \eta) &= \frac{a^2 + b^2 - 2a \sin \theta \tan \theta + \sin^2 \theta \tan^2 \theta}{a^2 + b^2 + 2a \sin \theta \tan \theta + \sin^2 \theta \tan^2 \theta} F_{\perp}(\theta, \eta) \\
\tan \delta_{\perp} &= \frac{2 \cos \theta}{\cos^2 \theta - a^2 - b^2} \\
\tan \delta_{\parallel} &= \frac{2b \cos \theta [(n^2 - k^2)b - 2nka]}{(n^2 + k^2)^2 \cos^2 \theta - a^2 - b^2}
\end{aligned} \tag{2.2}$$

with

$$\eta = n + ik \text{ (the complex IOR)}$$

$$2a^2 = \sqrt{(n^2 - k^2 - \sin^2 \theta)^2 + 4n^2 k^2} + n^2 - k^2 - \sin^2 \theta$$

$$2b^2 = \sqrt{(n^2 - k^2 - \sin^2 \theta)^2 + 4n^2 k^2} - n^2 + k^2 + \sin^2 \theta$$

If only the intensity of light reflection and transmission are of interest, as it is typically the case in a non-polarising renderer, the second pair of equations for  $\delta_{\perp}, \delta_{\parallel}$  is indeed irrelevant. However, for a polarising renderer, it is essential that this information is also computed – after all, phenomena like the conversion of linearly polarised incident light into elliptically polarised reflected light, which can occur on metallic surfaces, is caused by a phase shift that happens during reflection. In section 4.3.4, we show how these four quantities  $F_{\perp}, F_{\parallel}, \delta_{\perp}, \delta_{\parallel}$  can be assembled into a Mueller Matrix, which can then be used in actual light transport computations.




Figure 2.5: The same plots as figure 2.4 for the scenario where light is *leaving* a dielectric with IOR of 1.52, i.e. when it is traveling from a denser medium to a less dense one. Note that there are two critical angles in this case: Brewster's angle (cyan, dashed), and the angle at which total internal reflection (TIR) starts (grey, dashed). Up to TIR, the interface behaves like a normal reflective interface: it linearly polarises reflected light, and induces a phase shift of either  $\pi$  or 0. In the area of TIR above that angle, though, light is perfectly reflected and no linear polarisation is induced. But note that for TIR angles, the interface is capable of inducing a non-integer phase shift between the incident wave-train components!

The equations we quote here are the Fresnel equations for a dielectric–complex interface, which is arguably the most general usage case, since only one of two media at an interface can be conductive (and hence opaque). A dielectric–dielectric interface with two real–valued indices of

refraction can also be described by this formalism:  $\eta$  is real-valued in this case ( $n > 0, k = 0$ ), and the equations can be simplified accordingly.

## Transmission

The intensity of the transmission along the refracted ray  $T$  is  $1 - F$ , which of course also applies to the individual components as well:

$$\begin{aligned} T_{\perp} &= 1 - F_{\perp} \\ T_{\parallel} &= 1 - F_{\parallel} \end{aligned} \quad (2.3)$$




Figure 2.6: The same plots as figure 2.4 for a typical conductor. Notable points are that overall reflectance is much higher, that linear polarisation of reflected light is comparatively weak, and that over the entire range of reflection angles a non-integer phase shift is induced.

In figures 2.4, 2.5 and 2.6, we show plots of these functions for three typical scenarios: two that involve a dielectric interface (once from outside the dielectric, and once from the inside), and one conductor. Since there is no macroscopically relevant transmission through a conductor, we only show transmission plots for the first two cases.




Figure 2.7: Side-by-side plots of Fresnel retardance  $\delta_{\parallel}$  (green) and  $\delta_{\perp}$  (blue) together with the overall retardance (red, dashed) for the three typical cases shown in figures 2.4, 2.5 and 2.6.

It should be noted that as discussed in [15], the Fresnel terms in the form that is cited here only apply to optically isotropic media. In bi-refractive media (which are not that uncommon in nature – a large number of minerals form bi-refractive crystals) the situation for refracted rays is considerably more complicated. We refer readers to [15] for details.

### 2.2.2 Reflection from Glossy Surfaces

In the previous section, we saw that for a perfectly planar phase boundary, the polarisation aspects of the interaction between light and the phase boundary are well covered by the Fresnel equations. However, in reality, most surfaces are not perfectly smooth, but still exhibit some degree of reflective polarisation. Contrary to "normal", non-polarising BRDF models, there is further research needed in the area of semi-gloss polarisation-capable BRDF models. Currently, there are two options available to model the behaviour of such a surface: one can use polarising Fresnel facets in a microfacet-based BRDF model, or the He-Sillion-Torrance-Greenberg model. However, to date neither of these has received the sort of verification efforts with respect to their polarisation predictions that "normal" BRDFs have with respect to normal reflectance predictions, so the output of both models should, at this point, be taken with a grain of salt.




Figure 2.8: Polarising Torrance-Sparrow in a simple path tracer via brute force evaluation. For a given incoming direction  $I$ , a reflection direction  $R$  is randomly chosen according to the microfacet PDF, and the halfway vector  $H$  is assumed to be the normal vector of the microfacet responsible for this particular interaction (green). As long as only forward rays are being cast, and no evaluation of the entire BRDF is needed, such a simplistic approach works. It has the advantage to naturally yield polarisation behaviour if the polarising form of the Fresnel terms is used for each microfacet interaction. Note that even layered surfaces would work in this case, if transmission rays into additional layers are permitted.

#### Polarising Torrance-Sparrow

Since the Torrance-Sparrow is based on Fresnel microfacets, a stochastic renderer can use a comparatively straightforward technique to obtain a polarised version of the model. As each ray gets reflected in a random direction according to the microfacet PDF, the halfway vector between the incoming and outgoing direction is assumed to be the microfacet normal, and the polarising



version of the Fresnel terms can be used accordingly. As figure 2.8 shows, this approach could even be used for layered surfaces [14].

However, this has the obvious drawback that only very inefficient simple path tracing without light source sampling can be used. It is not immediately obvious how a closed form evaluation of such a polarised BRDF should best be done. Building a look-up table would work, and might well be the simplest way to obtain acceptable results in this regard. However, it is also not clear yet how realistic this entire approach is with respect to the polarisation characteristics of reflected light – verification measurements would be needed in order to properly assess this.

### The He-Sillion-Torrance-Greenberg BRDF

A BRDF that was designed from the outset to describe reflective polarisation off glossy surfaces is that of He et al. [6]. However, to the knowledge of the authors, it has not been used for this purpose so far within graphics, not even by the original authors<sup>3</sup>. As with the simplistic approach discussed in the previous section, the authors are not aware of any verification measurements that were aimed at the polarisation capabilities of this model. Given the considerably more involved nature of the model than simple Torrance-Sparrow, it is to be expected that its accuracy is higher, at least in some cases.

## 2.3 Emission Polarisation

As discussed in more detail in [18], glowing objects can be a source of fairly strongly polarised light. This effect is directly linked to the mechanisms behind reflective polarisation via Kirchhoff's law, which states that [13]:

*At thermal equilibrium, the emissivity of an object equals its absorbance.*

In this context, the absorbance is the fraction of incident light that is absorbed by the body/surface. One particular consequence of this law is that in order to describe a glowing object, one has to attenuate the blackbody radiation associated with its temperature with the absorbance  $a$  of the surface at the point of interest:

$$L_C(x, \omega) = I(T) \cdot a(x, \omega) + \int_{\Omega} L'(x', \omega') \cdot \rho(x, \omega, \omega') d\omega'$$

Since the conservation of energy dictates that the energy reflected and absorbed by an object have to be equal to the flux of incident energy, the absorbance  $a$  is easily computed as  $1 - f$  if one knows the reflectivity  $f$  of an object. For smooth surfaces with the two Fresnel reflectance components being given as  $F_{\perp}$  and  $F_{\parallel}$ , the absorbance can be computed as  $1 - F_{\perp}$  and  $1 - F_{\parallel}$ .

---

<sup>3</sup>The statement that not even the original publication actually used the polarising formulas they cite in the paper appendix is based on a personal communication from one of the authors of [6], whom the authors of this tutorial contacted about this.




Figure 2.9: Comparison plots which show the complementary behaviour of reflection and emission polarisation for dielectrics and conductors. The lower set of curves is the inversion of the reflection functions:  $1 - F$ . How much a given material polarises incident or exitant radiation is not solely determined by the difference between the two Fresnel components (cyan bars), but by the ratio of this difference to the unpolarised component of the interaction (orange bars).

This of course means that if the reflection from an object induces polarisation of some sort, any emission from the same surface will also be polarised – but in the opposite way as the reflection. And emission polarisation is indeed a characteristic feature of thermally induced emission from smooth to moderately rough surfaces [10, 9], and an example photograph can be seen in Figure 3.2, and a plot that visualises the fact that emission and reflective polarisation are indeed opposites of each other can be seen in figure 2.9. In the infrared region of the spectrum, where practically all objects emit some sort of “glow” even at room temperatures, the characteristic differences in polarisation behaviour between dielectric and conductors can even be used to identify the basic material class seen in an IR image [19], and the development of accurate models for the thermal emission polarisation of rough surfaces [1] is an active research topic.

## 2.4 Polarisation by Scattering

There are quite a number of known scattering phenomena. For the purposes of computer graphics, and in particular, polarisation ray tracing, Mie and Rayleigh scattering are those which are of main interest. Very briefly put, Mie scattering occurs on particles that are considerably larger than the wavelength of light, such as dust particles. Rayleigh, or molecular, scattering is caused by individual molecules (such as the Oxygen and Nitrogen atoms in the air), the size of which is on the same order of magnitude as the wavelength of the electromagnetic radiation that is being scattered.




Figure 2.10: Rayleigh scattering is a phenomenon which yields light that is strongly linearly polarised in the plane perpendicular to the direction of propagation. As can be seen in the scattering diagram on the right, the perpendicular polarisation component does not exhibit an angular dependency, while the parallel one goes to zero for scattering events that are 90 degrees from the direction of propagation.

Rayleigh scattering has two key properties: it is strongly wavelength dependent (it is much more likely to occur for energy-intensive short wave radiation), and it strongly polarises the scattered radiation – at least for some scattering directions. Figure 2.10 shows, on the right, a typical scattering plot for Rayleigh scattering which visualises this behaviour.

## Chapter 3

# Graphics Problems Where Polarisation Matters

Conventional engineering wisdom within the computer graphics community seems to have it that polarisation does not contribute enough to the appearance of an average scene to warrant its inclusion in rendering systems. For a large number of scenes, this assertion actually holds – but there are some exceptions. In this chapter, we give an overview of scenes where visual differences can occur, and list some object configurations where light propagation can be mis-predicted by a non-polarising renderer.

### 3.1 Computer Graphics vs. Physics

In this context, it is worth noting that as e.g. pointed out in the preface of [5], the view of a physicist on the issue of light polarisation is probably an fairly different one than that of a graphics engineer. One does not have to venture all that deeply into the realm of optics to realise that in nature, totally unpolarised light, and non-polarising interactions between light and matter (which are tacitly assumed in practically all of "normal" Computer Graphics), are actually both fairly rare occurrences. In fact, a wide variety of typical interactions between light and matter routinely cause more or less intense polarisation of light – *(partially) polarised light is literally almost everywhere!*

However, the overall effect of these frequent polarisation changes on the resulting radiant intensity (which is what our eyes ultimately perceive) is very often negligible. Since the human eye is, at best, only marginally capable of directly registering light polarisation, this means that most of these more subtle aspects of light propagation go unnoticed by human observers. Which allows computer graphics, with its focus on only computing those aspects of light propagation that are perceptually relevant, to (in most cases, justifiably) ignore these aspects of the interaction between light and matter. The only scenarios that are relevant for the purposes of graphics

engineering are those where light polarisation ends up being responsible for perceptible visual artefacts. One possible approach at systematising these is to group them into the following categories:

1. Scenarios where a single, final interaction of light with matter is so dependent on polarisation information that directly visible artefacts occur if this is not taken into account. Perhaps the most prominent example of this kind is the prediction of glare removal via polarising sunglasses – for this type of scene, the last specular reflection before the eye is reached is the main determining factor. Additional accuracy can be gained by extending the polarisation calculations beyond the first visible reflection, but the main part of any visible difference between polarising and non-polarising renderings originates from the final, directly viewed reflection.
2. Scenes where visible artefacts are caused by multiple interactions of light with matter, where each single one might not be perceptible by itself – but where the overall effect is noticeable. Examples of this are changes in the appearance of crystals (where light bounces multiple times between specular facets, before re-emerging from the stone), or the rainbow colours seen in stressed plexiglass when a polarised light source is viewed through it.
3. Special applications of rendering technology to realms with needs that are different from normal human vision. An example would be renderings used by experimental biologists to assess the visual cues that are present in a given scene for insects with polarisation vision.

These categories are partially overlapping (especially the first two), but could serve as a starting point into this area of study. However, there is also another approach that one can take – and that is to directly list the circumstances and effects that can give rise to visual artefacts. In the following sections, we do just that: we list a number of scenarios where visible differences that fall into the first two categories can occur.




Figure 3.1: A test image with a horizontal (left) and vertical (right) linear polariser applied before tone mapping. Compare the main effects of the polariser (i.e. the suppression and enhancement of certain specular highlights on the dielectrics) with the coloration in the degree of polarisation visualisation in the centre of figure 5.2.

## 3.2 Reflections from Specular Surfaces

Perhaps the most widely known area where "polarisation matters" is that of scenes which involve multiple light bounces via specular surfaces. Such scenes are susceptible to exhibit visible differences between polarising and non-polarising implementations, since reflection from specular surfaces can both give rise to significant polarisation, and alter an existing polarisation state. This is a direct consequence of the split between the parallel and orthogonal components of the Fresnel terms, and figure 3.1 shows a test scene with specular surfaces that exhibit more or less strongly polarised reflections. In reality, the most practical way to detect linearly polarised light is to view the scene through a polarisation filter, rotate the filter, and observe the changes in reflection intensity. Figure 3.1 shows that much the same thing can be done with polarised renderings, although, as very briefly discussed in section 5.2, dedicated visualisations are much better suited for diagnostic and debugging purposes.

The fact that such reflections are more or less strongly polarised is probably the most widely known polarisation-related phenomenon. It is directly useful in a number of settings, such as photography, where linear polarisation filters are routinely used to enhance or reduce specular reflections, polarising sunglasses (which are designed to reduce glare by eliminating, or at least reducing, bright specular highlights) to sophisticated engineering solutions such as the face scanning set-up presented in [4].

## 3.3 Polarised Area Light Sources

Currently, the most widespread sources of strongly polarised artificial light are flat-screen displays of all sorts. By their design, many such devices emit almost perfectly linearly polarised light as a by-product of image generation. Since the human eye is only marginally capable of directly perceiving the polarisation state of light, this is not relevant for the actual appearance of such devices. However, the specular reflection of such screens can be noticeably different, depending on the angle the reflection of the screen is being viewed under. Also, this property makes a flat-screen an ideal background to see stress visualisations with – see section 3.6.

## 3.4 Glowing Specular Surfaces

Figure 3.3 shows the overall degree of polarisation for two materials – black quartz and silver – with and without ambient illumination. It is worth noting that the presence of ambient illumination of roughly similar strength to the glow of the objects (as seen in Figure 3.3b) cancels out much of the intrinsic emission polarisation, and that the emission is only clearly polarised if the ambient illumination is switched off (as in Figure 3.3d). Also note that one can clearly see something that is mentioned in polarimetry literature, and that is evident from the plots shown

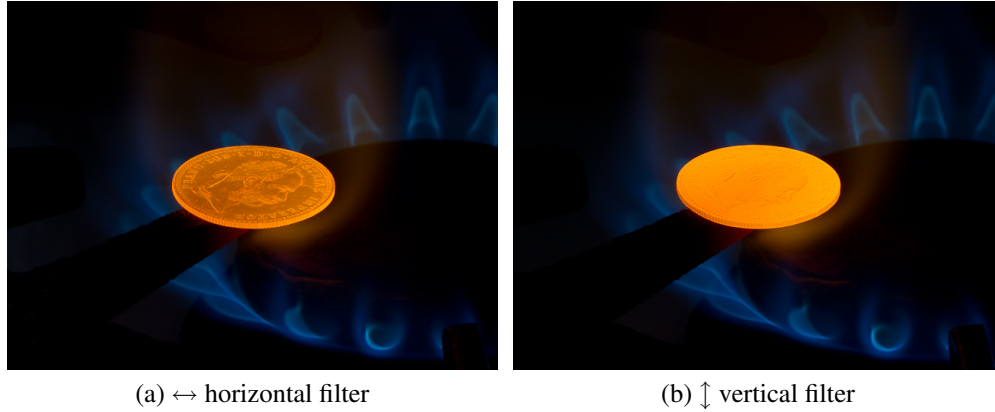


Figure 3.2: Photograph of a polished, glowing gold coin taken through a linear polarisation filter. **(a)** a horizontally oriented filter cancels out part of the emission, while **(b)** a vertically oriented one enhances it. This is exactly the opposite effect that a linear polarisation filter has on reflected light, and consistent with the plots shown in figure 2.9.

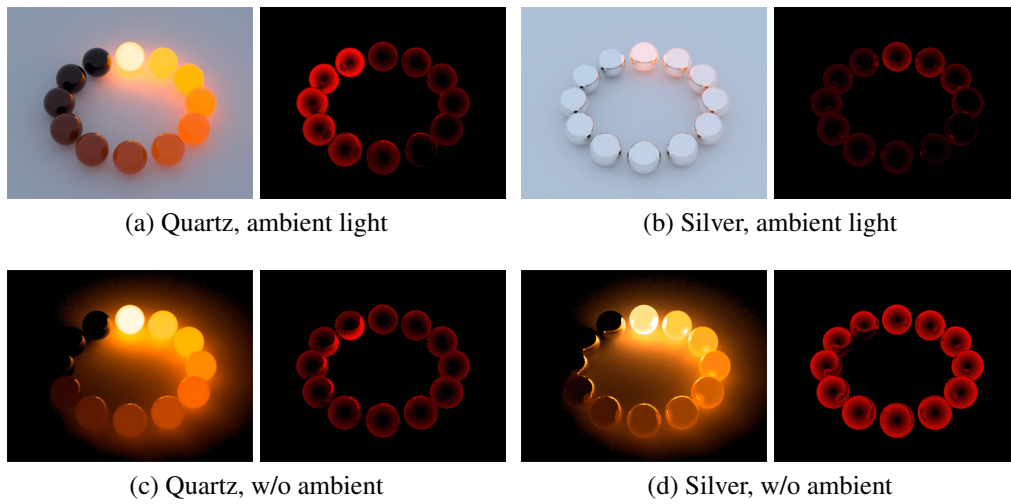


Figure 3.3: Visible light images (left), and polarisation visualisations according to [17] (right). Left two columns: black quartz. Right two columns: silver. Note how reflected ambient light is capable of cancelling out emission polarisation.

in figure 2.9: the emission polarisation of metals is stronger than that of dielectrics because the reflective polarisation is weak, and vice versa.

### Specular Reflections of Glowing Objects

One of the follow-on effects of this phenomenon can be seen in Figure 3.4. Since reflections from smooth phase boundaries are governed by the Fresnel terms which, amongst other things, also predict varying reflectance based on input polarisation, it stands to reason that the mirror image




Figure 3.4: Reflections of glowing spheres in a black glass surface without (a) and with (b) polarisation taken into account. If polarisation is taken into account, the appearance of the reflections is noticeably altered. The weak differences on the glowing spheres themselves are due to direct inter-reflections, which are also slightly altered. (c) shows the difference image.

of a glowing object should look different when polarisation is taken into account. Figure 3.4 shows that under certain circumstances, this effect can be strong enough to be directly visible. This fact is important in the prototyping of e.g. lamps where a bright light source is put into a strongly reflective environment.




Figure 3.5: A qualitative sketch of the linear polarisation patterns that can be seen on a clear sky.

## 3.5 Atmospheric Scattering

Scattering interactions can give rise to polarised light, which is e.g. the mechanism behind the in certain cases quite strong polarisation of the blue sky. Figure 3.5 gives an overview of the polarisation distribution found on clear skies, while figure 3.6 shows photographs of such skies.

A simple model for this behaviour is presented in [16], but very probably, more sophisticated simulations will emerge in the foreseeable future.





Figure 3.6: Photographs of sky-dome polarisation patterns as they can be seen in nature. **Left:** a fish-eye view towards the zenith, with a linear polariser placed in the lens so that the pattern on the right of figure 3.5 becomes visible. **Middle:** extreme wide-angle shot of an evening sky being reflected in a still lake. Observe how the blue of the sky is considerably darker than it should be in the middle of the reflected sky-dome – and in particular, how the white clouds, which are unpolarised, remain bright even in reflection. **Right:** the same kind of scene photographed with a linear filter placed in the lens. The correspondence between the polarised pattern on the sky-dome and the reflection becomes obvious.

### 3.6 Strain Visualisation – Stress-Induced Bi-refringence

One colourful phenomenon that involves polarised light is stress visualisation in materials which exhibit weak, wavelength-dependent bi-refringence. The typical material that such demonstrations are done with is plexiglass: figure 3.7 shows photographs of a plexiglass object in front of an LCD monitor.




Figure 3.7: Stress visualisations via pressure-induced bi-refringence in plexiglass. **Left:** the object in question in front of an LCD monitor, which, by design, emits almost perfectly polarised light. A uniformly white background image is shown on the monitor. **Center and right:** the same object viewed through a polarising filter, which is rotated by  $90^\circ$  between the two images. The colourful patterns are due to internal stress of the cheaply manufactured plexiglass box, which causes weak wavelength-dependent bi-refringence, the exact amount of which is stress-dependent.

It is noteworthy that the colourful patterns that one can see in the rightmost two images in figure 3.7 only appear if linearly polarised light traverses the material, *and then passes through yet another polariser!* The mechanism behind this is as follows: if the strain direction within the

material is at an angle to the polarisation direction of the incident light, the polarisation plane of the light is gradually rotated as it traverses the material. By how much it is rotated is dependent on the amount of stress in the material. Also, the polarisation of different wavelengths is usually rotated at different rates as light traverses such a material. As a result, for some wavelengths the linear polarisation of the traversing light ends up having been completely flipped during the traversal, while for others, it stays the same (usually via having been rotated a multiple of  $2\pi$  radians). If such a set-up is then viewed through a linear polarisation filter, one can observe the banded rainbow colours typical for such images. In this context, it should be noted that even the left image in figure 3.7 shows weak colour phenomena. These are due to inter-reflections within the plexiglass box, which act as weak polarising filters of sorts on the reflected light.

### 3.7 Organic Solutions

In certain transparent substances, such as solutions of chiral molecules such as sucrose (i.e. sugar), solids with rotated crystal planes such as quartz, and spin-polarized gases of atoms or molecules, the phenomenon of *optical rotation* can be observed. It amounts to the substance in question being able to alter the orientation of linearly polarised light that traverses it. This behaviour has several practical applications, such as being able to determine the concentration of a sugar solution in a non-invasive way by measuring the polarisation change that is induced by a sample of given size.

## Chapter 4

# Mathematically Describing Polarised Light

Physicists have developed several mathematical formalisms to describe the polarisation state of a given "light value"  $L_\lambda$  at a given frequency  $\lambda$ . Of these, the Stokes Vector and Coherency Matrix formalisms are the ones that are of interest to graphics: all extant graphics publications that deal with polarisation use one of these two techniques. Both have their relative merits and disadvantages, but from a graphics engineering standpoint, it would seem that on balance, Stokes Vectors are somewhat more convenient. This point could of course be disputed with some justification by someone who has had prior exposure to Coherency Matrices, but who has no working knowledge of Stokes Vectors yet. As is discussed later, Coherency Matrices can be used as drop-in replacements for Stokes Vectors, so to some extent, the choice between them amounts to personal preference, and practically everything what is being said in this tutorial about how one should use Stokes Vectors also applies to Coherency Matrices. The main reasons the authors of this tutorial favour the Stokes Vector formalism over Coherency Matrices are

1. that at least to them, the meaning of the components of Stokes Vectors is more intuitively obvious (which can be advantageous during debugging, but is also helpful during general interaction with such a system), and
2. that this formalism only involves real arithmetic, and no complex numbers.

For reasons of brevity, this chapter therefore exclusively focuses on Stokes Vectors, and their companion formalism, Mueller Matrices. We also briefly discuss one issue that is common to all such formalisms, namely that of coordinate system tracking. For conversion between Stokes Vectors and Coherency Matrices, we refer the reader to literature, such as [5].

### 4.1 Stokes Vectors

This particular technique to characterise the polarisation state of light was introduced by Sir George Gabriel Stokes in 1852. He found that for a given wavelength, four numbers are sufficient

to describe the polarisation state of a given transversal wave. The entire group of parameters is usually referred to as *Stokes Vector*  $S$ , while an individual parameter is referred to as *Stokes Component*  $S_i$ . This description only uses real-valued terms to describe all polarisation states of optical radiation. It also uses a noncomplex description of ray weights, or attenuation factors, in the form of Müller matrices [12]; these are discussed in the next section.




Figure 4.1: The Stokes Vector formalism: four real numbers –  $S_0$  through  $S_3$  – are used to describe the polarisation state of a wave at a given wavelength. The first component  $S_0$  encodes the radiant intensity, components  $S_1$  and  $S_2$  encode linear polarisation in two different reference frames rotated by  $45^\circ$ , and  $S_3$  is used to describe the circular polarisation component. Since  $S_0$  directly corresponds to the radiant intensity one already stores for each spectral sample in a normal spectral renderer, adding polarisation support just amounts to storing three additional real numbers per spectral sample in this case.

Actually, only three real-valued parameters would be required to describe a general polarisation ellipse. The slightly redundant, but very convenient, four value formalism with parameters  $\{S_0, S_1, S_2, S_3\}$  originated, and proved itself, in the optical measurements community, to which Stokes belonged. For graphics purposes, it has the key advantage that the first component  $S_0$  of this 4-vector is the unpolarised radiant intensity of the light wave in question, i.e. the same quantity that a nonpolarising renderer uses. Components  $S_1$  and  $S_2$  describe the preference of the wave towards linear polarisation at zero and 45 degrees, respectively, while the fourth,  $S_3$ , encodes preference for right-circular polarisation. While the first component is obviously always positive, the values for the three latter parameters are bounded by  $[-S_0, S_0]$ ; e.g. for a radiant intensity of  $S_0 = 2$ , a four-tuple of  $\{2, 0, 0, -2\}$  would indicate light which is totally left circularly polarised. Figure 4.1 gives a graphical representation of the individual Stokes Components. Note that the parameters  $S_1$  to  $S_3$  are also under the following constraint:

$$S_0 \geq \sqrt{S_1^2 + S_2^2 + S_3^2}. \quad (4.1)$$

In this tutorial, we will sometimes also refer to several Stokes Vectors at once, e.g. when discussing operations performed on several instances, such as in section 4.4. In these cases, we use the notation  $S_i(L)$  to denote the Stokes component  $i$  of a particular light value  $L$ .

## 4.2 Mueller Matrices

In a rendering system, we not only need a data structure to store light information, but also one to describe the *attenuation* of light<sup>1</sup>. In conventional rendering systems, this distinction is frequently blurred, since plain RGB or spectral values are normally used for both purposes. However, for a polarisation renderer, different structures are needed for the two data types.

In physics, the "attenuation counterpart" to Stokes Vectors are the so-called *Mueller matrices* –  $4 \times 4$  matrices which can describe all changes that the intensity and polarisation state of incident light can be subjected to.

## 4.3 Common Mueller Matrices

In this section, we quote a number of Mueller Matrices that are needed in the context of a polarisation-capable renderer.

### 4.3.1 Plain Attenuation

In the context of a renderer, "plain attenuation" means an idealised interaction of light with matter, where light passes in a straight line through a thin object without having its polarisation state altered – although attenuation can occur. A reasonable approximation for this behaviour are thin sheets of transparent, coloured plastic at normal incidence, e.g. when a scene is viewed through such a sheet. As it passes through such a material, the intensity of light is uniformly attenuated, but its characteristics are otherwise left unchanged. For an attenuation factor of  $\alpha$ , the corresponding Mueller matrix is given as

$$M_{attenuation} = \begin{bmatrix} \alpha & 0 & 0 & 0 \\ 0 & \alpha & 0 & 0 \\ 0 & 0 & \alpha & 0 \\ 0 & 0 & 0 & \alpha \end{bmatrix}. \quad (4.2)$$

### 4.3.2 Depolariser

Perfectly diffuse surfaces and transmissive materials act as *depolarisers*: the incident light intensity is attenuated to some degree, but any polarisation that is present in the incident light is

---

<sup>1</sup>In the context of a renderer, "attenuation" is a general term which encompasses all possible modifications made to travelling light, e.g. via reflection, scattering, absorption, fluorescent re-radiation, or any other event along its path through a scene.

destroyed. For an attenuation factor of  $\alpha$ , such a depolarising Mueller matrix is given by

$$\mathbf{M}_{depolariser} = \begin{bmatrix} \alpha & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \quad (4.3)$$

This is the Mueller matrix which has to be used for e.g. Lambertian or Oren-Nayar BRDFs.

### 4.3.3 The Ideal Linear Polarisation Filter

For a plane of linear polarisation that lies at an angle of  $\phi$  against the  $x$  axis, a perfect linear polariser has the following form ([12], p. 12):

$$\mathbf{M}_{linear} = \frac{1}{2} \cdot \begin{bmatrix} 1 & \cos 2\phi & \sin 2\phi & 0 \\ \cos 2\phi & \cos^2 2\phi & \sin 2\phi \cdot \cos 2\phi & 0 \\ \sin 2\phi & \sin 2\phi \cdot \cos 2\phi & \sin^2 2\phi & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad (4.4)$$

### 4.3.4 The Mueller Matrix for Fresnel Reflectance

Fresnel reflection off a perfectly smooth phase boundary, which we briefly discuss in section 2.2.1, is one of the two corner cases of reflectance where an exact formulas for the resulting Mueller Matrix is available (with a Lambertian surface as perfect depolarised being the other one). For this case, we can determine the Mueller matrix as being

$$\mathbf{M}_{Fresnel} = \begin{bmatrix} A & B & 0 & 0 \\ B & A & 0 & 0 \\ 0 & 0 & C & S \\ 0 & 0 & -S & C \end{bmatrix}, \quad (4.5)$$

with

$$\begin{aligned} A &= \frac{F_{\perp} + F_{\parallel}}{2} \\ B &= \frac{F_{\perp} - F_{\parallel}}{2} \\ C &= \cos(\delta_{\perp} - \delta_{\parallel}) \cdot \sqrt{F_{\perp} \cdot F_{\parallel}} \\ S &= \sin(\delta_{\perp} - \delta_{\parallel}) \cdot \sqrt{F_{\perp} \cdot F_{\parallel}} \end{aligned} \quad (4.6)$$

An explanation of the reasoning that allows one to arrive at this formula is given in chapter 8 of [5]. As per the definitions used in section 2.2.1,  $\delta_{\perp} - \delta_{\parallel}$  is the total retardance the incident wave-train is subjected to. The only difference between dielectrics and conductors is that the

formulas used for  $F_{\parallel}, F_{\perp}, \delta_{\parallel}, \delta_{\perp}$  simplify considerably if the IOR  $\eta$  is only real-valued – the form of the resulting Mueller Matrix given in equation 4.6 remains the same.

## 4.4 Stokes Vector and Mueller Matrix Calculus

In this section we briefly discuss which operations are permissible on instances of Stokes Vectors, and Mueller Matrices. In normal graphics systems, "anything goes" with respect to what mathematical operations one can apply to the colour values or spectral representations that are used to store light and attenuation<sup>2</sup> – component-wise scaling, addition or division (by a non-zero scalar) always work, and operations on these data types are commutative. However, for polarisation-aware systems, many of these assumptions do *not* hold!

In the following sections, we list those operations which are well-defined for each data type – preceded by a short section on the coordinate systems needed for these operations.

### 4.4.1 Coordinate System Tracking – Reference Frames

Regardless of which formalism one uses to describe light and attenuation, there is always a secondary engineering issue associated with polarisation-aware rendering: namely, that one must, at all times, track the coordinate system of each light or attenuation value one is describing.




Figure 4.2: Reference frames for Stokes Vectors and Mueller Matrices. Case **a**) shows a Stokes Vector, with its single reference frame. Case **b**) shows a Mueller Matrix for a reflection, where entry and exit frames are not the same. However, the two frames do not have to be dissimilar: case **c**) shows light passing through a polariser  $P$  at normal incidence, a scenario where both reference frames of a Mueller Matrix are identical.

Images of wave-trains propagating through space, such as figure 2.1, usually show a coordinate system in which the oscillations of the wave vector are taking place, and mathematical expressions like equation 2.1 implicitly assume a particular local coordinate system. The four values of

<sup>2</sup>In many rendering systems, not even an explicit distinction between light and attenuation values is made at the level of data modelling: typically, RGB values (or some spectral data type) are used for both.

a Stokes Vector, which we use to quantitatively describe such an oscillation state, do not make sense without the coordinate system they are defined in – so we have to store such a *reference frame* for each and every Stokes Vector we store!

Attenuation values such Mueller Matrices take this one step further, in that they require us to store *two* reference frames: one for the incoming light (which we, from the viewpoint of the attenuator, refer to as the *entering* light), and one for the reflected light (the *exiting* light). Figure 4.2 illustrates this: if we were to compute the Mueller Matrix for this particular reflection by using the Fresnel Terms, and equation 4.3.4, the resulting matrix only makes sense for light that enters via the direction of  $I$ , and exits via  $R$  – for any other set of directions, this matrix is simply meaningless!

Note that there are some interactions for which Mueller Matrices can be given, but which do not necessarily involve a change of direction, such as light simply passing through an idealised linear polarisation filter. From a data modelling standpoint, this does not matter, though: nothing prevents the entry and exit directions to be the same.

#### 4.4.2 Operations on Stokes Vectors

There is only a single operation that involves two light values  $L_0$  and  $L_1$  which makes sense from a physics viewpoint: their addition. All others, such as multiplication of two light values, do not correspond to any meaningful operation. Due to the principle of superposition, the addition of two light values can be done in a simple component-wise fashion:

$$S_i(L_r) = S_i(L_0) + S_i(L_1) \quad (4.7)$$

Since light addition is based on the principle of superposition, it is a commutative operation – the order in which light contributions are added to each other does not affect the end result.

For the sake of convenience, there are of course some other simple numerical operations which one can add to a library that deals with light representations that are based on Stokes Vectors. Uniform scaling of a given light (which corresponds to a plain attenuation Mueller matrix – see section 4.3.1), or component-wise exponentiation, are examples of this.

#### **Caveat: when adding Stokes Vectors is permissible, and when it is not**

It is crucial to note that even such an addition of Stokes Vectors is permissible if – and *only* if – the reference frame of the two light values match *exactly* – case **a**) in figure 4.3! If they do *not* match, we have to distinguish two scenarios:

The first scenario is where the two reference frames are already collinear, but not rotationally aligned yet – case **b**) in figure 4.3. In that case, the operation is meaningful: they can be brought






Figure 4.3: Addition of Stokes Vectors. Case **a)** works "out of the box": the two reference frames are aligned. In case **b)**, the reference frame of the second Stokes Vector is only collinear with the frame of the first, so it has to be rotated before the addition can take place. Case **c)** shows a meaningless operation: adding the Stokes Vectors of these two light waves, which propagate in different directions, makes no sense.

to a common reference frame, and then added. For an angle  $\phi$  between the two reference frames, the re-alignment operation that has to be performed prior to addition is as follows ([12], p. 25):

$$\begin{aligned}
 S_0(L_r) &= S_0(L_0) \\
 S_1(L_r) &= \cos(2\phi) \cdot S_1(L_0) + \sin(2\phi) \cdot S_2(L_0) \\
 S_2(L_r) &= -\sin(2\phi) \cdot S_1(L_0) + \cos(2\phi) \cdot S_2(L_0) \\
 S_3(L_r) &= S_3(L_0)
 \end{aligned} \tag{4.8}$$

The second scenario is that the two reference frames are completely askew (i.e. not even collinear) – case **c)** in figure 4.3. In that case, combining the two light values via addition is a semantically meaningless operation, and has to result in an error. If one encounters this kind of error in a rendering system, one should consider whether one is using the right data type for the computations at hand, and if not a light intensity data type should be used instead – see section 4.4.3.

### 4.4.3 Light vs. Light Intensity

In this context, it should be noted that in any practical polarisation rendering system, one will normally retain a separate data type that *only* describes the radiant intensity of light, and not its polarisation state as well. In our own rendering system, we refer to the data type that can potentially be polarised as `Light`, and to the storage type for radiant intensity as `LightIntensity`.

The latter type is the proper thing to use in all those cases where light *energy* is being stored, such as in a photon map, or a light map. As outlined in the previous section, adding the various non-collinear polarisation states of incident photons in a photon map is a semantically meaningless

operation – in those cases, all one is interested in is the non-polarised radiant intensity (a single value per wavelength, as in a non-polarising renderer), which can be safely added.

#### 4.4.4 Operations on Mueller Matrices

Attenuation elements like Mueller Matrices are the mirror image of Stokes Vectors insofar, as for these entities component-wise addition is an operation that is meaningless from a physics viewpoint. The only way to combine several Mueller Matrices  $M_0, M_1, \dots, M_n$  to a single result matrix  $M_r$  is via matrix multiplication (which is meaningless for Stokes Vectors):

$$M_r = M_0 \cdot M_1 \cdot \dots \cdot M_n \quad (4.9)$$

Matrix multiplications are of course not commutative, and the result matrix  $M_r$  corresponds to the matrices  $M_0, M_1, \dots, M_n$  being traversed in exactly the order in which they are being multiplied together!




Figure 4.4: A sequence of two polarising filters  $P_0$  and  $P_1$ , each of which is described by a Mueller Matrix  $M_0$  and  $M_1$ . If these are concatenated by multiplying  $M_0$  and  $M_1$ , the resulting matrix  $M_r$  has the entry frame of  $M_0$ , and the rotated (!) exit frame of  $M_1$  (which in this case is of course identical to the entry frame). Here, it is of course not necessary to perform an actual transformation on either of the reference frames of  $M_1$ , since one just uses the reference frames of  $M_0$ . As figure 4.5 shows, the same sort of consideration has to apply even if entry and exit frames are not identical – and there, an actual rotation is necessary!

As with the addition of Stokes Vectors, the multiplication of Mueller Matrices is only permissible if their reference frames match. However, since Mueller Matrices have two reference frames (entry and exit – see section 4.4.1), one has to be even more specific here: the exit frame of the first matrix has to match the entry frame of the second one, otherwise the matrix multiplication is again not semantically meaningful. Figure 4.4 shows a simple case of this in graphical form, and the rotation of a matrix  $M_0$  by an angle of  $\phi$  is given by:

$$M_{rotated} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos 2\phi & -\sin 2\phi & 0 \\ 0 & \sin 2\phi & \cos 2\phi & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdot M_0 \cdot \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos 2\phi & \sin 2\phi & 0 \\ 0 & -\sin 2\phi & \cos 2\phi & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (4.10)$$




Figure 4.5: A sequence of two specular reflections which are described by Mueller Matrices  $M_0$  and  $M_1$ . As with the linear filter sequence shown in figure 4.4, the exit frame of the Mueller Matrix that is being rotated (in this case,  $M_1$ ) has to be transformed for the result  $M_r$  to have a correct exit frame – but in this case, the frame has to be actually transformed, since it is not identical to one of the frames of  $M_0$ . The necessary rotation has to be performed around the propagation direction of the light at the interface under consideration.

As shown in figures 4.4 and 4.5, when multiplying two Mueller Matrices  $M_0$  and  $M_1$  with each other, one of them has to be rotated so that the exit frame of  $M_0$  and the entry frame of  $M_1$  are aligned. For the one that is being rotated, it is crucial that *both* its entry and exit frames are rotated by the same amount that is applied to the actual Mueller Matrix. In practice, if  $M_1$  is being rotated, the vectors of its entry frame do not have to be transformed, since they are not used after the operation anyway. But the exit frame of  $M_1$  will become the exit frame of the result  $M_r$ , so it has to be explicitly rotated around the direction in which light propagates when leaving  $M_1$ . As figure 4.4 shows, this is a fairly obvious thing to do for Mueller Matrices that describe an attenuation operation which does not involve a direction change. However, as one can see in figure 4.5, even reflections (and of course transmission events as well) have to perform a re-alignment of the exit frame of the transformed matrix, otherwise the end result would not be given in a correct reference coordinate system.

#### 4.4.5 Operations on Stokes Vectors and Mueller Matrices

There is of course one operation that takes both data types as input: if an attenuation element is applied to some incident light, a multiplication of the Stokes Vector  $S_0$  with the corresponding Mueller Matrix  $M_0$  is carried out to obtain the attenuated light  $S_r$ :

$$S_r = M_0 \cdot S_0 \quad (4.11)$$

Again, care has to be taken that the reference frame of the Stokes Vector, and the entry frame of the Mueller matrix match. The result Stokes Vector  $S_r$  will be given in the exit reference frame of the Mueller matrix.

## Chapter 5


# Building and Debugging a Polarisation Ray Tracer

### 5.1 Extending a Conventional Ray Tracer

The workings of a conventional ray or path tracer which are assumed as a given here – if not, see one of the many books which explain the underlying algorithms [11, 8]. If one takes the information given in the preceding chapters into account, there are several modifications that have to be made to such a conventional ray-based rendering system so that it can handle polarisation:

1. The data structures used for light and attenuation have to be replaced with data structures that contain Stokes Vectors and Mueller Matrices; figure 5.1 shows the data structures we use in our own polarisation ray tracer.
2. Certain easy optimisations ought to be built in from the get-go: such as each light and attenuation data structure carrying a flag which indicates whether the instance in question is actually polarising, or not. This information can be used to considerably speed up a polarisation ray tracer by avoiding unnecessary matrix multiplications during filter concatenations.
3. Reference frame tracking has to be implemented for all paths that are traced. This is a comparatively simple geometric problem, as one only has to store information which is a by-product of the normal ray-object intersections anyway. It does involve considerable extensions to all functions which deal with ray propagation, though.
4. All BRDF and scattering models have to be altered so that they yield the correct Mueller Matrix upon evaluation. As indicated in section 2.2, this is not an easy task, since we only have reliable BRDF models for trivial corner cases (diffuse and perfectly specular) so far. However, as also discussed in that section, intermediate solutions for glossy surfaces do exist.

```
typedef struct ArSVLight
{
    unsigned int    polarised;
    ArStokesVector * stokesVector;
    ArReferenceFrame referenceFrame;
}
ArSVLight;
```



```
typedef struct ArMMAAttenuation
{
    ArMMAProperties    properties;
    ArMuellerMatrix * muellerMatrix;
    ArReferenceFrame referenceFrameEntry;
    ArReferenceFrame referenceFrameExit;
}
ArMMAAttenuation;
```




Figure 5.1: C data structures used within ART (our own polarisation-capable rendering research toolkit) to describe Stokes Vectors and Mueller Matrices. The actual spectrally resolved data structures for Stokes Vectors and Mueller Matrices are encapsulated within a struct that contains reference frame information, and flags – such as whether the Stokes Vector in question is actually polarised, or not.

5. A workflow that involves a polarisation capable image raw format has to be implemented. Since no standardised formats of this kind exist yet, one has to develop one's own solution here at the moment. Within an already existing rendering framework, a technically easy way to do this in a temporary fashion would be to save and read four spectral images in the existing ray image file format of the toolkit in question – one for each Stokes Component.

## 5.2 Visualisation Tools as Part of the Working Environment

Amongst other things, it is helpful to have proper visualisation tools so one can diagnose what the system is doing. In [17], a set of such conventions is presented and discussed: figure 5.2 shows samples of this sort of visualisation for a test scene. Even though it would be desirable if practitioners in this area used similar conventions for such visualisations, the exact form of these visualisation is not as important as the fact that such a functionality be available in a polarisation ray tracing system. Without it, debugging such a system is much more of a hit-and-miss affair than without.




Figure 5.2: A test scene (left) with an overlay degree of polarisation plot (center), and a false colour plot that visualises the orientation of the linear reflective polarisation (right). Such false-colour renderings can be invaluable tools to determine whether polarisation is present in the right places in an image, and whether it has the right form and orientation.

# Bibliography

- [1] Changhyuk An and Kyle Zeringue. Emission polarization from rough surfaces. *Polarization: Measurement, Analysis, and Remote Sensing VIII*, 6972(1):69720O, 2008.
- [2] Max Born and Emil Wolf. *Principles of Optics*. The Macmillan Company, 1964.
- [3] Kate Devlin, Alan Chalmers, Alexander Wilkie, and Werner Purgathofer. Tone reproduction and physically based spectral rendering. In Dieter Fellner and Roberto Scopigno, editors, *STAR Proceedings of Eurographics 2002*, Saarbruecken, Germany, September 2002. Eurographics Association.
- [4] Abhijeet Ghosh, Tim Hawkins, Pieter Peers, Sune Frederiksen, and Paul E. Debevec. Practical modeling and acquisition of layered facial reflectance. *ACM Trans. Graph.*, 27(5):139, 2008.
- [5] Dennis Goldstein. *Polarized Light*. Marcel Dekker Inc., New York, USA, 2003.
- [6] Xiao D. He, Kenneth E. Torrance, François X. Sillion, and Donald P. Greenberg. A comprehensive physical model for light reflection. *Computer Graphics*, 25(4):175–186, July 1991.
- [7] G.P. Koennen. *Polarized Light in Nature*. Cambridge University Press, 1985.
- [8] Matt Pharr and Greg Humphreys. *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2004.
- [9] Andrew Resnick, Chris Persons, and George Lindquist. Polarized emissivity and Kirchhoff’s Law. *Appl. Opt.*, 38(8):1384–1387, 1999.
- [10] Oscar Sandus. A review of emission polarization. *Appl. Opt.*, 4(12):1634–1642, 1965.
- [11] Peter Shirley. *Realistic ray tracing*. A K Peters, 2000.
- [12] John B. Shumaker. Distribution of optical radiation with respect to polarization. In Fred E. Nicodemus, editor, *Self-Study Manual on Optical Radiation Measurements, Part 1: Concepts*. Optical Physics Division, Institute for Basic Standards, National Bureau of Standards, Washington, D.C., June 1977.

- [13] Robert Siegel and John R. Howell. *Thermal Radiation Heat Transfer, 4th Edition*. Taylor & Francis, New York, NY, 2001.
- [14] Andrea Weidlich and Alexander Wilkie. Arbitrarily layered micro-facet surfaces. In Andrew Rohl, editor, *GRAPHITE*, pages 171–178. ACM, 2007.
- [15] Andrea Weidlich and Alexander Wilkie. Realistic rendering of birefringency in uniaxial crystals. *ACM Trans. Graph*, 27(1), 2008.
- [16] A. Wilkie, C. Ulbricht, Robert F. Tobler, G. Zotti, and W. Purgathofer. An analytical model for skylight polarisation. In Alexander Keller and Henrik Wann Jensen, editors, *Eurographics Symposium on Rendering*, pages 387–397, Norrköping, Sweden, 2004. Eurographics Association.
- [17] Alexander Wilkie and Andrea Weidlich. A standardised polarisation visualisation for images. In *Proceedings of the Spring Conference on Computer Graphics 2010 (SCCG 2010)*, May 2010.
- [18] Alexander Wilkie and Andrea Weidlich. A physically plausible model for light emission from glowing solid objects. In Erik Reinhard and Ravi Ramamoorthi, editors, *Eurographics Symposium on Rendering*, Prague, Czech Republic, 2011. Eurographics Association.
- [19] L. B. Wolff, A. J. Lundberg, and R. Tang. Image understanding from thermal emission polarization. In *Image Understanding Workshop*, pages 747–754, 1998.
- [20] Lawrence B. Wolff and David Kurlander. Ray tracing with polarization parameters. *IEEE Computer Graphics & Applications*, 10(6):44–55, November 1990.