A Model for Anisotropic Reflection

Pierre Poulin
Alain Fournier

Department of Computer Science
University of Toronto
Toronto, Ontario, Canada
M5S 1A4
Currently:
Department of Computer Science
University of British Columbia
Vancouver, BC, Canada
V6T 1W5
{poulin | fournier}@cs.ubc.edu

Abstract

A reflection and refraction model for anisotropic surfaces is introduced. The anisotropy is simulated by small cylinders (added or subtracted) distributed on the anisotropic surface. Different levels of anisotropy are achieved by varying the distance between each cylinder and/or rising the cylinders more or less from the surface. Multidirectional anisotropy is modelled by orienting groups of cylinders in different direction. The intensity of the reflected light is computed by determining the visible and illuminated portion of the cylinders, taking self-blocking into account. We present two techniques to compute this in practice. In one the intensity is computed by sampling the surface of the cylinders. The other is an analytic solution. In the case of the diffuse component, the solution is exact. In the case of the specular component, an approximation is developed using a Chebyshev polynomial approximation of the specular term, and integrating the polynomial.

This model can be implemented easily within most rendering system, given a suitable mechanism to define and alter surface tangents. The effectiveness of the model and the visual importance of anisotropy are illustrated with some pictures.

CR Categories and Subject Descriptors: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism.
General Terms: Algorithms.
Additional Key Words and Phrases: Shadowing, surface mapping, Chebyshev polynomials, hair rendering, scientific visualization.

1 Introduction

Many objects in nature are visually very complex. Crow, as quoted in [13], suggested having different levels of geometric models to capture this complexity. However, as the representation of an object requires the addition of smaller geometric models to define smaller details, the extra cost of modelling and rendering these details can become very high. Moreover such an object representation would be prone to aliasing problems.

To overcome these problems, approximations to some geometric models have been proposed. If we note D_g the level of the geometry at which objects are usually defined (i.e. polygons or parametric surfaces), then displacement and bump mappings [4] [16] approximate what the geometric model D_{g+1} would define at a higher resolution or magnification. As the details become smaller, even surface mapping becomes highly subject to aliasing problems. At this geometric level, D_{g+2}, the details cannot be seen individually and appear only because of the way they modulate the light reflected off them. This is the reason why these details must be captured by local reflection models.

Traditionally, reflection models have been divided into two components: *diffuse* and *specular* [24] [3] [8] [2]. The diffuse component takes into account the light that interreflects onto elements of a same surface and is reemitted equally in all directions. To model the specular reflection, Torrance and Sparrow [24] assume that the surface is made of highly reflective microscopic facets distributed in v-grooves. If the facets are randomly distributed over the surface, shadowing and masking functions can be statistically estimated, and, for a given distribution of slopes of the facets, the light reflected in a particular direction can be approximated.[1] This type of surface is qualified as *isotropic* because the reflected light intensity at a given point is independent of the surface orientation along its normal at this point.

Many surfaces, however, cannot be modelled by a random isotropic distribution of facets. If an element of such surfaces is rotated around its normal while the light and viewer directions remain unchanged, the light intensity reflected to the viewer will vary. These surfaces are called *anisotropic*. Surfaces made of fur or burnished metal are examples where there is a strong correlation between the orientations of facets and the orientation of the surface. Brennan et al. [6] measured several natural surfaces and showed that surfaces like clouds, forests, oceans and even sand exhibit anisotropy. Anisotropy can be caused by a collection of strongly oriented elements, such as hair in fur or blades

[1] All the reflection models discussed in this paper ignore the wavelength. In the facets model, this means that the size of the facets is assumed much greater than the wavelength of the light source.

of grass in a meadow, or by the selected action of an external force. Many natural factors can contribute to preferred orientations of the facets like the wind, the sun's position for the orientation of leaves or the shape of the underwater terrain for waves. Since anisotropy is so prevalent, and our visual system use it as a cue for orientation and depth, a wide range of anisotropic surfaces needs to be accurately rendered to obtain more realistic images.

Moreover, the use of anisotropy is not restricted to the production of realistic images. With the recent emphasis on scientific visualization techniques, more and more information needs to be communicated to the viewers from a single image. The shape, colour, texture, opacity and surface roughness are some properties that can carry information. Anisotropy adds a new dimension to the information transmitted. For instance, a simple vector field over a surface can be associated to a given orientation of the facets on this surface, freeing colour or texture for other information.

This paper first briefly surveys and discusses current anisotropic reflection models. The model based on cylinders is then presented in detail, with its parameters and solutions for the reflected light. Finally images illustrating applications of anisotropy in computer graphics are presented and commented.

2 Previous Work

Kajiya [13] attempted to compute analytically the reflected intensity from a continuous surface. He bases his approach on the general Kirchhoff solution for scattering of electromagnetic waves [1] [2]. For given incident and reflected directions, the intensity reflected by a surface is computed. However, the method has its restrictions. For instance, the Kirchhoff solution is valid only if self-shadowing and multiple scattering are negligible. Even if this limitation is not considered, the size of the surface required (Fresnel zones) by the Kirchhoff solution and the stationary phase method used to approximate the integral [5] introduce new problems that are dependent of each surface type.

Cabral et al. [7] go down to the facet level. Facets are created from a height field and the reflection off each facet is studied, including the blocking factor for incident and reflected light. This method, which can be qualified as *brute force*, is computationally expensive. The computation of the blocking factor is done via a modification of Max's method for the self-shadowing of bump maps [16]. This method can seriously alias the shadows for surfaces which exhibit high frequency behaviour.

In these two approaches, the reflection intensities are computed once and kept in tables; interpolation is used for fast rendering. However, a new table needs to be computed each time another type of surface is being rendered, which involves a few hours of CPU time (≈ 12 hours on an IBM4341 for Kajiya's and around the same time for Cabral's on a VAX 11/785).

In Phong's reflection model, the distribution of the specular intensity is symmetric around the reflection direction. Ohira [29] defines it as elliptic with the intention to simulate the reflection from scratches-like surfaces. However, the definition of the ellipse has no physical motivation and can lead, according to Yokoi et al., to "unnatural images". Takagi et al. [29] extend this idea of an ellipsoid to the reflection model of Torrance and Sparrow [24]. For a surface covered with scratches, the distribution of the normals of the facets can be approximated by an ellipsoid elongated in the preferred orientation of the scratches. For given light and viewer directions, the bisector between these two directions is computed. Then the intensity reflected specularly to the viewer is considered directly proportional to the ratio of facets with their normal in this bisector direction. These last two models are only valid for "scratched" surfaces with negligible self-shadowing and multiple scattering.

Miller [17] introduced the use of cylinders to simulate anisotropy. When projecting a cylinder onto the viewing screen, only half of it is visible. The intensities at equal intervals of this projected cylinder are computed and averaged. Later, at the rendering stage, if scratches are oriented in the direction of one cylinder, the computed intensity is used to simulate the effect of the anisotropy of the surface.

Miller also adapts the *reflectance sphere* [28] for faster rendering. However, the technique inherits the limitations of the reflectance sphere. Thus, light sources are expected to be far from all the objects in a scene. The projection of the cylinder onto the viewing screen is also restricted to perspective rendering of a scene and therefore, unsuitable for a rendering technique like ray tracing. Another disadvantage is that each cylinder is treated individually, which obliges Miller to neglect the self-shadowing and the inter-reflection of one cylinder onto its neighbours. Therefore in Miller's model, for a given orientation of a cylinder, the same intensity will be computed, whatever the orientation of the surface normal.

3 Cylindrical Scratches

The model proposed here builds on the concept of Miller's cylinders. Scratches are represented by a large number of adjacent small cylinders. The intensity reflected by all the cylinders in one direction is approximated by the reflection off only a single cylinder. This approximation is accurate if many cylinders cover every sampled region of the surface. Moreover, since the cylinders have a very small radius in comparison to their length, the intensity reflected off a cylinder can be approximated by the reflection off only one cross section of this cylinder.

Consider the hierarchy of geometric models discussed previously. At the geometric level D_{g+2}, the details cannot be seen individually and are captured by the reflection model. Including our anisotropic reflection model corresponds to inserting two new geometric levels between the mapping (displacement or bump) and the isotropic reflection model, now identified as D_{g+4}. The isotropic reflection model D_{g+4}, like the facets model of Torrance and Sparrow, characterizes the surface nature of each cylinder. A group of adjacent cylinders oriented in one direction defines the D_{g+3} level while a set of groups of adjacent cylinders represents the D_{g+2} level. Figure 1 illustrates the interaction between these various levels.

To orient the cylinders, surface frame bundles (to use the terminology of [13]) need to be established. One axis is the normal at the surface, N. Another axis is the tangent to the surface, T (thus perpendicular to the normal). This tangent allows the specification of the orientation of the cylinders on the surface. The last axis defined is the binormal, B, formed by the cross product $T \times N$. The cross section of the cylinder of interest thus lies on the plane NB.

Since the normals are already defined for most of the surfaces used in computer graphics, the extra information required for introducing anisotropy is simply a tangent at each point of the surface. Many techniques have been developed to map textures [11] and perturbations of normals [4] onto objects; these techniques can easily be extended to map tangents onto objects.

The cylinders are too small to be seen individually. For each cylinder, the viewer position is so far away relative to the cylinder's radius that its direction can be assumed to be

Figure 1: Hierarchy of Geometric Models

constant over the cylinder. By the same token, the light direction is also assumed constant over each cylinder. Groups of adjacent cylinders are also assumed to be small enough to not be seen individually. Moreover, the surface covered by parallel and adjacent cylinders (within one group) is assumed to be so large in comparison with the surface covered by non-parallel and adjacent cylinders (cylinders on the edges of two groups with cylinders oriented differently) that the self-shadowing and the inter-reflection from one group on the other will be considered as negligible.

We will now examine in details the light reflected by individual cylinders. The result is the heart of the anisotropic model introduced in this paper.

3.1 Controlling the Anisotropy

The anisotropy in this model is created because the distribution of the surface normals along the cross section of the cylinders is different from the distribution along their axis; this difference determines the level of anisotropy. Two parameters are used to control the anisotropy. One is the distance d between the centers of two adjacent cylinders. Without lost of generality, assume a unit radius for each cylinder. This means that an intersection point on a cylinder defines directly the normal at this point. If the distance $d \in [0, \infty)$ is 0, only the surface normal is taken (i.e. no variation of normals) and the resulting effect is given simply by the underlying reflection model (Torrance and Sparrow in this case). As d goes from 0 to 2, the anisotropy increases to its maximum (the variation of normals cannot be greater than π). If the distance is greater than 2, a *floor* between adjacent cylinders appears. The resulting anisotropy is then a combination of the varying normals across the cylinders and the length of the floor where the surface normal does not vary.

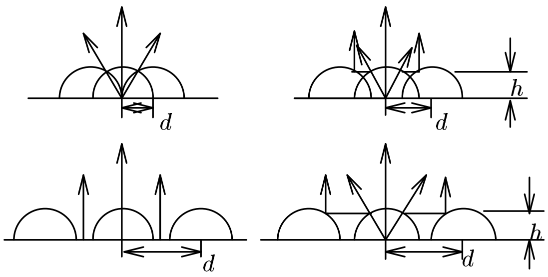

Figure 2: Controlling the Level of Anisotropy

The anisotropy is also controlled by a second parameter, the *height* $h \in [0, 1]$ by which the floor is raised. A higher floor reduces the variation of the normals on the cylinders and increases the length of the floor. Figure 2 illustrates the control of the anisotropy by these two parameters, d and h.

Defining an anisotropic surface relies on our intuitive knowledge of the surface. If the anisotropy is less pronounced (like in the scratched stainless steel surfaces shown in Image III and on the knobs of Image IV), a small d and no floor produce these characteristics. For surfaces like hair and textile (see the drapery of Image III), d is expected to be closer to 2. The roof of barns are often covered with corrugated iron that corresponds well to cylinders spaced by a large floor, $d > 2$.

A complementary model using negative (transparent) cylinders has also been developed. This model allows for (a) large grooves with relatively sharp edges or (b) for sharp grooves spaced by a longer floor. Conceptually, this model is more suitable for (a) surfaces like water with waves forming crests or (b) records (see the record on the turntable in Image IV). The equations relative to the model are given in appendix A. Although the concepts for both models are quite similar, the rest of the paper will treat only positive cylinders in order to simplify explanations.

3.2 Shadowing and Hiding

Consider one cylinder. Because of the parameters d and h, the normals on the cylinder can vary only within a certain angle θ_M from the surface normal:

$$\theta_M = \min(\phi_d, \phi_h)$$

where $\phi_h = \cos^{-1}(h)$

$$\phi_d = \begin{cases} \sin^{-1}(d/2) & \text{if } d < 2 \\ \frac{\pi}{2} & \text{otherwise.} \end{cases}$$

Depending on the orientation of the light vector L, some parts of the cylinder might not be lit. Let θ_L be the angle between the light vector projected onto the NB plane and the surface normal. The self-shadowing angle θ_{ss} corresponds to the angle from the surface normal at which the cylinder starts blocking light onto itself. The shadowing angle θ_s corresponds to the angle at which the neighbouring cylinder starts casting a shadow onto the cylinder. For a floor of length f, some length f_i is lit and some length f_s is in the shadow of the cylinder. All of these values can be computed as follows:

$$\theta_{ss} = \frac{\pi}{2} - \theta_L$$
$$f = \max(d - 2\sqrt{1 - h^2}, 0)$$

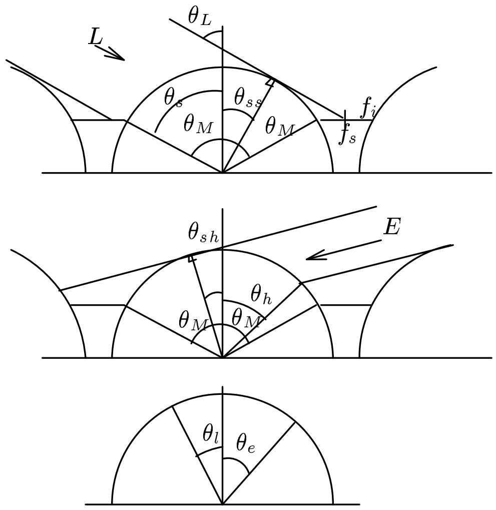

Figure 3: Some Variables for the Positive Cylinders

if $(\theta_{ss} > \theta_M)$
{
$\theta_s = \theta_{ss} = \theta_M$
$f_i = f; f_s = 0$
}
else
{
if $(\frac{1}{\cos\theta_L} < (d-1))$
$\theta_s = \frac{\pi}{2}$
else
$\theta_s = \theta_L + \sin^{-1}(d\cos\theta_L - 1)$
$\theta_s = \min(\theta_s, \theta_M)$
$f_s = \min(\frac{1}{\cos\theta_L} - \sqrt{1-h^2} - (\frac{h\sin\theta_L}{\cos\theta_L}), f)$
$f_s = \max(f_s, 0)$
$f_i = f - f_s$
}

The same logic is used to obtain the self-hiding angle θ_{sh} (the angle at which the cylinder starts hiding itself from the viewing position), the hiding angle θ_h (the angle at which the neighbour cylinder starts hiding the cylinder from the viewing position), the length of the hidden floor f_h and the length of the visible floor f_v. It is only necessary to replace θ_L by θ_E, the angle between the viewing vector projected onto the NB plane and the surface normal. Figure 3 identifies the various angles for given light and viewer positions.

3.3 Sampling as Seen by the Viewer

At this point, all the information about the visible and illuminated portions of the floor and the cylinder is known. It is now necessary to compute the light intensity reflected in the viewing direction. Since the floor has a constant normal and the viewer position is assumed at infinity, the intensity reflected by the floor is computed via any favorite reflection model.

Computing the light reflected off the cylinder to the viewer is more involved. One solution consists of sampling the cylinder at regular intervals as seen by the viewer. Figure 4 illustrates this process and figure 5 shows how the sampled normals are computed. Thus, the visible and illuminated arc is projected onto the surface plane and this

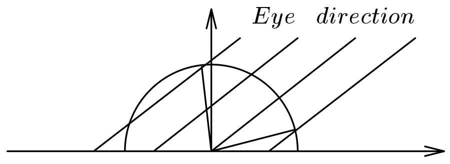

Figure 4: Regular Sampling from the Viewer

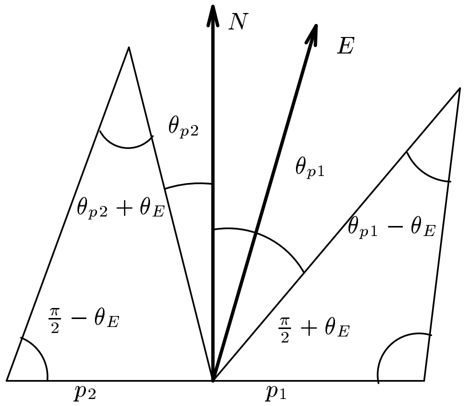

Figure 5: Finding the Sampling Angles θ_p

projection is subdivided equally. For each subdivision, the corresponding angle on the cylinder is

$$\begin{aligned} \theta_{p_1} &= \sin^{-1}(p_1\cos\theta_E) + \theta_E \\ \theta_{p_2} &= \sin^{-1}(p_2\cos\theta_E) - \theta_E \end{aligned}$$

depending on which side of N the projection of the sampling angle lies.

The intensity reflected to the viewer is averaged from the intensities computed at all sampling normals. If the anisotropic surface is reflective, these normals are used to determine the directions in which reflected rays are generated. For transparent objects, the same normals are used. To reduce the strong aliasing patterns susceptible to appear, the new rays are simply jittered. Finally, the new rays might intersect the neighbouring cylinders. In this case, only the colour of the surface is considered.

The intensities computed so far are only for the visible and illuminated arc I_a and floor I_f. However visible portions of the surface can be in shadow. For these portions, the diffuse and specular intensities are zero. In order to be able to treat the floor and the cylinder in the same units, it is necessary to project the cylinder onto the surface plane. Then, the lengths of the visible part of the cylinder l_v and of the visible and illuminated part l_{vi} are given by

$$l_v = \frac{\sin(\theta_h - \theta_E) + \sin(\theta_{sh} + \theta_E)}{\cos\theta_E}$$

$$l_{vi} = \begin{cases} \frac{\sin(\theta_l - \theta_E) + \sin(\theta_e + \theta_E)}{\cos\theta_E} & L,E \text{ same side of } N \\ \frac{\sin(\theta_e - \theta_E) + \sin(\theta_l + \theta_E)}{\cos\theta_E} & \text{otherwise} \end{cases}$$

where θ_l and θ_e define the visible and illuminated arc, θ_l being the angle between the surface normal and the point where the visible and illuminated portion of the cylinder begins from the side of the light source, and θ_e being defined like θ_l but from the other side (see figure 3).

Now, the intensity I_r reflected to the viewer can be normalized by

$$I_r = \frac{(I_a \times l_{vi}) + (I_f \times f_{vi})}{l_v + f_v}$$

where I_r is the fraction of the incident light that is reflected to the viewer.

3.4 Analytic and Approximate Solutions

Sampling the reflected intensity has some advantages, especially in the context of a rendering method such as ray tracing. In this case, when the surface is reflective or refractive, the same sampled normals can be used to compute the directions of the reflected or refracted rays. Sampling, however, has many drawbacks. It is usually prone to aliasing. As the specular coefficient n of a reflection model like Phong's grows, the peak of the intensity reflected becomes thinner and could easily be missed when sampling the cylinder. As a result it is difficult to determine the correct number of samples for a visually satisfactory result at a reasonable cost.

In order to avoid sampling problems, an analytic solution for the relevant integrals would be useful. An analytic solution for the diffuse term of the reflectance model is introduced in this section. The solution for the specular term is not as easy, but a solution is presented where the reflection function is approximated by Chebyshev polynomials and an analytic solution for the integral of this approximation is used.

3.4.1 Reflected Diffuse Intensity

In many reflectance model, the term $(N \cdot L)$ determines the light reflected diffusely. If a coordinate system in two dimensions is established in the NB plane[2], summing the contributions as a function of $(N \cdot L)$ is expressed as:

$$\int_{\phi_i}^{\phi_s} (N \cdot L) \cos(\phi_E - \phi) d\phi \tag{1}$$

where the term $\cos(\phi_E - \phi)$ is a correction factor due to the viewing position. In the NBT coordinate system, L is expressed as (L_n, L_b, L_t) and N as $(\sin\phi, \cos\phi, 0)$. The solution to equation 1 is

$$\frac{L_n}{2}\left(\begin{array}{l}\cos\phi_E(\sin^2\phi_s - sin^2\phi_i) + \\ \sin\phi_E\,(\phi_s - \phi_i - \sin\phi_s\cos\phi_s + \sin\phi_i\cos\phi_i)\end{array}\right) + \frac{L_b}{2}\left(\begin{array}{l}\sin\phi_E(\sin^2\phi_s - sin^2\phi_i) + \\ \cos\phi_E\,(\phi_s - \phi_i + \sin\phi_s\cos\phi_s + \sin\phi_i\cos\phi_i)\end{array}\right)$$

Since it was necessary to include a correction factor due to the viewer's position, this solution must be normalized by dividing by

$$\int_{\phi_i}^{\phi_s} \cos(\phi_E - \phi) d\phi = \begin{array}{l}\cos\phi_E(\sin\phi_s - \sin\phi_i) + \\ \sin\phi_E(\cos\phi_i - \cos\phi_s)\end{array}$$

3.4.2 Reflected Specular Intensity

If the term $(N \cdot H)^n$ is used to determine the light reflected specularly[3], an analytic solution similar to the one above could be computed assuming n integer, but it degenerates into an impressive number of terms for each intersection point.

If the curve of $\cos^n\phi$ is plotted as a function of ϕ, one can observe that it is symmetric about the Y axis and relatively smooth. Thus, this curve can easily be approximated by one or more polynomials which are simple to integrate.

It is important to mention that the approximating polynomials need only be computed once for a given n. The polynomials are thereafter only referred to at the rendering stage.

One approximation is based on Chebyshev polynomials. If regular samples are taken on the curve, Chebyshev approximation is assured to return the best (min-max) polynomial of a given degree. In order to reduce the degree of the polynomials for a given tolerance, some simplifications are necessary. First, since the curve is symmetric about the Y axis, the approximation can be limited to the positive side of the X axis. It is also known that the reflected intensity for an angle larger than $\frac{\pi}{2}$ is zero. Moreover, since we deal with discrete intensities at a pixel level, we can use this fact to bound the error of the polynomial approximation. If the intensity of an incident light ranges from 0 to 1, [4] the difference in intensity between the real value and the approximate one can be at most $\frac{1}{256}$ if we want to be at ± 1 rgb value. Then the approximate polynomial can cover the domain $[0,\phi]$ where $\phi = \cos^{-1}(\frac{1}{256}^{1/n})$ without violating the $\frac{1}{256}$ accuracy. With this assumption, several curves with various n have been approximated and, by experience, a degree 6 polynomial has always been sufficient to respect this criterion.

A short piece of code returning Chebyshev polynomials is given in appendix B. The advantage of this technique is that not only can the curve $(N \cdot H)^n$ be approximated easily, but so can any smooth curve such as most of the curves given by other reflection models.

The solution for $(N \cdot H)^n$ follows. To find the intensity reflected specularly to the viewer, an integral of the form

$$\int_{\phi_i}^{\phi_s} (N \cdot H)^n \cos(\phi_E - \phi) d\phi \tag{2}$$

has to be solved. Figure 6 illustrates the variables used to solve equation 2; N' is a sampling normal.

Equation 2 can be rewritten as

$$(H\cos\phi_p)^n \int_{\alpha_0-\phi_s}^{\alpha_0-\phi_i} \cos^n\alpha \cos(\phi_E + \alpha - \alpha_0) d\alpha$$

for which the integral part is evaluated as

$$\begin{array}{l}\cos\phi_E\left(\cos\alpha_0 \int\cos^{n+1}\alpha d\alpha + \sin\alpha_0 \int\cos^n\alpha\sin\alpha d\alpha\right) - \\ \sin\phi_E\left(\cos\alpha_0 \int\cos^n\alpha\sin\alpha d\alpha - \sin\alpha_0 \int\cos^{n+1}\alpha d\alpha\right)\end{array}$$

In this solution, $\int\cos^{n+1}\alpha d\alpha$ is approximated via Chebyshev polynomials and

$$\int \cos^n\alpha \sin\alpha d\alpha = \frac{-\cos^{n+1}\alpha}{n+1}$$

There is another, simpler method to approximate the integral needed for the specular contribution. The first observation is that the Phong model is itself only an approximation. A very similar distribution is obtained when, as

[2] ϕ in this coordinate system is calculated from the positive axis of B. Then ϕ_E in this coordinate system corresponds to $\frac{\pi}{2} - \theta_E$ or $\frac{\pi}{2} + \theta_E$, depending on which side of the normal N the viewer is.

[3] H being the bisector vector between L and E and n Phong's specular coefficient.

[4] This is an unrealistic limitation but many current rendering systems use this range.

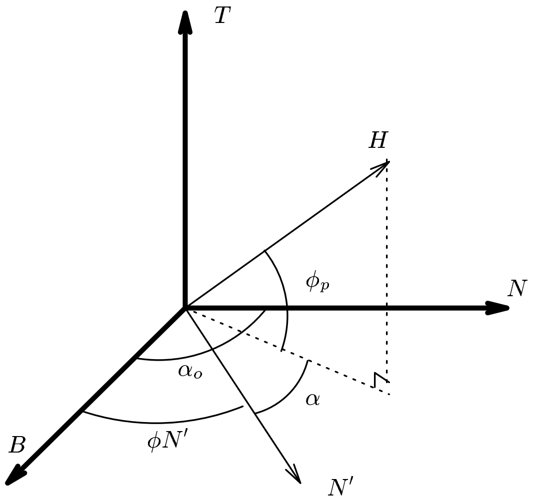

Figure 6: Integrating $(N \cdot H)^n$

in the Torrance-Sparrow reflection model [24], one takes a Gaussian distribution around the highlight: $e^{-(k\phi)^2}$. As suggested by Blinn [3] the two distributions can be fitted to each other by specifying a common angle β for which the intensity falls to one half of the peak value. In this case, the Phong exponent n is given by:

$$n = -\frac{\log 2}{\log(\cos \beta)}$$

and the coefficient k is given by:

$$k = \frac{\sqrt{\log 2}}{\beta}$$

We can therefore replace the Phong expression by the Gaussian distribution with the same characteristics using only one parameter. The result of the total integral is immediate from the knowledge of the integral of the Gaussian:

$$\int_{-\infty}^{\infty} e^{-(k\phi)^2} d\phi = \frac{2}{k\sqrt{\pi}}$$

In general what is needed is some definite integral. In this case one can replace the Gaussian by its classic piece-wise approximation by a cubic spline, also known as the Parzen window (scaled to have a maximum of 1). The polynomials are:

$$\begin{cases} \frac{1}{4}(2+u)^3 & \text{when} \quad -2 \le u < -1 \\ \frac{1}{4}(4-6u^2-3u^3) & \text{when} \quad -1 \le u < 0 \\ \frac{1}{4}(4-6u^2+3u^3) & \text{when} \quad 0 \le u < 1 \\ \frac{1}{4}(2-u)^3 & \text{when} \quad 1 \le u < 2 \\ 0 & \text{when} \quad |u| > 2 \end{cases}$$

The parameter u is scaled as a function of β:

$$u' = \frac{0.7223517u\pi}{2\beta}$$

For the usual range of n, from 4 and up, the integral computed this way is never more than 1% away from the corresponding Phong integral. If n is less, then the approximation of the Gaussian integral can deviate by up to 10% from the Phong integral. It must be remembered, however, than one has as much claim to the truth as the other.

It is not critical for the present problem to have a fast method to replace the Phong expression by an integrable function, since it is done only once per distinct value of n. Once the polynomials are obtained, computing the integrals is quite close in cost for the two methods. The second method will be of interest in other applications (such as normalizing the energy redistributed by a specular reflector) where it is important to have a fast approximation of the total integral.

4 Rendering Anisotropic Surfaces

The model of cylindrical scratches presented in this paper can be used within several rendering techniques. The only requirements are to know the surface normals, tangents and the viewer and light positions. The images provided in this paper have been produced via ray tracing which is very attractive for this model because of its "ray" nature.

Techniques relying only on scanline conversion are also suitable to render anisotropic surfaces through this model. However, the reflective and transparent anisotropic surfaces have to be treated differently. For instance, if some environment is mapped onto the surface, the characteristics of the cylinders in the model can be used to determine the colour reflected/refracted to the viewer. This colour will therefore be a weighted average of some portion of the environment map.

Techniques intended to solve the global illumination of scenes can also include anisotropic surfaces. Path tracing [14], combinations of radiosity and ray tracing [12][22] and the light driven approach of FIAT [9] are some examples of techniques where this model can be applied, demonstrating its flexibility.

The equations used in this paper include many trigonometric functions. The equations have been presented as such in order to simplify their understanding and make them as general as possible. However many times the angles themselves do not have to be computed since only their sines and cosines appear in the actual code. In that case the cosines (and consequently sines) can be evaluated via dot products. Finally rules to decompose trigonometric functions can be used.

Image	Description	Relative Time
Image Ia and Ib	Varying d and h	1.27
Image II	Reflective surfaces	
	isotropic	1.00
	bumps	1.09
	longitudinal	14.55
	latitudinal	14.75
Image III	Kitchen utensils	1.06
Image IV	Sound Design	1.03

The table given above indicates the additional rendering cost for the images presented in this paper. The numbers given are times relative to rendering the same scene without anisotropy. Note that for the longitudinal and latitudinal anisotropy of Image II, no jittering was used and as much as 90 reflected rays where shot in the environment.

In the current implementation of Optik, our local ray tracer, the rendering of a scene with every surface anisotropic shows an increase on average of 25% over the same scene without anisotropy, assuming ray casting. As can be seen for Images III and IV, the ratio can be much lower for complex scenes.

5 Images

On Image Ia and Ib, twelve spheres with different levels of anisotropy are rendered. A directional light source is used and all the spheres are projected orthogonally onto the screen in order to keep a constant viewing direction. The anisotropy over the spheres is defined by longitudinal scratches; on some spheres the position of one of the two poles is clearly visible as a black region.

Image Ia shows spheres with the anisotropy simulated by positive cylinders. From left to right and top to bottom, the first sphere has a $d = 0.0001$ and $h = 0$. This is practically undistinguishable from an isotropic reflection model. On the next sphere, d is increased to 0.5; the highlight starts forming a ring. The highest anisotropy produced by the cylindrical model is shown on the next sphere where $d = 2.0$ while h is still 0. Note how the highlight ring is clearly defined and how dark regions appear, due to the shadowing between cylinders. On the next sphere (second row), d is raised to 5.0, leaving a floor of length 3.0 between the cylinders. Note how the isotropic highlight reappears due to this floor region while the highlight ring becomes dimmer because the cylinders occupy a smaller region (proportionally to the floor). The floor is raised to $h = 0.86$ on the next sphere, and most of the dark regions disappear. Finally, the cylinders are spaced by only 2.0 with the same $h = 0.86$. The shadowing effect, while less pronounced on this last sphere, still makes the highlight ring dimmer.

On Image Ib, negative cylinders are used, with the same values for d and h. Most of the spheres are quite similar to the ones with positive cylinders, but notice how the self-shadowing effect is more pronounced on the third sphere (with $d = 2$ and $h = 0$).

The anisotropy on a surface also influences the way the environment is reflected onto it. Image II shows four views of the same highly reflective sphere above the unavoidable checkerboard. The surface is not specular so that highlights do not complicate the picture. A directional light source points down on the sphere. On the top left corner, we used a perfect mirror-like reflection. By adding high frequency bumps onto the surface, we get the image on the top right corner; it just adds some noise to the reflection. On the bottom left corner, the surface is given longitudinal anisotropy, and on the bottom right corner the surface exhibits latitudinal anisotropy. Note how the red squares are stretched horizontally on the left while on the right the reflected squares are stretched vertically.

Images III and IV show some objects which we normally associate with anisotropy. In Image III, the objects have been created by sweeping a spline curve, and the triangles produced have their tangents approximated via the same technique than used for the surface normals. In Image IV, the sound system has many anisotropic knobs abd other parts, The hanging cloth on the right is modelled as an isotropic surface. The motifs on the cube in the foreground were created using a 3D texture which provides the various orientations for the surface tangents.

6 Conclusion

We presented an anisotropic reflection model based on covering the surface with groups of microscopic cylinders. Anisotropy is caused by the directional distribution of the normals along the cylinders. Multidirectional anisotropy can be achieved by defining groups of cylinders oriented in specific directions and given specific weights. Two parameters, the distance between cylinders and the height of the floor, provide qualitative and quantitative control over the anisotropy. These parameters can easily be tuned to simulate a specific surface or to produce the desired effect. Slightly different effects can also be produced by "negative" cylinders.

The model takes into account hiding and shadowing between cylinders. The reflected intensity can be determined from regular samples (as seen by the viewer) along the cylinders surface which are then averaged. The sampling angles can also be used to obtain the colour reflected/refracted to the viewer via the anisotropic surface. An analytic solution to the diffuse intensity over the cylinders was developed, as well as an approximation to the specular intensity using Chebyshev polynomials. The Chebyshev approximation can also be used to solve similar problems encountered by Saito et al. [21] for the reflection off rounded edges and Kajiya and Kay [15] for the rendering of fur. It has also been used to solve the specular reflection off a surface lit by a linear light source [19].

There are several related problems yet to solve. One such problem concerns the trade-off between the accuracy of the Chebyshev approximation and the degree of the resulting polynomial. It might be possible to reduce the degree of the polynomial with a more careful study of the maximal difference between the curves such that, for a given display system, no observer could distinguish between the two resulting intensities.

The Chebyshev approximation can also be used to estimate the inter-reflection between cylinders. In the current model the inter-reflection is ignored, but it would be interesting to see if inter-reflection improves the realism of the model enough to justify the extra processing.

The authors also believe this model can be used to provide an additional tool for visualizing information on data like orientation of simple vector fields. However in order to obtain understandable results, some motion must be applied to the surface or to the light source. At the current stage of our implementation we are quite far from real-time motion.

7 Acknowledgement

The authors would like to thank Avi Naiman, Andrew Woo, Mikio Shinya, Peter Cahoon and especially John Buchanan for their help. We also thank Frederic Taillefer for the drapery model used in Image IV. The authors are grateful to NSERC, which provided an operating Grant and a Graduate Scholarship, to Ontario ITRC for their financial support, and to the generous support of the University of British Columbia where the final phase of this work took place.

Appendix

A Model of Negative Cylinders

A negative cylinder is a cylinder subtracted from the surface. The mathematics involved with this type of cylinder are simpler because a negative cylinder cannot be hidden from the viewer by another negative cylinder. Thus, only the self-shadowing and self-hiding have to be considered. The same parameters d and h are used to control the anisotropy, however the floor is pushed downwards by a distance h instead upwards from the surface in the case of positive cylinders.

For negative cylinders, the angles are measured from the surface plane and they represent the opposite of what they did for the positive cylinders[5] (see figure 7, 8 and 9). The

[5] The main reason of this redefinition is to generate simpler mathematical expressions.

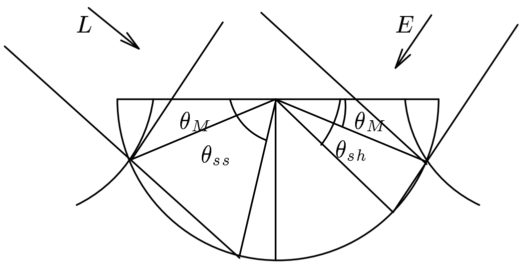

Figure 7: Some Variables for the Negative Cylinders

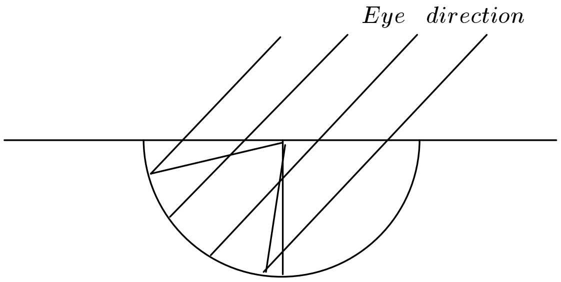

Figure 8: Regular Sampling from the Viewer

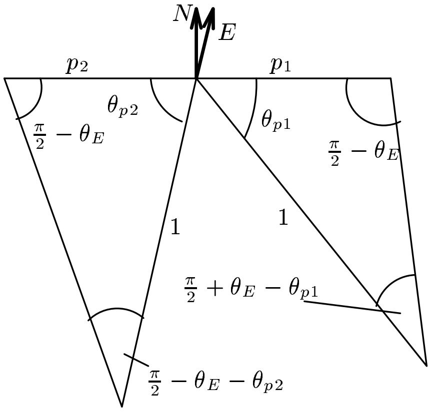

Figure 9: Finding the Sampling Angles θ_p

variation of the normals is then within $[\theta_M, \pi - \theta_M]$, where θ_M is given by

$$\theta_M = \min(\phi_d, \phi_h)$$

and $\phi_h = \sin^{-1}(h)$

$$\phi_d = \begin{cases} \cos^{-1}(d/2) & \text{if } d < 2 \\ 0 & \text{otherwise} \end{cases}$$

The angles at which the self-shadowing θ_{ss} and the self-hiding θ_{sh} stop are then respectively

$$\begin{aligned} \theta_{ss} &= \max(2\theta_L - \theta_M, \theta_M) \\ \theta_{sh} &= \max(2\theta_E - \theta_M, \theta_M) \end{aligned}$$

Simplifications also occur because the floor is always completely visible and illuminated. The sampling angles are computed as

$$\begin{aligned} \theta_{p_1} &= \frac{\pi}{2} + \theta_E - \sin^{-1}(p_1 \cos\theta_E) \\ \theta_{p_2} &= \frac{\pi}{2} - \theta_E - \sin^{-1}(p_1 \cos\theta_E) \end{aligned}$$

It is worthwhile to note that the analytic solution for the diffuse term as well as the Chebyshev approximation for the specular term are the same as for positive cylinders.

B Chebyshev Approximation

Pseudo-code is provided to compute the Chebyshev polynomials approximating the function $\cos^n \phi$, where $\phi \in [0, x_M]$. The array $a[0..D(a)]$ will contain the coefficients of the polynomial approximating the curve. For more information on the properties of the Chebyshev polynomials for approximation, consult [20].

$D(a)$ is the degree of the polynomial approximation

```
Approximate (n, D(a), a[])
x_M : cos^-1 0.0039^(1/n)
for i = 0 to D(a)
    sum = 0;
    for k = 0 to D(a)
        x_cheby = cos(kπ/D(a))
        x = x_M × (x_cheby + 1)/2
        if (k = 0 or k = D(a))
            sum = sum + 0.5 cos^n(x)× Cheby(i, x_cheby)
        else
            sum = sum + cos^n(x)× Cheby(i, x_cheby)
        endfor
    b[i] = 2/D(a) × sum
    endfor
for i = 0 to D(a)
    if (i = 0 or i = D(a))
        a[i] = 0.5b[i]× Cheby(i, x_cheby)
    else
        a[i] = b[i]× Cheby(i, x_cheby)
    endfor

Cheby (degree, x)
    if (degree = 0) return (1)
    if (degree = 1) return (x)
    return (2x Cheby(degree − 1, x) − Cheby(degree − 2, x))
```

References

[1] Bass, F. G. and Fuks, I. M., *Wave Scattering from Statistically Rough Surfaces.* Pergamon Press Ltd., 1979.

[2] Beckmann, P., Spizzichino A., *The scattering of electromagnetic waves from rough surfaces.* Artech House, Inc., 2^{nd} ed., 1987.

[3] Blinn, J., Models of Light Reflection for Computer Synthesized Pictures. Proceedings of SIGGRAPH'77, In *Computer Graphics* 11, 2 (July 1977), 192-198.

[4] Blinn, J., Simulation of wrinkled surfaces. Proceedings of SIGGRAPH'78, In *Computer Graphics* 12, 3 (August 1978), 286-292.

[5] Born, M. and Wolf, E., *Principles of Optics*. Pergamon, Oxford, 1975.

[6] Brennan, B., Bandeen, W., Anisotropic Reflectance Characteristics of Natural Earth Surfaces. *Applied Optics* 9, 2, (February 1970).

[7] Cabral, B., Max, N., Springmeyer, R., Bidirectional Reflection Functions from Surface Bump Maps. Proceedings of SIGGRAPH'87, In *Computer Graphics* 21, 4 (July 1987), 273-281.

[8] Cook, R., Torrance, K., A Reflectance Model for Computer Graphics. Proceedings of SIGGRAPH'81, In *Computer Graphics* 15, 3 (August 1981), 307-316.

[9] Fournier, A., Fiume, E., Ouellette, M. and Chee, C., FIAT: Light Driven Global Illumination. DGP Technical Memo DGP89-1, Dynamic Graphics Project, University of Toronto, 1989.

[10] Gouraud, H., Continuous Shading of Curved Surfaces. *IEEE Transactions on Computers* 20, 6 (June 1971), 623-628.

[11] Heckbert, P., Survey of Texture Mapping. *IEEE Computer Graphics and Applications*, 6, 11 (November 1986), 56-67.

[12] Immel, D., Cohen, M. and Greenberg, D., A Radiosity Method for Non-Diffuse Environments. Proceedings of SIGGRAPH'86, In *Computer Graphics* 20, 4 (August 1986), 143-150.

[13] Kajiya, J., Anisotropic Reflection Models. Proceedings of SIGGRAPH'85, In *Computer Graphics* 19, 3 (July 1985), 15-21.

[14] Kajiya, J., The Rendering Equation. Proceedings of SIGGRAPH'86, In *Computer Graphics* 20, 4 (August 1986), 143-150.

[15] Kajiya, J. and Kay, T., Rendering Fur with Three Dimensional Textures. Proceedings of SIGGRAPH'89, In *Computer Graphics* 23, 3 (July 1989), 271-280.

[16] Max, N., Horizon mapping: shadows for bump-mapped surfaces. *The Visual Computer*, Vol. 4, 1988, 109-117.

[17] Miller, G., From Wire-Frames to Furry Animals, Proceedings of *Graphics Interface '88*, 1988, 138-145.

[18] Phong, B., Illumination for Computer Generated Pictures. *Communications of the ACM* 18, 6, (June 1975), 311-317.

[19] Poulin, P. and Amanatides, J., Shading and Shadowing with Linear Light Sources. Proceedings of *Eurographics 90*, (September 1990).

[20] Ralston, A., *A First Course in Numerical Analysis.* McGraw-Hill, 1965.

[21] Saito, T., Shinya, M., Takahashi, T., "Highlighting Rounded Edges", Proceedings of CG International 89. In *New Advances in Computer Graphics* 1989, 613-629.

[22] Sillion, F. and Puech, C., A General Two-Pass Method Integrating Specular and Diffuse Reflection. Proceedings of SIGGRAPH'89, In *Computer Graphics* 23, 3 (July 1989), 335-344.

[23] Torrance, K., Sparrow, E., Polarization, Directional Distribution, and Off-Specular Peak Phenomena in Light Reflected from Roughened Surfaces. *J.Opt.Soc.Am.* 56, 7, 1966.

[24] Torrance, K., Sparrow, E., Theory for Off-Specular Reflection from Roughened Surfaces. *J.Opt.Soc.Am.* 57, 9, 1967, 1105-1114.

[25] Trowbridge, T., Reitz, K., Average Irregularity Representation of a Roughened Surfaces for Ray Reflection. *J.Opt.Soc.Am.* 65, 3, 1967.

[26] Watanabe, Y., Suenaga, Y., Hair Animation in Backlight. Proceedings of CG International 89. In *New Advances in Computer Graphics* 1989, 691-700.

[27] Whitted, T., An Improved Illumination Model for Shaded Display. *Communications of the ACM* 23, 6, (June 1980), 343-349.

[28] Williams, L., Pyramidal Parametrics. Proceedings of SIGGRAPH'83, In *Computer Graphics* 17, 3 (July 1983), 1-11.

[29] Yokoi, S., Toriwaki, J., *Computer Science and Technologies, Ch. 3.7 Realistic Expression of Solids with Feeling of Materials*, JARECT, Vol. 18, Ohmsha Ltd., 1988.