The Pyramidic Microfacet BRDF

Rendering pyramidically textured photovoltaics

Paul E.A. Adriaanse

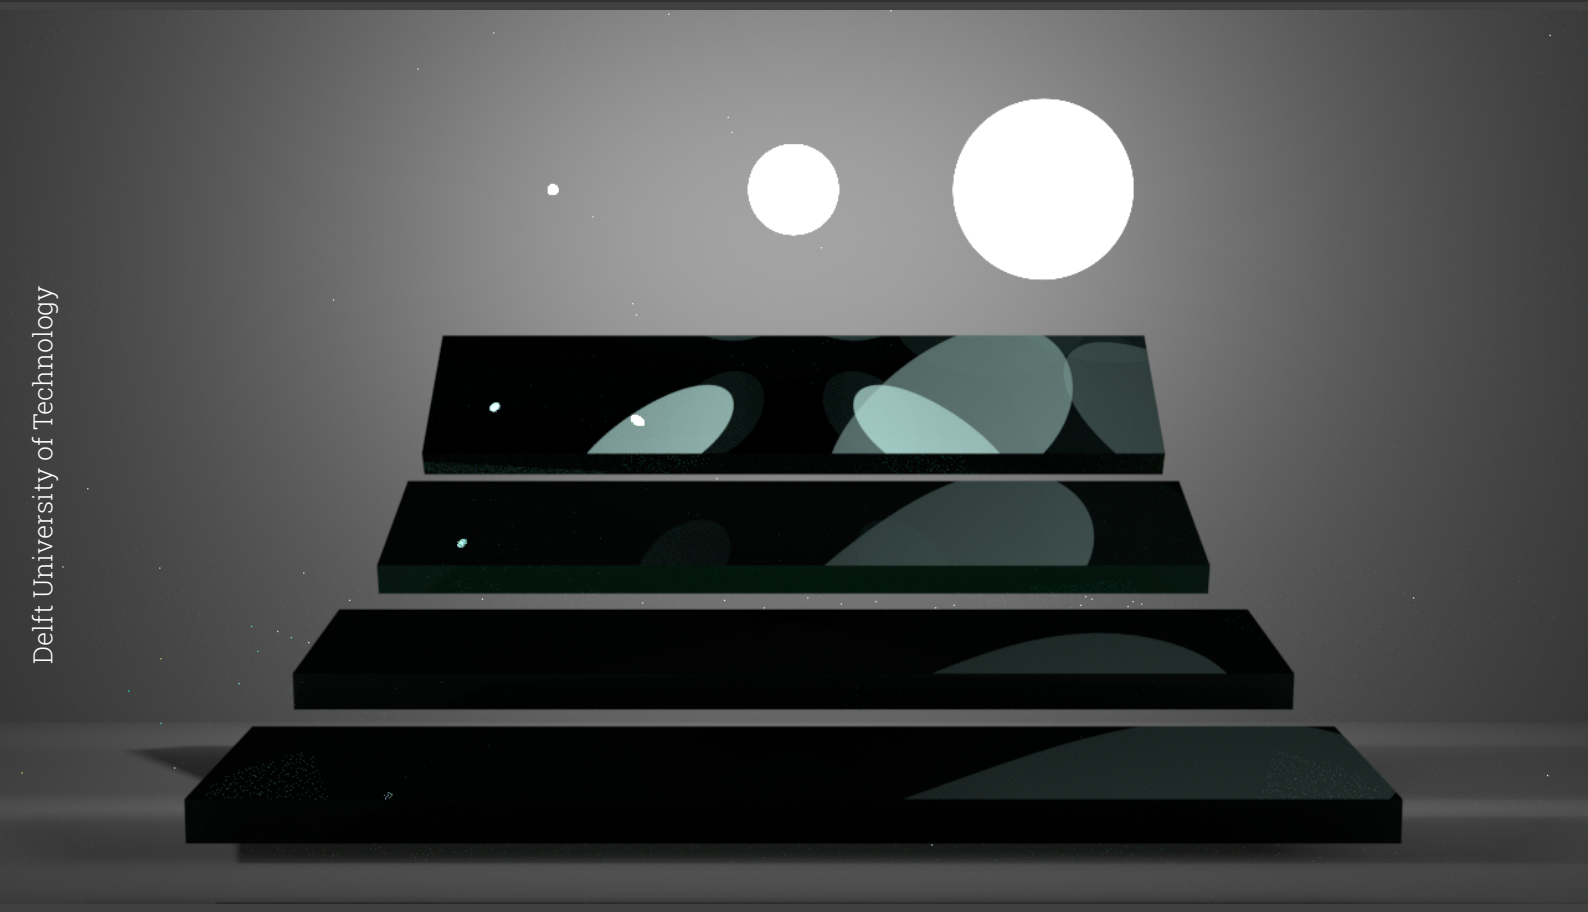

Delft University of Technology

TU Delft

The Pyramidic Microfacet BRDF

Rendering pyramidically textured photovoltaics

by

Paul E.A. Adriaanse

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Friday July 4, 2025 at 11:00 AM.

Student number:	4721454	
Thesis committee:	Prof. Ricardo Marroquim,	Thesis Advisor
	Ir. Mark van de Ruit,	Daily Co-Supervisor
	Prof. Qing Wang,	Committee Chair
	Dr. Rudi Santbergen,	Committee Member

Cover: Modified version of 'MIS' scene by Bitterli [1], based on original by Veach [2], rendered using the coated pyramidic BRDF.

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.
The implementation is available at `https://github.com/pea-adriaanse/Thesis-PyramidicBRDF`.

TU Delft

Abstract

To incentivize solar panel adaptation, a multitude of solutions are actively being researched. Among these solutions are colored solar cells, using coatings or filters to open the door for more architectural expression. In order to help architects explore these possibilities, it is necessary to enable flexible, fast, and realistic visualization. This work explores the modeling of color-coated photovoltaic cells using a physically based bidirectional reflectance distribution function (BRDF), with a focus on the pyramidically textured structures embedded inside many of these cells. The BRDF is analytically derived, modeling light interactions as recursively specular multiple-scattering. The model is parametric, characterizing the surface using its pyramid density and pyramid slant angle, making it generally applicable to homogeneous pyramidic surfaces with uniform pyramid heights. Evaluation indicates the model closely approximates the behavior of generated references at steep viewing angles. Clear avenues of improvement, including corrections at shallow angles, are discussed within the context of future work.

Contents

Nomenclature

Notation

Notation	Definition
$\vec{v}$	Vector
$\hat{v}$	Normalized/Unit vector
$\underline{M}$	Matrix
$\|C\|$	Number of elements in set C (cardinality of C)

Abbreviations

Abbreviation	Definition
BRDF	Bidirectional Reflection Distribution Function
NDF	Normal Distribution Function
PDF	Probability Density Function
CDF	Cumulative Distribution Function

Symbols

Dimensionless units are denoted using [-]. Symbols with no quantity/units are denoted not applicable (N/A).

Symbol	Definition	Unit
ρ	Pyramid density	[*pyramids* $/\mu m^2$]
α	Slant angle	[$^\circ$]
θ	Polar angle	[rad]
ϕ	Azimuth	[rad]
z	Depth	[μm]
A	Area	[μm^2]
$w(z)$	Pyramid width	[μm]
$A_B(z)$	Pyramid base surface area	[μm^2]
C	Set of pyramids on surface	N/A
A_T	Total surface area of reference plane region defining the pyramidic surface	[μm^2]
$P_C(A)$	Pyramid spawning probability	[-]
$P_0(A)$	Pyramid non-spawning probability	[-]
$P_B(z)$	Pyramid burying probability (surface height CDF)	[-]
$P_E(z)$	Pyramid non-burying probability	[-]
$E(z)$	Surface height PDF	[-]

Symbol	Definition	Unit
E, N, W, S	Pyramid faces	N/A
$D_m(\hat{\omega})$	Microfacet normal distribution function	$[1/sr]$[1]
$D_n(\hat{\omega})$	Visible normal distribution	$[1/sr]$
$D(\hat{\omega})$	Sum of dot-products between pyramid faces and viewing direction $\hat{\omega}$	[-]
$P_S(A)$	Shadowing probability	[-]
$G_1(\hat{\omega})$	Shadowing function	[-]
$G_L(\theta)$	Shadowing function derived by de Haas [4]	[-]
$P_{BS}(\hat{\omega})$	Back-scattering probability	[-]
R	Reflectance	[-]
λ	Wavelength	$[nm]$
Ω	Upper hemisphere	N/A
$\delta(x)$	Dirac delta function	[2]

[1] The units of the NDF and its interpretation, also defined as $[m^2/sr]$ [3], merit discussion beyond the scope of this work.
[2] The unit of the dirac delta function is the inverse of that of its argument.

1 Introduction

Globally, solar power is the fastest growing individual source of renewable energy [5]. However, it remains outpaced by fossil fuels in absolute growth terms in 2024 [6]. Renewable energy adaptation thus remains a steppingstone to achieving climate goals. Besides improved solar power efficiency, factors such as improved aesthetics can contribute to increased adaptation. Several improvement approaches exist, such as the incorporation of colored coatings in or on the cell. Alongside incorporation of solar cells into roof tiles, facades and windows, such colored cells enable architects to express novel designs. Here, architectural or predictive rendering is an important tool in the design stage.

A common approach to solar cell design involves a surface texture composed of microscopic pyramids. These serve to strategically increase the amount of light captured by the surface. In rendering, surface materials are described by the bidirectional reflectance distribution function (BRDF). However, research into BRDFs mimicking such pyramidical surfaces is limited in applicability or render performance [7, 8].

This work introduces a statistical model of the microscopic surface textures commonly used in solar cell design, and from there analytically derives a functional BRDF for rendering purposes. Subsequently, an optical filter coating is introduced to reproduce common solar cell appearances.

First, some background and related work is introduced (Chapter 2). Following, the textured surface model is discussed and statistical properties are derived (Chapter 3). Subsequently, the pyramidic BRDF is described (Chapter 4). Afterwards, relevant components are evaluated (Chapter 5). Finally, a conclusion is provided (Chapter 6), including discussion about future work.

2

Background and Related Work

In the following, solar cell texturing and coloring are shortly introduced alongside related work (Section 2.1). Subsequently, a description of microfacet BRDFs and related work is provided (Section 2.2).

2.1. Solar Cells

Surface Texturing Surface texturing is a common technique utilized in the optimization of solar cell efficiency. For this, numerous methods exist, with square pyramids the focus of this work. A large body of work investigates these pyramidically textured solar cells, focusing on manufacture [9, 10, 11] or efficiency optimization [12, 13, 14, 15].

Alternatives, however, such as triangular pyramids and V-grooves exist [16, 17]. Nano-scale structures employing the moth-eye effect [18] and other biologically inspired designs have also been proposed [19], as have texture combinations, such as nano-scale whiskers deposited onto macro-scale pyramids, improving infrared absorption [20]. For a wider overview, refer to Hussein and Ismael [21].

The pyramidically textured cells under consideration are spatially homogeneous, with randomly distributed aligned pyramids [9, 10]. The pyramid heights on any particular surface may vary, with different manufacturing approaches yielding differing pyramid height distributions [12, 13]. Generated examples of these surfaces are shown in Figure 2.1, with varying pyramid heights provided by Procel *et al.* [15]. For consistency, constant values for the pyramid slant angle $\alpha = 54.7°$ and density $\rho = 0.6$ *pyramids*$/\mu m^2$ are used throughout this work, also following Procel *et al.* [15].

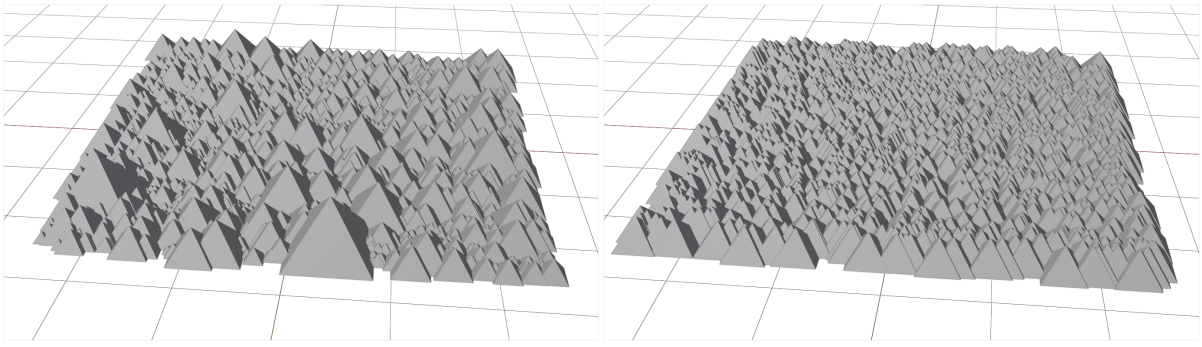

(a) Varying pyramid heights, provided by Procel *et al.* [15]. **(b)** Constant pyramid heights.

Figure 2.1: Generated pyramid texture examples, with slant angle $\alpha = 54.7°$ and pyramid density $\rho = 0.6$ *pyramids*$/\mu m^2$.

The shape of the square pyramids (Figure 2.2) is fully characterized by their height h and slant angle α.

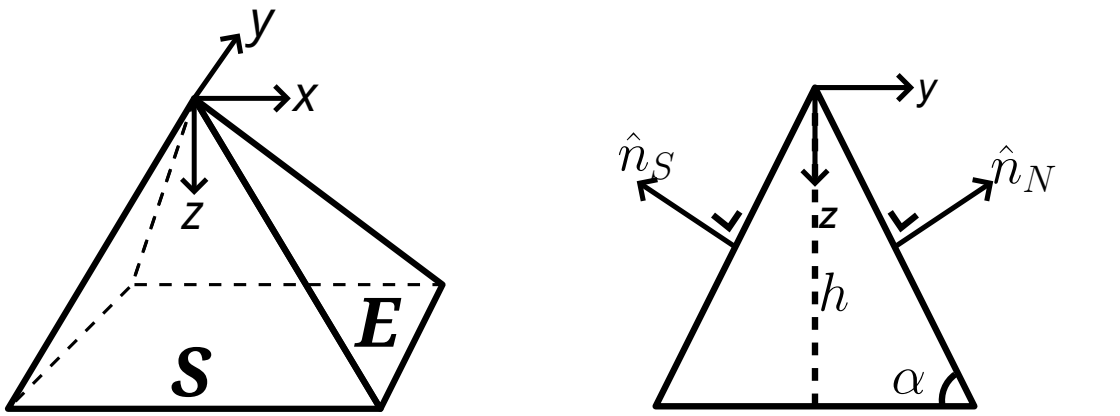

(a) 3D view of pyramid. East and south faces are labeled 'E' and 'S'.

(b) Side view of eastern face, with pyramid height h, slant angle α and face normals $\hat{n}_N$ and $\hat{n}_S$.

Figure 2.2: Local coordinate system of pyramid, with its peak centered at the origin and its base aligned to the x- and y-axes.

Coloring Photovoltaic cell coloring can be achieved in several ways. Coated glass layers can be added [22], underlying surfaces can be colored and even luminescence has been investigated. Block *et al.* [23] provide an overview of different approaches. This work focuses on interference filters (Section 4.7), applied as layered coatings above the pyramidic texture [24].

2.2. Microfacet BRDFs

BRDFs $f(\hat{\omega}_o, \hat{\omega}_i)$ describe the ratio between the reflected radiance (in direction $\hat{\omega}_o$) and irradiance (from direction $\hat{\omega}_i$) of a material. Microfacet BRDFs model the surface of this material using many flat microfacets. Specular microfacet BRDFs are typically described in the form of Equation 2.1 [3, 25]. It contains Fresnel, shadowing-masking function and microfacet normal distribution terms (F, G_2 and D), alongside the microfacet normal $\hat{n}_m$ and global surface normal $\hat{n}_g$, corresponding to $-\hat{z}$ in this work.

$$f(\hat{\omega}_o, \hat{\omega}_i) = \frac{F(\hat{\omega}_i, \hat{n}_m) G_2(\hat{\omega}_o, \hat{\omega}_i, \hat{n}_m) D_m(\hat{n}_m)}{4|\hat{n}_g \cdot \hat{\omega}_i||\hat{n}_g \cdot \hat{\omega}_o|} \tag{2.1}$$

This description only considers the first reflection. This work, however, examines multiple reflections. Additionally, as material properties are out of scope in this work, Fresnel will not be discussed.

Shadowing Function Also referred to as the masking function, the shadowing function $G_1(\hat{\omega}_o, \hat{n}_m)$ is defined by Heitz [3] as the visible fraction of microfacets with normal $\hat{n}_m$ from direction $\hat{\omega}_o$. It is discussed in detail in Section 4.2. The shadowing-masking function $G_2(\hat{\omega}_o, \hat{\omega}_i, \hat{n}_m)$ is defined as the fraction of microfacets visible from both direction $\hat{\omega}_i$ and $\hat{\omega}_o$, often defined as the product $G_1(\hat{\omega}_o, \hat{n}_m) * G_1(\hat{\omega}_i, \hat{n}_m)$, assuming shadowing/masking from both directions is independent [3].

Normal Distribution Function In creating microfacet BRDFs, a start is typically made in defining the normal distribution function (NDF) $D_m(\hat{\omega})$. Described by Heitz [3] as a conversion tool between "spacial and statistical integrals", the NDF effectively yields the (differential) size of the microfacets with normal $\hat{\omega}$, relative to the ground plane surface area they cover. Although the explicit use of the NDF is not necessary for this work, its addition (Section 4.1) allows for the application of theory from previous work.

Visible Normal Sampling Normal sampling is a way of using a microfacet's scale from the perspective of a viewer, weighted by its visibility, to determine what microfacet is sampled. This can be derived directly but can also be determined using normal distribution and shadowing functions (Section 4.3, [3, 26]).

Prior Work The only prior work on pyramidic microfacets in a rendering context is that of de Haas [4]. The authors introduce a simplified model, deriving an azimuth-agnostic shadowing function, but do not apply this for rendering. Preliminary ray-traced measurements of generated surfaces, however, indicate the azimuth can introduce a shadowing difference of up to approximately 8.5% (Figure 2.3), when utilizing pyramid heights as provided by Procel *et al.* [15].

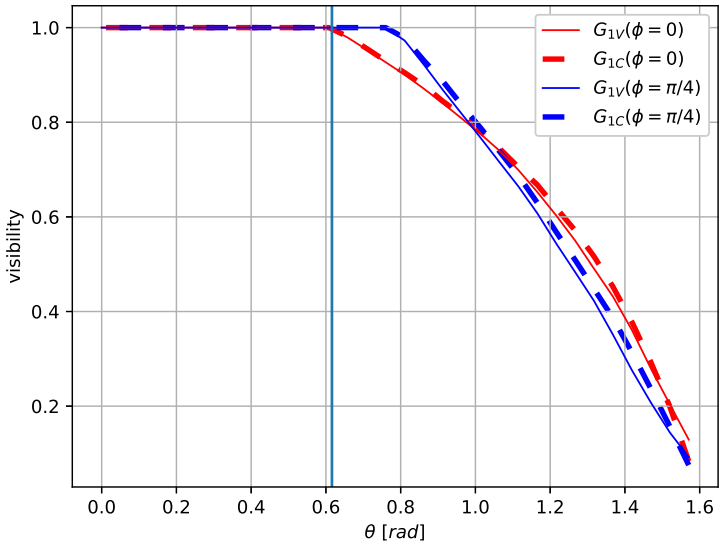

Figure 2.3: Ray-traced shadowing function measurements of generated textured surfaces with constant, G_{1C}, and variable pyramid heights, G_{1V}, at azimuth angles $\phi = 0$ and $\phi = \pi/4$, using 8000 samples per data-point. Pyramid density $\alpha = 0.6$ *pyramids*$/\mu m^2$ and slant angle $\alpha = 54.7°$ were used. The vertical line denotes the critical viewing angle.

In contrast, v-grooves are extensively applied [27, 28, 29, 30]. Scattering behavior for v-grooves is further studied outside rendering w.r.t. solar cell design [31]. Other research on solar cell rendering does not consider pyramidic microsurfaces or coated layers [7, 8].

Since the color coating referred to in this work employs layers, layered BRDFs also necessitate brief mentioning. As, although simplified and modeled as a single layer on top of the pyramidic BRDF in this work, alternatives for modeling several layers individually exist. An extensive overview of layered BRDF models is provided by Montes and Ureña [32]. For more on the topic of microfacet BRDFs, the reader is referred to Pharr *et al.* [33].

3

Pyramidic Surface Model

In order to derive a BRDF for the pyramidic surface (Section 2.1), a simplified model is required. Additionally, a description of both pyramid distribution and surface height distribution is necessary.

In this chapter, the model is first described and summarized (Section 3.1). Subsequently, the actual placement of pyramids on the surface is examined (Section 3.2). Then, the resulting height distribution describing the surface is derived (Section 3.3). Finally, surface gaps are discussed as a potential pitfall in the implementation and evaluation of finite depth surfaces (Section 3.4).

3.1. Model

In order to model the pyramid surface practically, several simplifications need to be performed. Firstly, the surface is modeled infinite in its extent. Likewise, (maximum) surface depth modeled as infinite. This approximation closely matches finite surfaces, assuming it is large compared to individual pyramids, with a maximum depth large enough to minimize surface gaps (Section 3.4), as all metrics in the rest of this chapter approach their minima/maxima at a relatively small scale.

Secondly, as the surface is homogeneous with randomly placed pyramids (Section 2.1), their distribution is modeled as a continuous uniform distribution (Section 3.2).

Finally, pyramid heights are taken to be constant, as this reduces complexity, while there is no universal pyramid height distribution, as these vary depending on the manufacturing process [10]. Interestingly however, although beyond the scope of this work, preliminary ray-traced measurements (Figure 2.3) indicate a constant pyramid height surface's shadowing function can closely approximates a variable height surface's shadowing function (using pyramid heights as provided by Procel *et al.* [15]).

Given this constant pyramid height model with infinite depth, the global coordinate system can be defined such that all pyramid peaks are placed on the reference plane $z = 0$ (Figure 3.1), with sides aligned to the x- and y-axes, equivalent to a shifted version of the pyramids' local coordinate systems (Figure 2.2). The ground plane is defined as the bottom of the material perpendicular to the reference plane. Although placed at $z = \infty$, it aids intuition.

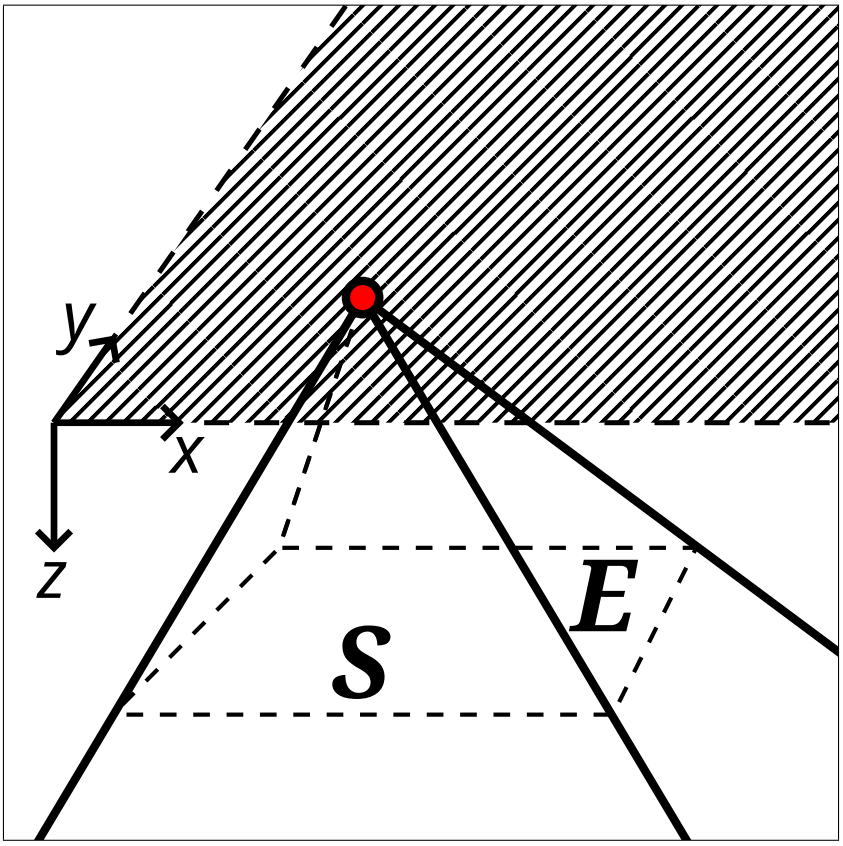

Figure 3.1: Global coordinate system, with infinite-depth axis-aligned pyramids, with peaks (red) on the reference plane $z = 0$.

Summary

Given this model, pyramid shape is dictated by slant angle α, while pyramid distribution is characterized by pyramid density ρ. The model is thus fully defined by α and ρ and has the following properties:

- The surface is infinite in extent and depth.
- The surface consists of identically shaped axis-aligned square pyramids.
- Pyramids are placed according to a continuous uniform distribution.
- Pyramid peaks reside on the reference plane ($z = 0$).

For consistency, constant slant angle $\alpha = 54.7°$ and pyramid density 0.6 *pyramids*$/\mu m^2$ are used throughout this work, following Procel *et al.* [15].

3.2. Pyramid Distribution

In this section the pyramid distribution is examined by treating the construction of the surface as a stochastic process. In particular, a relation is established between an arbitrary surface patch and number of pyramid placed within its extent.

Given pyramids are uniquely identified by their peaks, and all peaks reside on the reference plane, their distribution is reduced to a two dimensional problem. Consider the reference plane as a finite dimensional surface of area A_T upon which a set of pyramid peaks C is distributed according to a continuous uniform distribution (Figure 3.2). The number of peaks is then $|C| = \rho A_T$. Note the pyramid density ρ should be seen as a global average, as the local density may deviate due to the stochastic nature of the surface.

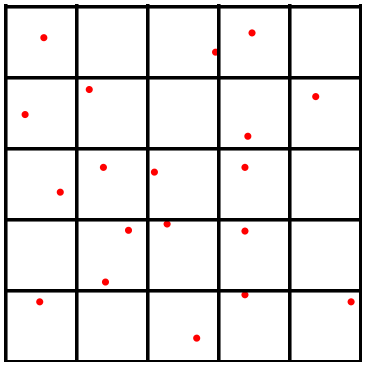

Figure 3.2: Points distributed over a two dimensional grid.

Given any particular peak, the probability $P_1(A, A_T)$ of finding this peak inside area $A \leq A_T$ is described by its ratio (Equation 3.1). The pyramid spawning probability $P_C(A, A_T)$ of finding any peak(s) from set C in area A then follows as the negation of the probability that no peaks are found in A (Equation 3.2).

$$P_1(A, A_T) = \frac{A}{A_T} \tag{3.1}$$

$$P_C(A, A_T) = 1 - (1 - P_1(A, A_T))^{|C|} = 1 - (1 - \frac{A}{A_T})^{\rho A_T} \tag{3.2}$$

Given $A_T \to \infty$, as the surface is modeled infinite in extent, and using the limit in Equation 3.3, as derived in Appendix A, the pyramid spawning probability $P_C(A, A_T)$ resolves to $P_C(A)$ (Equation 3.4, Figure 3.3), with non-spawning probability $P_0(A)$ being its complement (Equation 3.5).

$$\lim_{x \to \infty} \left(1 + \frac{k}{x}\right)^{mx} = e^{km} \tag{3.3}$$

$$P_C(A, A_T) \approx \lim_{A_T \to \infty} P_C(A, A_T) = P_C(A)$$

$$P_C(A) = 1 - e^{-\rho A} \tag{3.4}$$

$$P_0(A) = 1 - P_C(A) = e^{-\rho A} \tag{3.5}$$

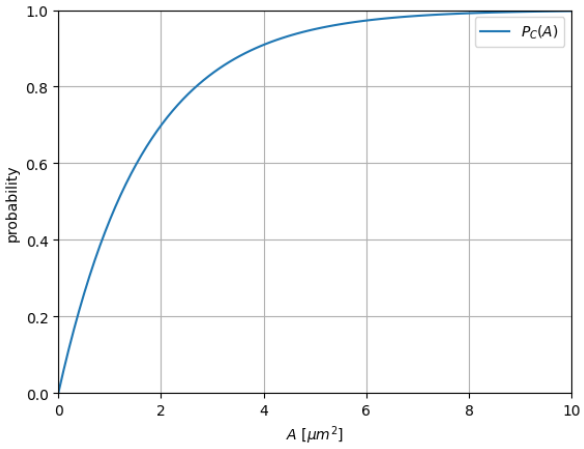

Figure 3.3: Pyramid spawning probability P_C as a function of area.

In practice, the error in this approximation (Equation 3.6) tends to zero as A_T increases, being negligible for practical values of ρ and A_T.

$$|P_C(A, A_T) - P_C(A)| = e^{-\rho A} - (1 - \frac{A}{A_T})^{\rho A_T} \tag{3.6}$$

3.3. Surface Height Distribution

Given the description of the pyramid distribution discussed above, the surface height distribution can be described using the probability density function (PDF) $E(z)$ and its corresponding cumulative distribution function (CDF) $P_B(z)$.

As a CDF, $P_B(z)$ describes the probability some point $\vec{p}$ at depth z is buried by the surface above, with as complement the non-burying probability $P_E(z)$. The non-burying probability is equivalent to the probability no pyramid spawns around $\vec{p}$ with it's base $A_B(z)$ overlapping at depth p_z. This is given by the pyramid non-spawning probability function $P_0(A_B(z))$ (Section 3.2). Here $w(z)$ describes the width of pyramid faces at depth z.

$$w(z) = \frac{2z}{\tan(\alpha)} \tag{3.7}$$

$$A_B(z) = w(z)^2 = \left(\frac{2z}{tan(\alpha)}\right)^2 \tag{3.8}$$

$$P_E(z) = P_0(A_B(z)) = e^{-\rho A_B(z)} = e^{-\rho 4z^2/\tan^2(\alpha)} \tag{3.9}$$

$$P_B(z) = 1 - P_E(z) \tag{3.10}$$

As the PDF, $E(z)$ follows as the derivative of the corresponding CDF $P_B(z)$ and should integrate to 1. As it describes the probability density of points on the surface occurring at depth z, it also follows as the weighted ratio of non-buried points on the surface of a pyramid.

$$E(z) = \frac{d(P_B(z))}{dz} = \frac{8\rho z}{\tan^2(\alpha)} e^{-\rho 4z^2/\tan^2(\alpha)} \tag{3.11}$$

$$E(z) = \frac{P_E(z)w(z)}{\int_0^\infty P_E(z)w(z)dz} \tag{3.12}$$

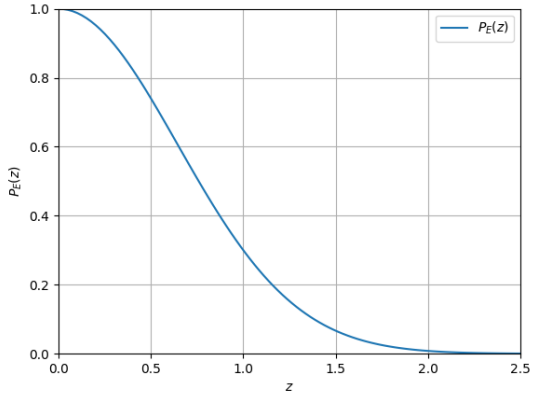

Figure 3.4: Non-burying probability as a function of depth.

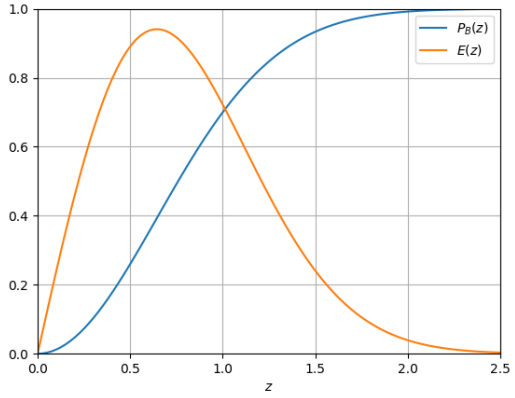

Figure 3.5: Surface height distribution, as described by its PDF $E(z)$ and CDF $P_B(z)$.

3.3.1. Plausibility Check

The following plausibility check can be performed using non-burying probability $P_E(z)$. The predicted ground plane surface area covered by the pyramid's surface, A'_P, should equal the actual average ground plane surface area per pyramid A_P (Equation 3.13), or equivalently, the predicted surface area of each face, A'_{AF}, should equal actual average A_{AF} (Equation 3.14).

$$A_P = \frac{A_T}{|C|} = \frac{A_T}{\rho A_T} = \frac{1}{\rho} \tag{3.13}$$

$$A_P = 4\cos(\alpha)A_{AF}$$

$$A_{AF} = \frac{1}{4\cos(\alpha)\rho} \tag{3.14}$$

The predicted face area A'_{AF} is described by the surface integral in Equation 3.16, where F is the surface of one of the pyramid's faces and z_F is the horizontal projection of z onto a face's surface. Since $P_E(z)$ is independent of x, this can be simplified to a single integral (Equation 3.17).

$$z_F = z/\sin(\alpha) \tag{3.15}$$

$$A'_{AF} = \iint_F P_E(z(z_F))dA \tag{3.16}$$

$$= \int_0^\infty P_E(z)w(z)\frac{dz_F(z)}{dz}dz \tag{3.17}$$

Reintroducing the $4\cos(\alpha)$ factor, the equality is then proven.

$$\begin{aligned} A'_P &= 4\cos(\alpha)\int_0^\infty P_E(z)w(z)\frac{dz_F}{dz}dz \\ &= 4\cos(\alpha)\int_0^\infty e^{-\rho 4z^2/\tan^2(\alpha)}\frac{2z}{\tan(\alpha)}\frac{1}{\sin(\alpha)}dz \\ &= \frac{8}{\tan^2(\alpha)}\int_0^\infty ze^{-\rho 4z^2/\tan^2(\alpha)}dz \\ &= \frac{8}{\tan^2(\alpha)}\frac{-\tan^2(\alpha)}{\rho 8}\left[e^{-\rho 4z^2/\tan^2(\alpha)}\right]_0^\infty \\ &= \frac{-1}{\rho}(0-1) = \frac{1}{\rho} = A_P \end{aligned} \tag{3.18}$$

3.4. Surface Gap

One peculiarity resulting from the model used and its height distribution is the fact surfaces with finite height/depth h_S may contain gaps. When performing simulations using generated surfaces, this has to be taken into consideration. In particular, the non-burying probability $P_E(z)$ is non-zero at a finite maximum depth. Statistically, the ground plane at this maximum depth can thus contains regions not covered by pyramids.

The average gap per pyramid, $A_G(h_S)$, can be derived from the average ground plane surface area covered by a pyramid $A_P(z)$ (Equation 3.19, the finite version of Equation 3.18), with fractional gap $A'_G(h_S)$ defined with respect to the average ground plane surface area per pyramid A_P (Equation 3.13). Note this fractional gap is identical to the non-burying probability $P_E(h_S)$ (Equation 3.9, Figure 3.4).

$$A_P(h_S) = 4\cos(\alpha)\int_0^{h_S} P_E(z)w(z)\frac{dz_F}{dz}dz = \frac{1}{\rho}\left(1 - e^{-\rho 4{h_S}^2/\tan^2(\alpha)}\right) \tag{3.19}$$

$$A_G(h_S) = A_P - A_P(h_S) = \frac{1}{\rho}e^{-\rho 4{h_S}^2/\tan^2(\alpha)} \tag{3.20}$$

$$A'_G(h_S) = \frac{A_G(h_S)}{A_P} = e^{-\rho 4{h_S}^2/\tan^2(\alpha)} \tag{3.21}$$

Although there is thus no guarantee a finite surface is fully covered, this means the fractional gap decreases exponentially as maximum height/depth increases, with a fractional gap of only $2.28 * 10^{-7}$ at $h_S = 6\mu m$ for surfaces with parameters $\rho = 0.6$ ***pyramids***$/\mu m^2$ and $\alpha = 54.7°$.

4

Pyramidic Surface BRDF

Given the pyramidic surface description (Chapter 3), a BRDF can now be constructed. Specifically, a sequence of localized light-surface interactions (ray reflections) will be analytically approximated to find the BRDF. As a localized approximation, scattered rays enter and exit the represented material at the same position.

At a scattering, the BRDF $f(\omega_o, \omega_i)$ will be modeled using specular microfacets following a specific distribution (Section 4.1). A shadowing function (Section 4.2) modulates ray exit probabilities at every interaction. Given such specularity and exit probability, sampling is necessary to incorporate the BRDF in a renderer. This work relies on visible normal sampling to consider relevant samples (Section 4.3). As pyramidic microfacets facilitate multiple scattering, the BRDF incorporates this (Section 4.4). Due to approximations introduced in the general model, several problems become apparent. These are addressed in Section 4.5 and Section 4.6. Finally, in Section 4.7, a colored optic filter is incorporated into the BRDF, implemented such that it matches prior measurements [24].

4.1. Microfacet Normal Distribution Function

Although not strictly required for the derivations in this work, having the microfacet NDF $D_m(\hat{\omega})$ at hand simplifies comparisons with previous work. Given its definition as (Section 2.2), it follows directly from as the ratio between the average face surface area A_{AF} and average covered ground plane surface area per pyramid A_P (Subsection 3.3.1). As this ratio is identical for all pyramid faces, this results in a sum of Dirac delta functions:

$$D_m(\hat{\omega}) = \sum_{F' \in \{E,N,W,S\}} \frac{A_{AF}}{A_P} * \delta(\hat{\omega}_m - \hat{n}_{F'}) \tag{4.1}$$

To show this is correct, a more conventional derivation follows. As the surface model is identical under quarter-turn rotations, with four unique microfacet orientations, the NDF can be described by a constant scaled Dirac delta functions with equal terms for all four pyramid faces.

$$D_m(\hat{\omega}) = \sum_{F \in \{E,N,W,S\}} a * \delta(\hat{\omega}_m - \hat{n}_F) \tag{4.2}$$

The scale follows from the condition the differential areas integrate, over the hemisphere Ω, to match the scale of the underlying ground plane (Equation 4.3).

$$\begin{aligned} \int_{\Omega} D_m(\hat{\omega}) \cos(\theta) d\omega &= 1 \\ \int_{\Omega} \sum_{F \in \{E,N,W,S\}} a * \delta(\hat{\omega} - \hat{n}_F) \cos(\theta) d\omega &= 1 \\ 4a * \cos(\alpha) &= 1 \\ a &= \frac{1}{4\cos(\alpha)} \end{aligned} \tag{4.3}$$

This yields the NDF (Equation 4.4), confirming the previous formulation (Equation 4.1). It is essentially equivalent to the NDF provided by de Haas [4], albeit derived with slightly different assumptions and justification.

$$D_m(\hat{\omega}) = \sum_{F \in \{E,N,W,S\}} \frac{\delta(\hat{\omega} - \hat{n}_F)}{4\cos(\alpha)} \tag{4.4}$$

4.2. Shadowing Function

The shadowing function $G_1(\hat{\omega}, \hat{n}_F)$ can be found by first examining the shadowing function, locally, considering individual points on a face (Subsection 4.2.1). The global shadowing function is then approximated as the average visibility across the surface of the face (Subsection 4.2.2).

Note, as the shadowing function will be identical for each face, due to the symmetry in the model, simplification is possible. By transforming the azimuth ϕ of the viewing direction $\hat{\omega}$ to a local azimuth ϕ' relative to the normal of the face under consideration, effectively rotating the model by quarter turns, only the shadowing function of the eastern face E needs to be considered to find the shadowing function as a whole. Apart from quarter turns, the model is also symmetric with respect to z-axis rotation around viewing angles with an azimuth of $\phi' = 45°$. As such only a $22.5°$ wide azimuthal region needs to be evaluated, represented by a transformed viewing direction $\hat{\omega}'$, with $\theta \in [0, \pi/2]$ and $\phi' \in [0, \pi/4]$. Putting this together the universal shadowing function $G_1(\hat{\omega})$ is defined.

$$\hat{\omega}' = (\theta, \phi') \tag{4.5}$$

$$\phi' = \begin{cases} \phi \mod \pi/2 & (\phi \mod \pi/2) \leq \pi/4 \\ \pi/2 - (\phi \mod \pi/2) & \text{else} \end{cases} \tag{4.6}$$

$$G_1(\hat{\omega}) = G_1(\hat{\omega}', \hat{n}_E) = G_1(\hat{\omega}, \hat{n}_F) \; \forall F \in \{E, N, W, S\} \tag{4.7}$$

4.2.1. Local Shadowing Function

In order to find the shadowing function, the local shadowing function $G_1(\vec{p}, \hat{\omega})$ at an individual point $\vec{p}$ on the eastern face of a pyramid's face must be considered. By inverting the problem from one of visibility to one considering the probability of shadowing $P_S(\vec{p}, \hat{\omega})$, one can imagine a region of space, R_S, within which any secondary pyramid will obscure point $\vec{p}$ with respect to viewing direction $\hat{\omega}$ (Figure 4.1).

$$G_1(\vec{p}, \hat{\omega}) = 1 - P_S(\vec{p}, \hat{\omega}) \tag{4.8}$$

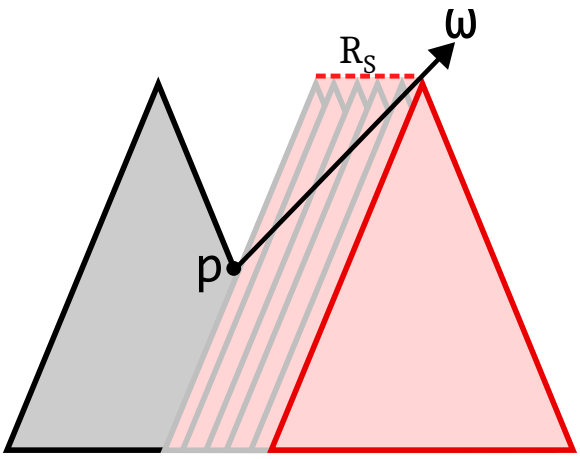

Figure 4.1: Local shadowing of point $\vec{p}$ by secondary pyramids, in viewing direction $\hat{\omega}$, with shadowing region R_S.

By considering the peaks of pyramids as sole identifiers, the region R_S can be defined by the two-dimensional region of space on the reference plane within which the secondary's peak would reside while shadowing point $\vec{p}$. Given the stochastic nature of the surface (Chapter 3), this can be described as the "shadow-spawning region", with reference plane surface area A_S.

As described in Section 3.2, the surface area $A_S(\vec{p}, \hat{\omega})$ can be used to derive the shadowing probability $P_S(\vec{p}, \hat{\omega})$ (Equation 4.9), as this equals the probability secondary peaks reside inside shadow-spawning region R_S.

$$P_S(\vec{p}, \hat{\omega}) = P_C(A_S(\vec{p}, \hat{\omega})) \tag{4.9}$$

Before going into more detail, a two-dimensional analogue is fist discussed in more depth.

2D Point Shadowing

The two-dimensional version of the problem (Figure 4.2), entails finding the local point shadowing probability $P_S(\vec{p}, \hat{\omega})_{2D}$ for point $\vec{p}$ and viewing direction $\hat{\omega}$, by finding the spawning region R_{S2D} and its one-dimensional length L_S.

The spawning region R_{S2D} is complemented by the "burying" region R_{B2D}. Any secondary pyramid's inside this region would not only shadow point $\vec{p}$ from the perspective of viewing direction $\hat{\omega}$, but also bury it below its surface. The union of these two regions can be considered the extended shadow-spawning region R_{ES2D}, delimited by two points.

The left-most delimiter is the point at which the two pyramids fully overlap, with peaks coinciding. The right-most delimiter depends on the viewing angle. At angles below the critical slant angle, at $\theta \geq \frac{\pi}{2} - \alpha$, the right-most delimiter is the point such that viewing direction $\hat{\omega}$ intersects both the secondary pyramid's peak and point $\vec{p}$. At a steeper viewing angle, at $\theta \leq \frac{\pi}{2} - \alpha$, the delimiting second pyramids face will intersects point $\vec{p}$. This is identical to the delimiting point between the burying and non-extended shadowing regions R_{B2D} and R_{S2D}. Points obscured at steeper angles will thus always be buried, with the burying region R_{B2D} coinciding with the extended shadow-spawning region R_{ES2D}. Put differently, shadowing does not occur at steeper viewing angles for non-buried points on the surface.

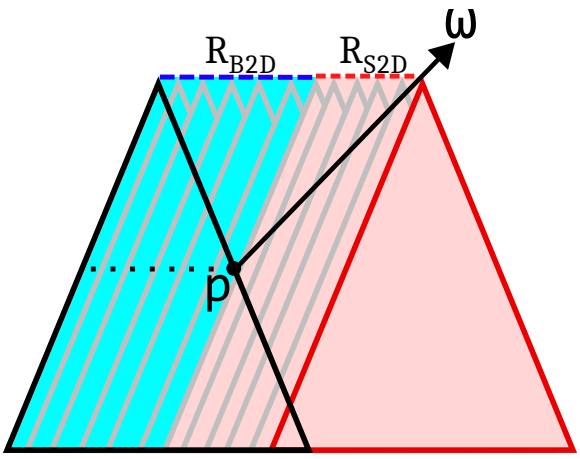

Figure 4.2: Local two-dimensional shadowing of point $\vec{p}$ by secondary pyramids, from viewing direction $\hat{\omega}$, with burying and shadowing regions R_{B2D} and R_{S2D}.

As can be seen from the dotted line in Figure 4.2, the burying region's surface length L_B is thus equal to the base length of the sub-pyramid that has $\vec{p}$ as its bottom-right corner (Equation 4.10, the two-dimensional equivalent of Equation 3.8).

$$L_B(p_z) = \frac{2p_z}{\tan(\alpha)} \tag{4.10}$$

Although enough information is given to derive the shadow-spawning region's size L_S, it is fruitful to mention one additional property that will carry over into the three-dimensional scenario. As visualized in Figure 4.3a, the size of the extended shadowing region, L_{ES}, is given by the orthogonal projection of the smaller p_z-delimited pyramid's face E, onto the horizontal reference plane, along viewing direction $\hat{\omega}$.

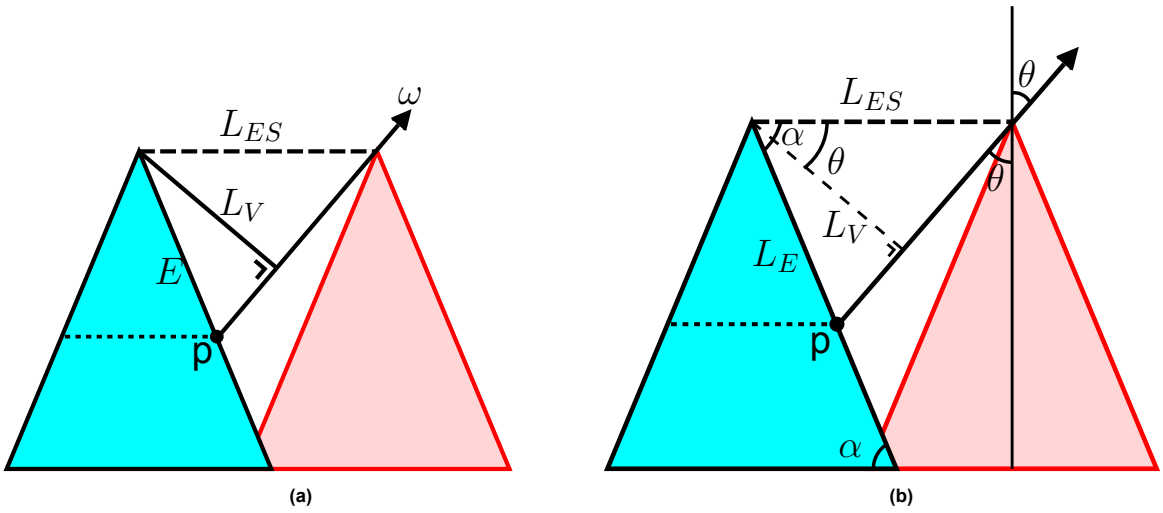

Figure 4.3: Two-dimensional projection of the p_z-delimited pyramid's east face E onto the image plane L_V and reference plane L_{ES}.

Using this property, with Figure 4.3b as a reference and L_E the length of face E, the shadow-spawning region length L_S is derived after finding L_{ES}. To simplify the derivation of this projection, the image plane perpendicular to the viewing direction functions as intermediary, providing intermediate view-space length L_V.

$$L_E(p_z) = \frac{p_z}{\sin(\alpha)} \tag{4.11}$$

$$\begin{aligned} L_V(p_z, \theta) &= \cos(\alpha - \theta) * L_E(p_z) \\ &= (\sin(\alpha)\sin(\theta) + \cos(\alpha)\cos(\theta)) * \frac{p_z}{\sin(\alpha)} \\ &= \left(\sin(\theta) + \frac{\cos(\theta)}{\tan(\alpha)}\right) p_z \end{aligned}$$

$$L_{ES}(p_z, \theta) = \frac{L_V(p_z, \theta)}{\cos(\theta)} = \left(\tan(\theta) + \frac{1}{\tan(\alpha)}\right) p_z \tag{4.12}$$

$$L_S(p_z, \theta) = L_{ES}(p_z, \theta) - L_B(p_z) = \left(\tan(\theta) - \frac{1}{\tan(\alpha)}\right) p_z \tag{4.13}$$

3D Point Shadowing

In accordance to the two-dimensional situation, the three-dimensional shadow-spawning region's area $A_S(\vec{p}, \hat{\omega})$ can be found by orthographic projection of the three-dimensional p_z-delimited sub-pyramid onto the reference plane along viewing direction $\hat{\omega}$ (Figure 4.5). The fact this is still valid is however less obvious.

Conceptually, with the shadowing pyramid is shifted around in space, every point of the p_z-delimited pyramid can uniquely shadow point $\vec{p}$ from the perspective of the viewer from direction $\vec{\omega}$, assuming it is on a front-facing face. Each of these shadowing points matches a single unique shadowing pyramid peak in region R_{ES} on the reference plane. There is thus a one-to-one mapping between front-facing points on the p_z-delimited pyramid and the peaks in the shadow-spawning region. Points lower on the shadowing pyramid are unable to shadow point $\vec{p}$ due to their lower z values, while back-facing points on the shadowing pyramid overlap with an already considered matching front-facing point.[1]

Considering the shadow-spawning peaks also do not overlap in view-space, the one-to-one mapping is retained in view-space. The shadow-spawning region can thus essentially be found in view-space by shifting the p_z-delimited pyramid's projection about point $\vec{p}$ and recording its peak's position, as visualized in Figure 4.4.

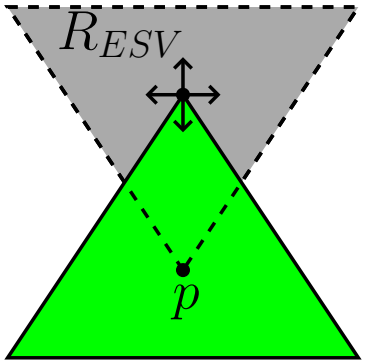

Figure 4.4: Tracing the view-space extended shadow-spawning region R_{ESV} by overlapping the p_z-delimited pyramid's projection (green) with $\vec{p}$ and tracking its peak's position, resulting in the projection's exact inversion (grey).

This view-space tracing results in a view-space shadow-spawning region R_{ESV} with the exact (inverted) shape of the p_z-delimited pyramid's projection onto the image plane. This inversion being the result of the way the shadowing pyramid's projection needs to be shifted left for $\vec{p}$ to be shadowed by a point on its right, for example. It has no further implications, however, as only the region's surface area is of concern.

[1] For non-convex shapes this mapping would only be one-to-one for top-most front-facing points in view-space. The modified argument would, however, still be valid.

The extended shadow-spawning region R_{ES} is then found by projecting the view-space region R_{ESV} back onto the reference plane along $\hat{\omega}$, finding the shadowing peaks in 'world-space' matching those found in view-space. The final shape of the extended shadow-spawning region thus takes on the stretched and inverted form of the p_z-delimited pyramid as projected onto the image plane (Figure 4.6).

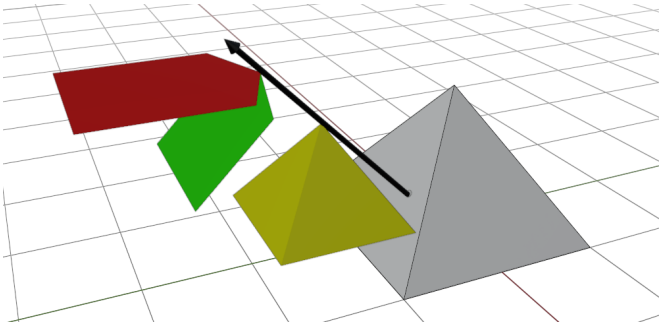

Figure 4.5: Projection of the p_z-delimited shadowing pyramid (yellow) onto image (green) and reference planes (red). The arrow originates from shadowed point $\vec{p}$ and points along viewing direction $\hat{\omega}$.

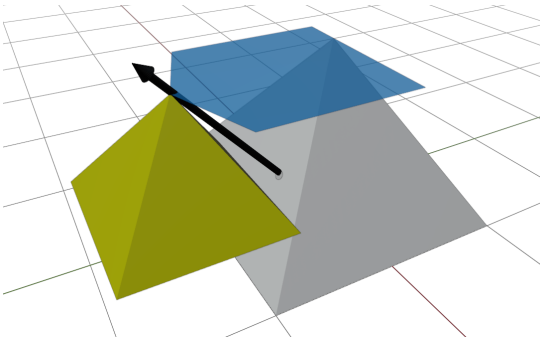

Figure 4.6: Extended shadow-spawning region R_{ES} (blue) above the shadowed pyramid (grey) and p_z-delimited shadowing pyramid (yellow). The arrow originates from shadowed point $\vec{p}$ and points along viewing direction $\hat{\omega}$.

The extended shadow-spawning region's surface area A_{ES} can now be derived. In order to find the scale of the view-space projection of the p_z-delimited pyramid A_V, individual contributions A_{VF} of all faces are firstly considered, with 'F' as a placeholder for any of the faces. The surface area of the full projection then follows as the sum of the individual contributions of front-facing faces, as their individual view-space shadowing regions R_{VF} cannot overlap, given the pyramid's convex shape.

As projection scales a surface's area by the cosine of the angle between the original and projected surface, A_{VF} is given by the normal vector dot-product scaling of the front-facing faces' surface areas A_F. Taking the summation A_V and back-projecting it onto the reference plane is done inversely, inversely scaling by $\cos(\theta)$, yielding A_{ES} (Equation 4.17).

$$A_F(p_z) = \frac{1}{2} * w(p_z) * \frac{p_z}{\sin(\alpha)} = \frac{p_z^2 \sin(\alpha)}{\cos^2(\alpha)} \tag{4.14}$$

$$A_{VF}(p_z, \hat{\omega}) = \max(0, \hat{\omega} \cdot \hat{n}_F) * A_F(p_z) \tag{4.15}$$

$$A_V(p_z, \hat{\omega}) = \sum_{F \in \{E,N,W,S\}} A_{VF}(p_z, \hat{\omega}) \tag{4.16}$$

$$A_{ES}(p_z, \hat{\omega}) = \frac{A_V(p_z, \hat{\omega})}{\cos(\theta)} \tag{4.17}$$

The shadow-spawning region's surface area A_S, and the local shadowing function G_1 as a whole, is then found by subtracting the burying region's surface area A_B, as the existence of the point on the surface should be considered as a precondition. Since $\vec{p}$ is buried by the p_z-delimited pyramid's base, centered at $\vec{p}$, the burying regions' surface area equals the p_z-delimited pyramid's base surface area $A_B(p_z)$ (Equation 3.8).[2]

$$A_S(p_z, \hat{\omega}) = A_{ES}(p_z, \hat{\omega}) - A_B(p_z) \tag{4.18}$$

$$G_1(p_z, \hat{\omega}) = P_0(A_S(p_z, \hat{\omega})) = e^{-\rho A_S(p_z, \hat{\omega})} \tag{4.19}$$

In order to simplify, it can be useful to define a z-independent dot-product sum $D(\hat{\omega})$ describing the view-space surface area as a factor of the faces' surface areas $A_F(p_z)$, as these are identical for all

[2] Here, A_B refers to both the burying region's surface area as well as the pyramid's base surface area, as they are equivalent.

faces. Beware this is distinct from the microfacet distribution $D_m(\hat{\omega})$ (Section 4.1).

$$D(\hat{\omega}) = \sum_{F \in \{E,N,W,S\}} \max(0, \hat{\omega} \cdot \hat{n}_F) \tag{4.20}$$

$$A_S(p_z, \hat{\omega}) = \frac{D(\hat{\omega}) A_F(p_z)}{\cos(\theta)} - A_B(p_z) \tag{4.21}$$

For the actual expressions of these dot-products, as well as the normals, view appendix B.

4.2.2. Global Shadowing Function

Given the local shadow function $G_1(p_z, \hat{\omega})$ (Equation 4.19), the global shadowing function $G_1(\hat{\omega})$ is found by determining the average visibility across the face under consideration. For this, the surface integral over the face is weighted using the surface height distribution PDF $E(z)$ (Equation 3.11). As discussed at the start of Section 4.2, only the eastern face 'E' needs to be considered.

$$\begin{aligned}
G_1(\hat{\omega}) &= \int_0^\infty G_1(z, \hat{\omega}) E(z) dz \\
&= \int_0^\infty e^{-\rho(D(\hat{\omega})A_F(z)/\cos(\theta) - A_B(z))} \frac{8\rho z}{\tan^2(\alpha)} e^{-\rho A_B(z)} dz \\
&= \frac{8\rho}{\tan^2(\alpha)} \int_0^\infty z e^{-\rho(D(\hat{\omega})A_F(z)/\cos(\theta)} dz \\
&= \frac{8\rho}{\tan^2(\alpha)} \int_0^\infty z e^{-\rho z^2 D(\hat{\omega}) \cos(\alpha)/\sin^2(\alpha)\cos(\theta)} dz \\
&= \frac{8\rho}{\tan^2(\alpha)} \frac{\sin^2(\alpha)\cos(\theta)}{-2\rho D(\hat{\omega})\cos(\alpha)} \left[e^{-\rho z^2 D(\hat{\omega})\cos(\alpha)/\sin^2(\alpha)\cos(\theta)} \right]_0^\infty \\
&= \frac{-4\cos(\alpha)\cos(\theta)}{D(\hat{\omega})}(0 - 1)
\end{aligned} \tag{4.22}$$

$$G_1(\hat{\omega}) = \frac{4\cos(\alpha)\cos(\theta)}{D(\hat{\omega})} \tag{4.23}$$

The shadowing function found is depicted in Figure 4.7, alongside ray-traced measurements.

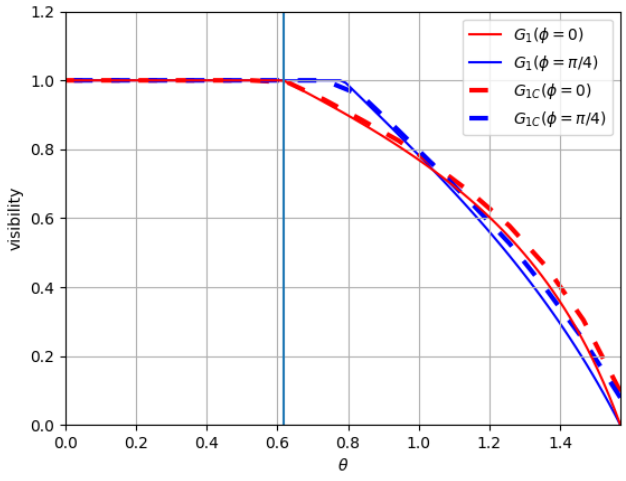

Figure 4.7: Shadowing function G_1 and constant-height shadowing function G_{MC} measured using 4000 ray-traced samples per data-point, using pyramid density $\rho = 0.6$ *pyramids*$/\mu m^2$ and slant angle $\alpha = 54.7°$. The vertical line denotes the critical viewing angle.

Interestingly, the resulting shadowing function can be shown to be identical to the generic shadowing function described by Heitz [3], originally derived by Ashikmin *et al.* [34, Equation 42] (Equation 4.24).

Here χ^+ is the Heaviside step function. It is explicitly derived under the assumption the microfacet normal/shadowing at any point is uncorrelated with respect to other points on the surface.

$$G_1(\hat{\omega}, \hat{n}_F) = \chi^+(\hat{\omega}, \hat{n}_F) \frac{\cos(\theta)}{\int_\Omega max(0, \hat{\omega} \cdot \hat{\omega}_m) D_m(\hat{\omega}_m) d\hat{\omega}_m} \tag{4.24}$$

Additionally, at an azimuth of $\phi' = 0$, the shadowing function is identical to the azimuth-agnostic shadowing function G_L found by de Haas [4][3]. This is shown by rewriting the dot-product sum $D(\hat{\omega})$ using individual terms (Equation B.3) at $\phi = 0$.

$$G_L(\theta) = \begin{cases} 1 & \theta \leq \pi/2 - \alpha \\ \frac{4\cos(\alpha)\cos(\theta)}{2\cos(\alpha)\cos(\theta)+\cos(\theta-\alpha)} & \text{else} \end{cases} \tag{4.25}$$

$$\begin{aligned} D_E(\theta) &= \cos(\theta - \alpha) \\ D_W(\theta) &= \cos(\theta + \alpha) \\ D_N(\theta) &= \cos(\alpha)\cos(\theta) \\ D_S(\theta) &= \cos(\alpha)\cos(\theta) \end{aligned} \tag{4.26}$$

$$D(\theta, \phi = 0) = \begin{cases} 4\cos(\alpha)\cos(\theta) & \theta \leq \pi/2 - \alpha \\ 2\cos(\alpha)\cos(\theta) + cos(\theta - \alpha) & \text{else} \end{cases} \tag{4.27}$$

By applying this to the global shadowing function (Equation 4.23) equality is proven.

$$G_1(\theta, \phi = 0) = \begin{cases} \frac{4\cos(\alpha)\cos(\theta)}{4\cos(\alpha)\cos(\theta)} = 1 & \theta \leq \pi/2 - \alpha \\ \frac{4\cos(\alpha)\cos(\theta)}{2\cos(\alpha)\cos(\theta)+\cos(\theta-\alpha)} & \text{else} \end{cases} \tag{4.28}$$

Normalization

As a requirement, the shadowing function should be normalized such that the total sum of visible portions equals the actual visible surface area. In particular, the visible portions should sum to equal the visible reference plane in view-space. This requirement is described by Equation 4.29, where the faces' surface areas are taken to be equal to their statistical average A_{AF} (Equation 3.14).

$$|C| * \sum_{F \in \{E,N,W,S\}} G_1(\hat{\omega}) \max(0, \hat{\omega} \cdot \hat{n}_F) A_{AF} = \cos(\theta) * A_T \tag{4.29}$$

This normalization requirement is then proven by elimination[4].

$$\begin{aligned} \rho \sum_{F \in \{E,N,W,S\}} \frac{4\cos(\alpha)\cos(\theta)}{D(\hat{\omega})} \max(0, \hat{\omega} \cdot \hat{n}_F) \frac{1}{\rho 4\cos(\alpha)} &= \cos(\theta) \\ \sum_{F \in \{E,N,W,S\}} \max(0, \hat{\omega} \cdot \hat{n}_F) &= D(\hat{\omega}) \\ D(\hat{\omega}) &= D(\hat{\omega}) \end{aligned}$$

This requirement is more commonly described by Equation 4.30, as described by Heitz [3, Sec. 2.4]. Although the latter is expressed infinitesimally, it can be shown to be equivalent by factoring out A_T and using the microfacet NDF derived in Section 4.1. Note the microfacet NDF $D_m(\hat{\omega}_m)$ here is distinct from the dot-product sum $D(\hat{\omega})$.

$$\int_\Omega D_m(\hat{\omega}_m) G_1(\hat{\omega}, \hat{\omega}_m) \max(0, \hat{\omega} \cdot \hat{\omega}_m) d\hat{\omega}_m = \cos(\theta) \tag{4.30}$$

[3] The expression provided in equation 4.25 is rewritten for simplicity but otherwise identical to the original.

[4] Note, as the shadowing function is equivalent to the generic shadowing function discussed by Ashikmin *et al.* [34] and Heitz [3], this is already guaranteed.

4.3. Visible Normal Sampling

In order to describe the BRDF, and allow for it to be sampled, it is necessary to find the probability a microfacet is hit given some incoming ray from direction $\hat{\omega}$. As the visible view-space surface area of particular faces attenuate the probability of hitting these faces, sampling of a particular face F, with probability $P_{VNS}(F)$, is performed proportional to this visible view-space surface area A_{VVF}, given by applying the shadowing function to the view-space surface area A_{VF} (Equation 4.15).

$$P_{VNS}(F, \hat{\omega}) \propto A_{VVF}(\hat{\omega}) = A_{VF}(h, \hat{\omega}) G_1(\hat{\omega}) \tag{4.31}$$

Since only the relative proportionality of these is of importance, it's simplest to factor out common parts and normalize. This yields the sampling probability:

$$P_{VNS}(F, \hat{\omega}) = \frac{\max(0, \hat{\omega} \cdot \hat{n}_F)}{D(\hat{\omega})} \tag{4.32}$$

This result is equivalent to the visible normal distribution $D_n(\hat{\omega})$ described by Heitz [3, Section 3.1] (Equation 4.33), as can be shown be substitution.

$$\begin{aligned} D_n(F, \hat{\omega}) &= \frac{G_1(\hat{\omega})}{\cos(\theta)} D_m(\hat{n}_F) \max(0, \hat{\omega} \cdot \hat{n}_F) \\ D_n(\hat{\omega}) &= \sum_{F \in \{E,N,W,S\}} D_n(F, \hat{\omega}) * \delta(\hat{\omega}_m - \hat{n}_F) \end{aligned} \tag{4.33}$$

4.4. BRDF Sampling

Provided with a shadowing function and a visible normal sampling description, a BRDF can now be constructed. The BRDF will need to be sampled by a renderer, however, given the specularity of the model. As such, the set of possible exit directions and their matching sampling probability are required.

As discussed at the start of this chapter, multiple reflections within the surface are considered. For this, no positional information is used, with shadowing modeled as a global average, as discussed in Subsection 4.2.2. Likewise, pyramid positions are not sampled or recorded, as a purely statistical approach is taken.

The resulting BRDF can be described recursively, with potential paths taken defined as the sequence of faces hit. The resulting exit direction follows trivially by recursively applying reflections to the input direction using the face normals along the path. The recursive algorithm for calculating the probability a particular path $P_P(\hat{\omega}_i, path)$ is taken is described below using pseudo-code. At every step the probability of reflecting off a specific face along the path is evaluated, given by the visible normal sampling probability P_{VNS}(Equation 4.32). The shadow function G_1(Equation 4.23) subsequently provides the probability the resulting ray exits the surface (or continues along the path). The actual BRDF values $f(\hat{\omega}_i, path)$, matching these paths, are then given by the path occurrence probability, scaled by the exiting cosine (Equation 4.34)[5]. The resulting BRDF is then given as the sum of Dirac delta functions (Equation 4.35).

$$f(\hat{\omega}_i, path) = \frac{P_P(\hat{\omega}_i, path)}{\cos(\theta_o)} \tag{4.34}$$

$$f(\hat{\omega}_i, \hat{\omega}_o) = \sum_{path \in paths} \delta(\hat{\omega}_o - \hat{\omega}_o(path)) * f(\hat{\omega}_i, path) \tag{4.35}$$

[5]This scaling is the result of the definition of BRDFs in the context of the rendering equation. For more the reader is referred to [33].

Algorithm 1 Recursive Scattering Path Probability

```
1: procedure PathProbability(ω̂_i, path)
2:     ω̂ ← ω̂_i
3:     prob ← 1
4:     exitProb ← 0
5:     for i ← 0 to length(path) do                ▷ Iterate over scattering events
6:         face ← path[i]
7:         prob ← prob * P_VNS(face, ω̂)            ▷ Visible Normal Sampling
8:         ω̂ ← Reflect(face, ω̂)
9:         shadowing ← G_1(ω̂)
10:        exitProb ← prob * shadowing
11:        prob ← prob − exitProb
12:    end for
13:    return exitProb
14: end procedure
```

Actual implementation of the BRDF can be performed in a multitude of ways. A rudimentary approach was taken in the evaluation of this work. Recursion is performed online (at run-time) to calculate path exit probabilities up to a limited path-length, given a viewing direction $\hat{\omega}_i$, with which sampling is performed. A more efficient approach would be to tabulate path probabilities offline, at the cost of some precision, and potentially use interpolation to combine table entries at run-time, sampling from the combined result. Where applicable, GPU hardware could even be leveraged, e.g. using tri-linearly interpolated 3D textures.

One drawback of the approach described thus far becomes evident once compared with ray-traced results over a generated surface. In particular, since no path history is maintained, no consideration is taken of previous scattering events. This results in at least two very noticeable deviations. The first is related to the disappearance of previously encountered pyramids (Section 4.5). The second is related to the introduction of pyramids that invalidate the path (Section 4.6).

4.5. Back-Scattering Correction

Since the previously described recursive algorithm is forgetful, rays will not take their own history into consideration. This will manifest itself in different ways, the most prevalent of which occurs when a ray is scattered back onto a previously visited face with high probability. As opposed to scattering off this face, the ray may incorrectly escape the surface directly.

The lack of this back-scattering can introduce a large error depending on the viewing direction, as shown in the exit-path histogram in Figure 4.8, comparing the uncorrected BRDF with a ray-traced reference. In this case, the model overestimates the probability of ”WE” paths occurring, paths that scatter of a west facing pyramid face followed by an opposing east face, while underestimating the exit probability of ”WEW” paths. In order to correct for this behavior in path p_1, the probability of back-scattering P_{BS} must be determined, with the back-scattering probability offsetting the previously calculated exit probability $P_P(p_1)$ onto back-scattered path p_2. This approach only considers the last scattering event. In the event of longer paths, intermediate probability variables should be updated instead.

$$P_P(p_1)' = P_P(p_1) - P_P(p_1) * P_{BS}(p_1) \tag{4.36}$$

$$P_P(p_2)' = P_P(p_2) + P_P(p_1) * P_P(p_1) \tag{4.37}$$

Geometrically speaking, P_{BS} represents the probability the last pyramid hypothetically sampled on the path occurred within a specific region of the reference plane within which the reflected ray would back-scatter on the previous pyramid. This region of space is categorized by the back-scattering spawning volume V_{BS} within which ray an outgoing ray would hit the face of the previous pyramid, as visualized in Figure 4.10. This corresponds to ray c_2 and the left pyramid in Figure 4.9.

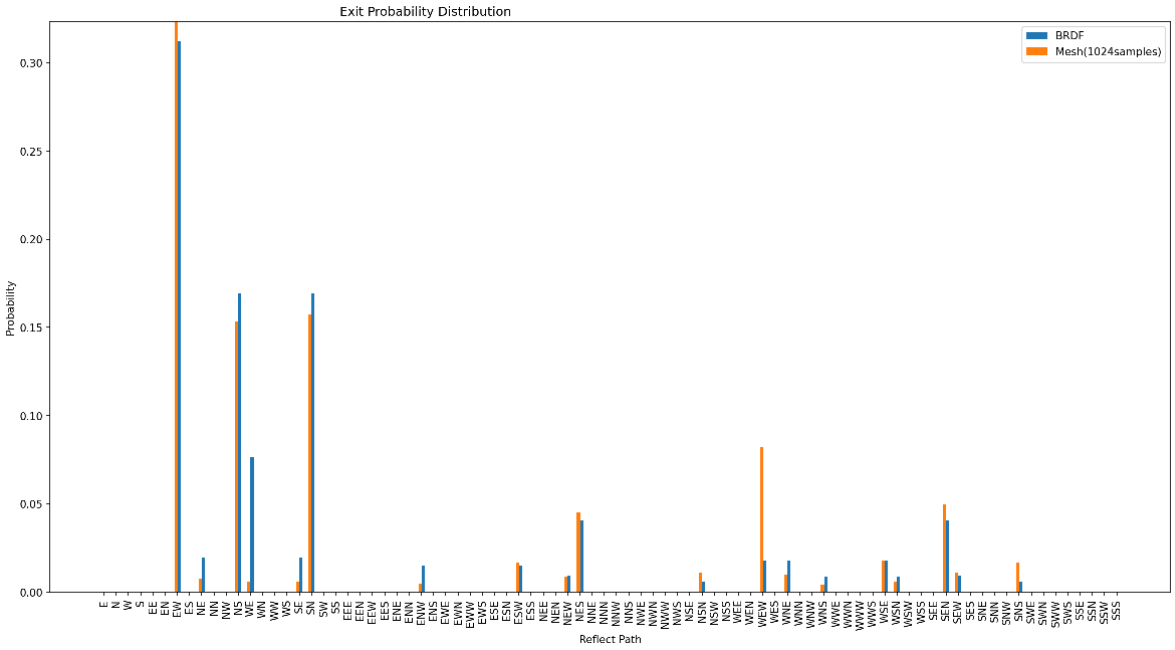

Figure 4.8: Exit-path histogram of uncorrected BRDF alongside ray-traced reference, using a generated constant pyramid height surface with pyramid density $\rho = 0.6$ *pyramids*$/\mu m^2$ and slant angle $\alpha = 54.7°$, with $\theta = 15°$ and $\phi = 0°$.

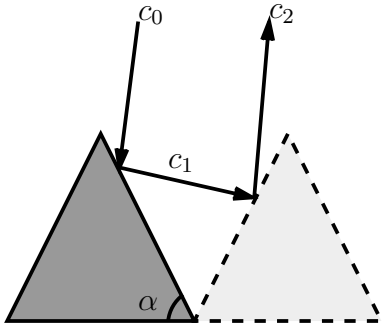

Figure 4.9: Reflections off of opposing faces.

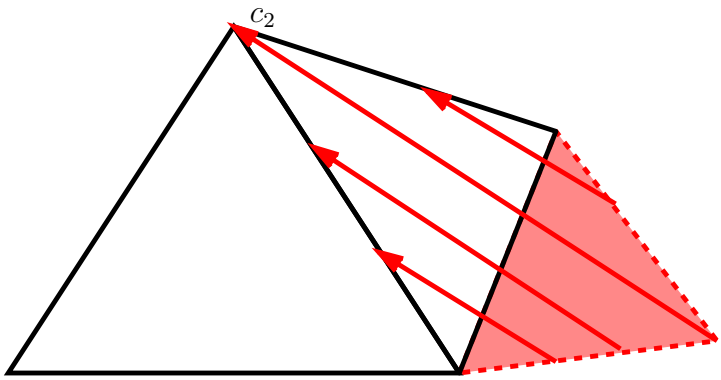

Figure 4.10: Back-scattering spawning volume V_{BS} under c_2.

In line with the integration approach taken in Subsection 4.2.2, an average probability of back-scattering can be found by taking the surface integral over the face hit in the event of back-scattering. In contrast to the approach taken before, the integrals involved are solved numerically, as their complexity makes algebraic solutions unfeasible. In addition, due to the high cost of this integration, the solutions are precomputed and tabulated offline, with interpolation performed online (at run-time).

The local back-scattering probability is first discussed (Figure 4.11), followed by the global solution (Subsection 4.5.2). In line with the approach taken in Section 4.2, the problem is examined for a ray scattered off of the eastern face, with the general solution following by quarter rotations of $\hat{\omega}$ for the other faces.

4.5.1. Local Back-Scattering

In the following, local back-scattering (Figure 4.11) is considered provided a given position $\vec{p}$ and incoming ray $\hat{c}_0$. Here, the probability of back-scattering is identical to the probability the face of a secondary pyramid is hit along the path exiting volume $V_B S$, starting from point $\vec{p}$, in direction $\hat{c}_1$.

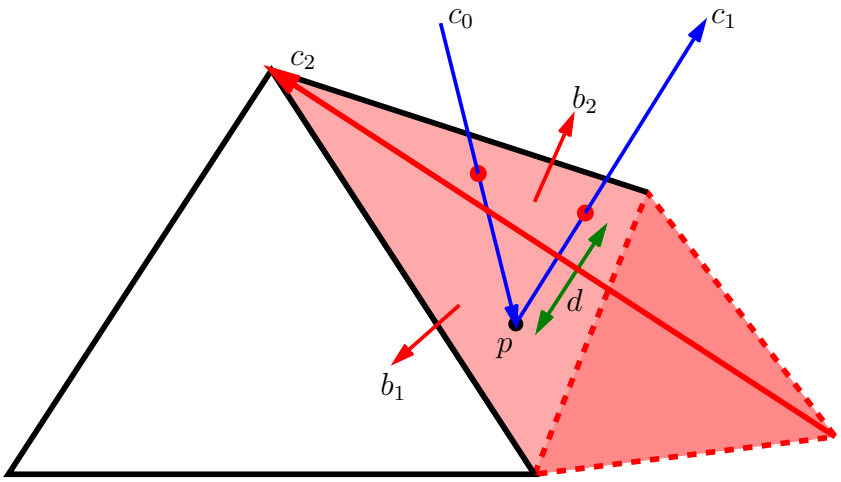

Figure 4.11: Back-scattering spawning volume V_{BS} (red) with normals $\hat{b}_1$ and $\hat{b}_2$, pierced by the path exiting from point $\vec{p}$ in direction $\hat{c}_1$, with exit travel-distance d.

The pyramid spawning region matching the potential hit-points along this path is found by integration along the path. The hypothetical secondary pyramid's face width determines the width of the spawning region at each position along the ray. The perpendicular size of the spawning region along $\hat{x}$ is derived in differential form ds following Figure 4.12. Here W_1 and W_2 are the faces of two hypothetical secondary pyramids, with the x-distance between their peaks forming the differential spawning region's size along $\hat{x}$. Rays that rise at a steeper rate than the slant angle are unable to hit a secondary peak, and are thus invalid candidates for back-scattering correction.

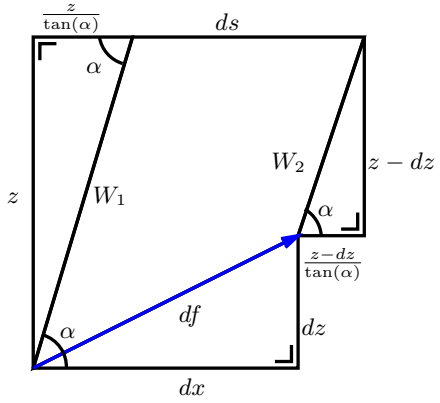

Figure 4.12: Exit path's differential spawning region's size ds along $\hat{x}$, with W_1 and W_2 representing the hypothetical faces hit along the path (blue).

Given the spawning region's differential size, its reference plane surface area can be calculated by integration along the path, yielding the local back-scattering probability using the pyramid spawning probability P_C (Section 3.2):

$$ds = dx - \frac{dz}{\tan(\alpha)} \tag{4.38}$$

$$\frac{dx}{df} = c_{1x}\ ,\ \frac{dz}{df} = c_{1z} \tag{4.39}$$

$$P_{BS}(p_z, \hat{c}_1, d) = P_C\left(\int_0^d w(p_z + f * c_{1z})\frac{ds}{df}df\right) \tag{4.40}$$

$$= P_C\left(\int_0^d w(p_z + f * c_{1z}) * \left(c_{1x} - \frac{c_{1z}}{\tan(\alpha)}\right)df\right) \tag{4.41}$$

The actual exit travel-distance d can be determined with explicit intersection calculation. Since the back-scattering spawning volume V_{BS} is delineated by two sides, this results in two exit travel-distances d_1 and d_2, with only the lowest positive travel-distance being the valid exit travel-distance d. In cases where both exit travel-distances are lower than or equal to zero, local back-scattering is guaranteed. Here, the east face's legs are defined as $\vec{l}_1$ and $\vec{l}_2$. In the following, their scales are not of importance.

$$\vec{l}_1 = \begin{pmatrix} 1 \\ -1 \\ -\tan(\alpha) \end{pmatrix} \qquad \vec{l}_2 = \begin{pmatrix} 1 \\ 1 \\ -\tan(\alpha) \end{pmatrix} \tag{4.42}$$

$$\vec{b}_1 = \vec{l}_1 \times \hat{c}_2 \qquad \vec{b}_2 = \vec{l}_2 \times \hat{c}_2 \tag{4.43}$$

$$d_1 = \frac{-\vec{p}\cdot\vec{b}_1}{\vec{c}_1\cdot\vec{b}_1} \qquad d_2 = \frac{-\vec{p}\cdot\vec{b}_2}{\vec{c}_1\cdot\vec{b}_2} \tag{4.44}$$

4.5.2. Global Back-Scattering

The surface integral required to calculate the average local back-scattering probability is complex due to the existence of two different exit travel-distances. In order to allow for integration, the surface integral has to be split into two parts, with either part corresponding to the use one of the exit travel-distance. This is done by recognizing the face under consideration can be split into two parts, divided by the line along which the two exit travel-distances are equal (Figure 4.13), defined by the y-coordinate Y_{mid}. As Y_{mid} scales along with the pyramid, it can be defined as a linear factor of z:

$$Y_{mid}(z, \hat{c}_0) = z * Y_{mid}(1, \hat{c}_0) \tag{4.45}$$

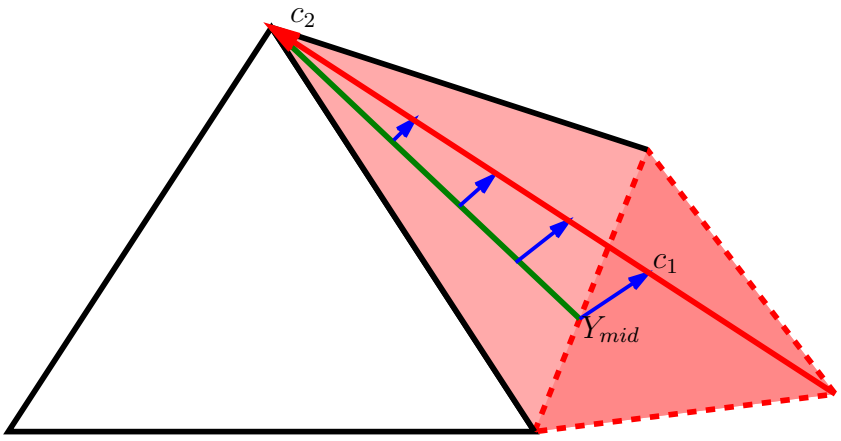

Figure 4.13: Back-scattering spawning V_{BS} (red), with delimiting line (green) on $Y_{mid}(z)$.

The surface integral averaging the local back-scattering probability (Equation 4.51) is expressed using intermediates I_1 and I_2. Here z_F is the horizontal projection of z onto a face's surface (Equation 3.15).

$$Y_{min}(z) = \frac{-z}{\tan(\alpha)} \tag{4.46}$$

$$Y_{max}(z) = \frac{z}{\tan(\alpha)} \tag{4.47}$$

$$Y'_{mid}(z, \hat{c}_0) = \max(Y_{min}(z), \min(Y_{max}(z), Y_{mid}(z, \hat{c}_0))) \tag{4.48}$$

$$I_1(z, \hat{c}_0) = \int_{Y_{min}(z)}^{Y'_{mid}(z,\hat{c}_0)} P_{BS}(z, \hat{c}_1, d_1)dy \tag{4.49}$$

$$I_2(z, \hat{c}_0) = \int_{Y'_{mid}(z,\hat{c}_0)}^{Y_{max}(z)} P_{BS2}(z, \hat{c}_1, d_2)dy \tag{4.50}$$

$$P_{BS}(\hat{c}_0) = \frac{\int_0^\infty (I_1(z, \hat{c}_0) + I_2(z, \hat{c}_0)\, P_E(z) \frac{dz_F}{dz} dz}{\int_0^\infty P_E(z) w(z) \frac{dz_F}{dz} dz} \tag{4.51}$$

Note, in contrast to Equation 4.23, $P_E(z)$ is used instead of $E(z)$, alongside normalization, as the face's width $w(z)$ is split into separate parts. Here, take note of Equation 3.12.

In order to find a numeric expression for the integral, Maple [35] was used. For details with regards to the numeric solutions please view appendix C. The resulting back-scattering probability tabulation is depicted in Figure 4.14.[6] Since the back-scattering probability is symmetric, as was the case for the shadowing function, the tabulation is calculated at a higher resolution, without additional costs, by considering only the subdomain $\theta \in [0, \pi/2]$ $\phi \in [0, \pi]$.

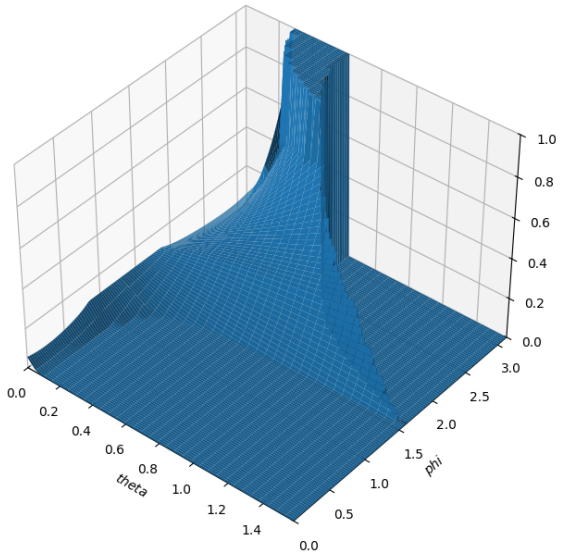

Figure 4.14: Improved back-scattering probability tabulation, with pyramid density $\rho = 0.6$ *pyramids*$/\mu m^2$ and slant angle $\alpha = 54.7°$.

While performing tabulation, it is important to take into consideration back-scattering edge-conditions. The path under consideration needs to be able to hit the opposing faces in succession. Additionally, heed needs to be taken of back-scattered paths that do not exit the surface. The input validity region is depicted in Figure 4.15.

[6] Interestingly enough, the naive use of dx instead of ds in calculating the local back-scattering probability yields an almost undistinguishable difference in tabulation results.

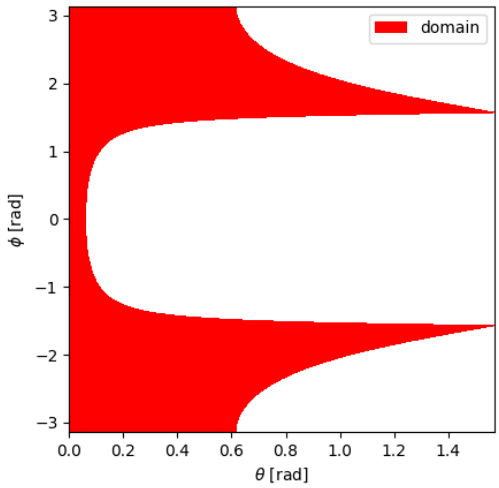

Figure 4.15: Valid back-scattering input region, with pyramid density $\rho = 0.6$ *pyramids*$/\mu m^2$ and slant angle $\alpha = 54.7°$.

Several small optimizations can be performed for numeric precomputations, one of which occurs when back-scattering is globally guaranteed, skipping integral computation. The exit travel-distances are individually guaranteed to be either positive or negative across the entire face. Given this property, the two exit travel-distance can be calculated at arbitrary positions on the face to test for guaranteed global back-scattering, as is the case when both exit travel-distances are negative or equal to zero.

Matching the exit-path diagram exemplifying the necessity of back-scattering correction (Figure 4.8), the exit-path diagram of the back-scattering corrected BRDF, compared with a ray-traced reference, is provided in Figure 4.16. With an identical viewing angle, this exhibits the desired behavior. More exhausted analysis in provided in Chapter 5.

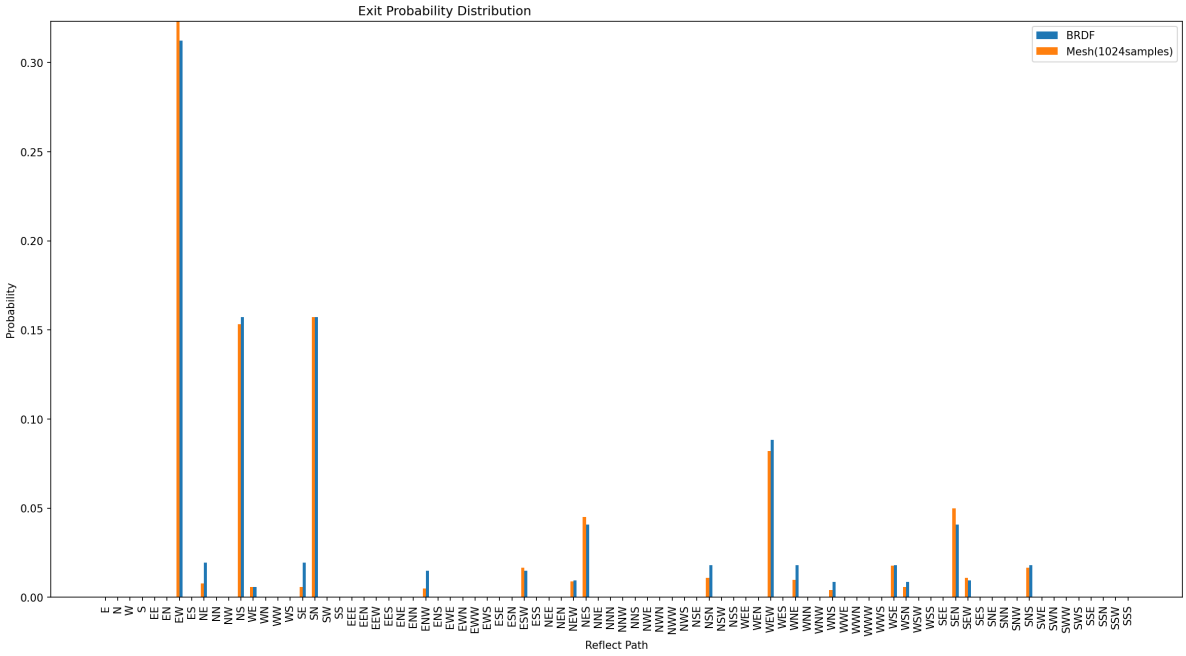

Figure 4.16: Exit-path histogram of back-scattering corrected BRDF alongside ray-traced reference, using a generated constant pyramid height surface with pyramid density $\rho = 0.6$ *pyramids*$/\mu m^2$ and slant angle $\alpha = 54.7°$, with $\theta = 15°$ and $\phi = 0°$.

4.6. Self-Burying

Apart from back-scattering, another closely related effect that introduces error in the taken approach is the way a subsequent pyramid may inadvertently invalidate the previous path's history by essentially burying the path. Although valid in isolation, the shadow-spawning region used in the shadowing function may thus contain an invalidation region, introducing additional scattering where none should occur. This is exemplified in Figure 4.17, comparing the exit-path histograms of the BRDF modeled thus far with a ray-traced reference. Note how additional scattering events are introduced along the "EW" exit path.

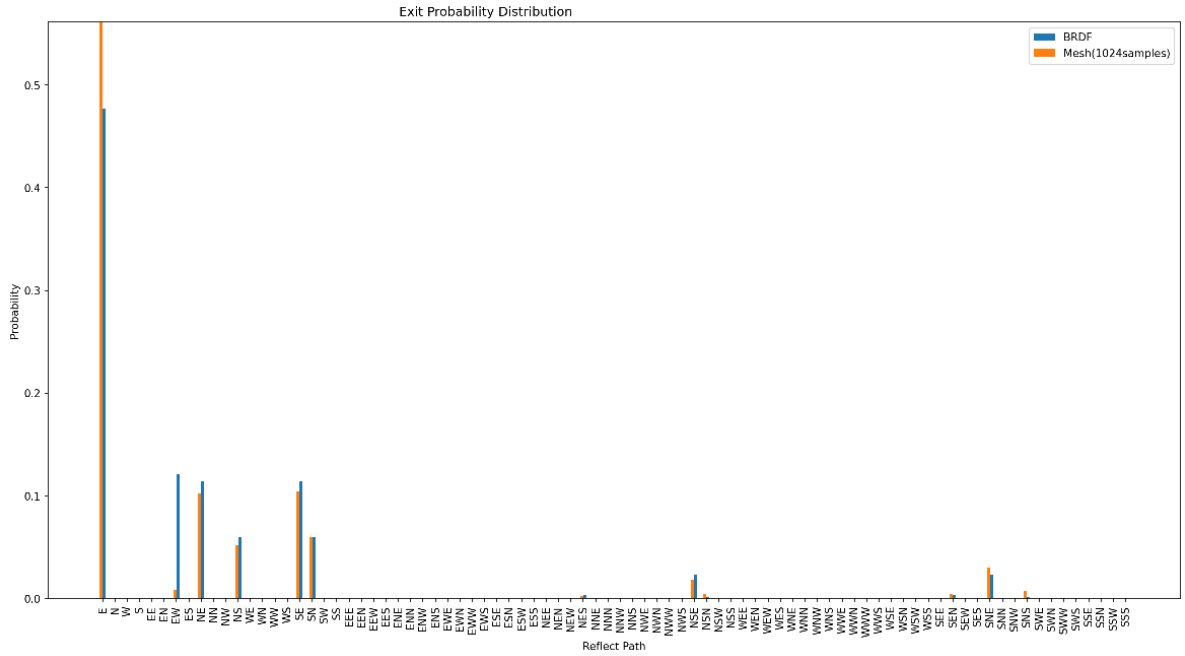

Figure 4.17: Exit-path histogram of back-scattering corrected algorithm alongside ray-traced reference, with $\theta = 54.46°$ and $\phi = 0°$.

This effect is especially prevalent at shallow viewing angles, as this results in larger shadow-spawning invalidation regions. A solution to this problem is not provided in this work.

4.7. Optical Filter Coating BRDF

The optical filter coating is modeled as a wavelength and angle dependent reflectance function on top of the textured surface. It is assumed to not impact the surface's geometry, as described in Chapter 3, and is likewise modeled using specular reflections. Given these assumptions the scattering behavior of the textured model can be maintained.

Assuming a reflectance profile is known, the reflectance value can be used to filter the different wavelengths compromising the ray. As each scattering event off of a pyramid's coated face entails to the application of the filter, the reflectance profile is sampled at each scattering event along the path (Section 4.4), with the throughput of each wavelength scaled by the resulting reflectance value. Expanding on the uncoated model, the BRDF along a path is thus scaled by both the shadowing function and reflectance value. Any light transmitted through the coating is assumed absorbed, either within the coating itself or by the textured pyramid. The reflectance profile used in this work was modeled after the reflectance profile simulated by Ortiz Lizcano *et al.* [24, Section 2.2 Figure 2.A]. The fitted model is depicted in Figure 4.18. Within the range of visible light, the main reflectance profile at constant angle can be characterized by a bell curve with a minimum lower-bound, with angle shifting and lowering the bell curve while raising the lower-bound. The bell curves are approximated using super-gaussian functions, while the angle-dependent behavior is approximated by linear interpolation.

The resulting reflectance expression R is described below, with λ the wavelength, R_B the base radiance, A_g the gaussian's amplitude, C_g the gaussian's center, and σ_g the gaussian's standard deviation. The values interpolated between are given in Table 4.1. Note angles beyond $90°$ are clipped to $80°$.

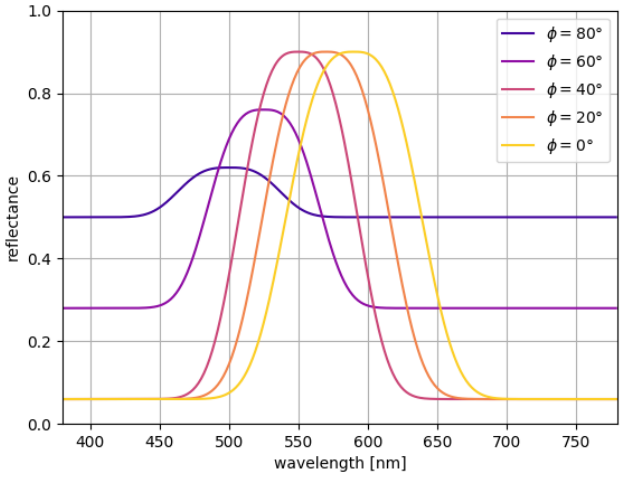

Figure 4.18: Modeled wavelength and angle dependent reflectance profile.

$$R(\theta, \lambda) = R_B(\theta) + A_g(\theta) * e^{\left(\frac{-(\lambda - C_g(\theta)^2}{2\sigma_g(\theta)^2}\right)^{1.5}} \tag{4.52}$$

θ	R_B	A_g	C_g	σ_g
$0°$	0.06	0.84	$590°$	$40°$
$40°$	0.06	0.84	$550°$	$35°$
$80°$	0.5	0.12	$500°$	$30°$

Table 4.1: Fitted reflectance profile interpolation values.

5

Results

In the following, the derived shadowing function (Section 5.1) and multiple forms of the BRDF (Section 5.2) are compared with reference results procured using generated constant pyramid height surfaces, with pyramid density $\rho = 0.6$ *pyramids*$/\mu m^2$ and slant angle $\alpha = 54.7°$. Renders comparing the BRDF with references are subsequently included (Section 5.3). Finally, a discussion is provided (Section 5.4).

5.1. Shadowing Function

The shadowing function found is compared with the shadowing function derived by de Haas [4], alongside reference ray-traced measurements (Figure 5.1).

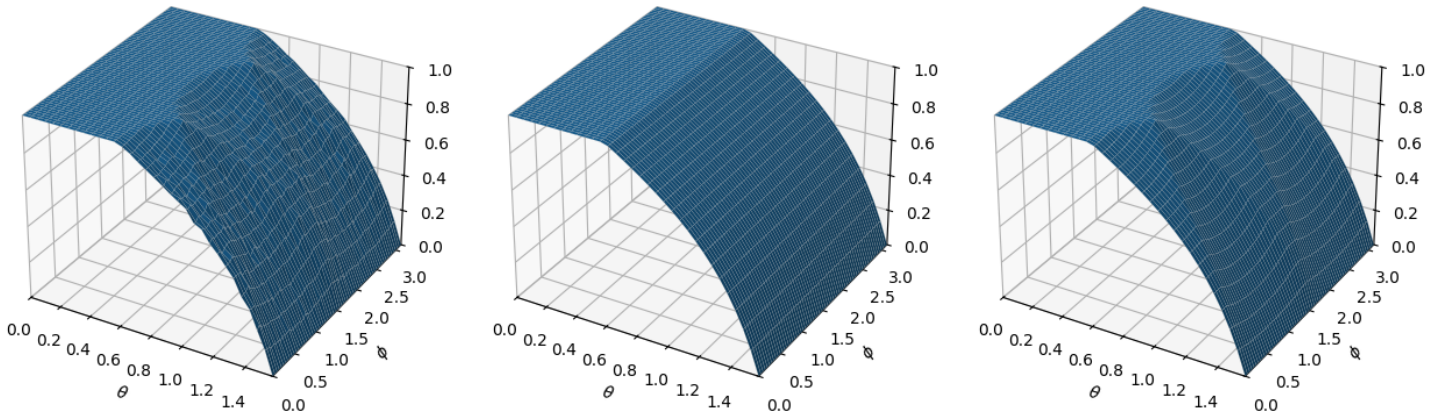

(a) Measured shadowing function of generated mesh with constant peak heights, pyramid density $\rho = 0.6$ peaks$/\mu m^2$ and slant angle $\alpha = 54.7°$, using 4096 samples per data-point.

(b) Shadowing function $G_L(\theta)$ (Equation 4.25) derived by de Haas [4].

(c) Shadowing function $G_1(\theta, \phi)$ (Equation 4.23) derived in this work.

Figure 5.1: Measured and derived shadowing functions.

5.2. BRDF

BRDFs are evaluated using an error metric. In particular, the uncorrected and back-scattering corrected BRDFs are compared. Additionally, an uncorrected BRDF using the shadowing function found by de Haas [4] is also included. As the BRDF needs to be sampled, the error metric is defined as the sum of the absolute differences between sampled path probabilities, using ray-traced measurements as the reference (Figure 5.2). For brevity, histograms of the errors found are also included. Additionally, the total throughput of these BRDFs is compared (Figure 5.3), alongside reference's ray-traced throughput.

Note that, by definition, back-scattering correction has no effect on throughput. Also note the error metric's maximum is two, in the case of fully disjoin path probabilities.

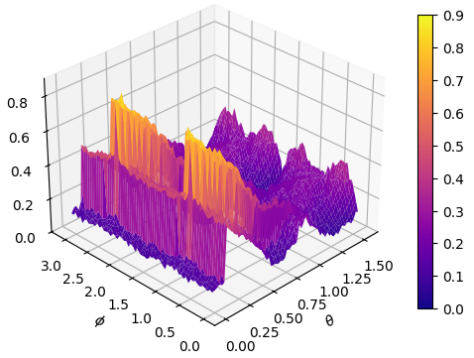

(a) Error metrics of back-scattering uncorrected pyramidic BRDF derived in this work, using shadowing function found by de Haas [4].

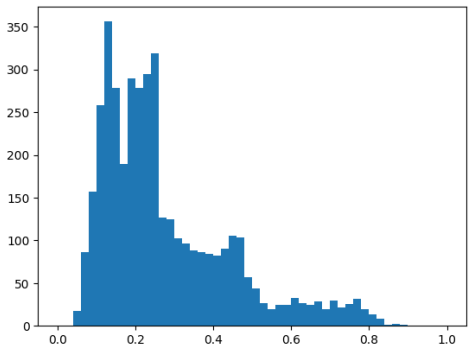

(b) Histogram of error metrics (Figure 5.2a), with an average error of 0.277257.

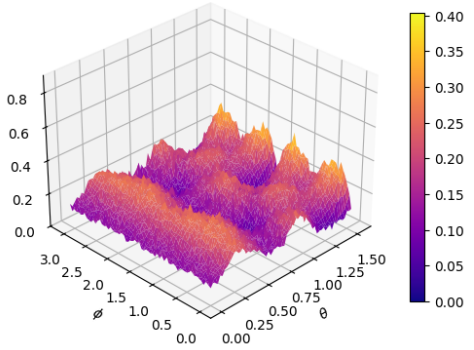

(c) Error metrics of back-scattering uncorrected pyramidic BRDF derived in this work.

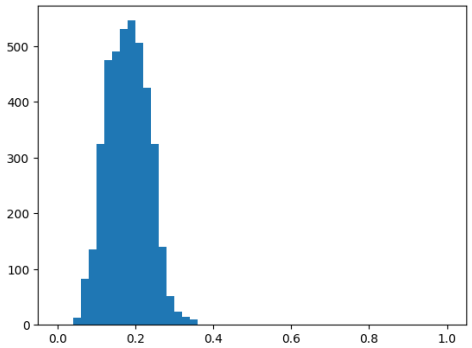

(d) Histogram of error metrics (Figure 5.2c), with an average error of 0.180246.

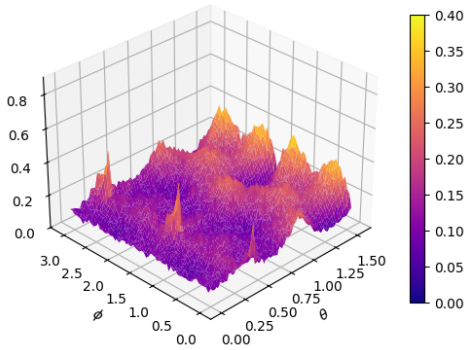

(e) Error metrics of back-scattering corrected pyramidic BRDF derived in this work.

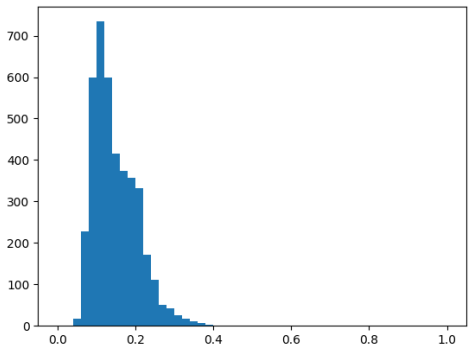

(f) Histogram of error metrics (Figure 5.2e), with an average of 0.148517.

Figure 5.2: BRDFs compared with ray-traced reference mesh using the sum of absolute path probability differences as error metric. The mesh was generated with constant peak heights, pyramid density $\rho = 0.6$ peaks/μm^2 and slant angle $\alpha = 54.7°$, and sampled with 2048 samples per data-point. Path lengths of the mesh and BRDFs were limited to three reflections.

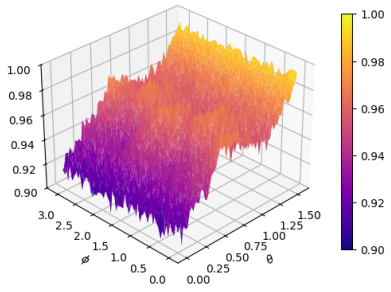

(a) Total throughput of ray-traced reference mesh at a maximum pathlength of three.

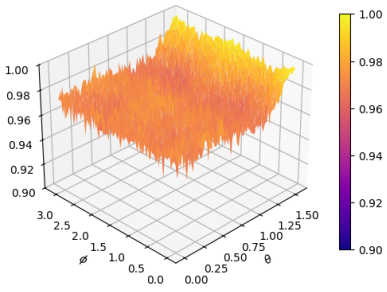

(b) Total throughput of ray-traced reference mesh at a maximum pathlength of four.

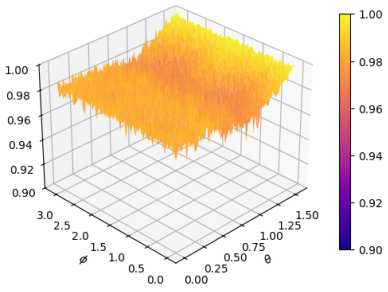

(c) Total throughput of ray-traced reference mesh at a maximum pathlength of six.

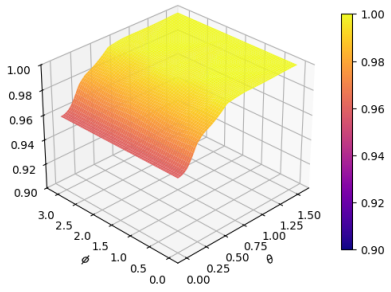

(d) Total throughput of BRDF at a maximum pathlength of three.

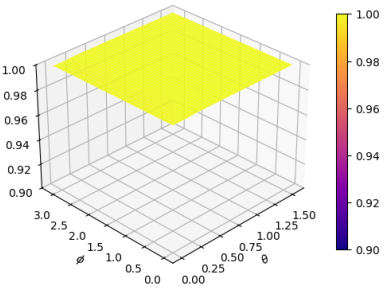

(e) Total throughput of BRDF at a maximum pathlength of four.

Figure 5.3: Total throughput, measured as the sum of path probabilities. Reference mesh was generated with constant peak heights, pyramid density $\rho = 0.6$ peaks$/\mu m^2$ and slant angle $\alpha = 54.7°$, and sampled with 1024 samples per data-point.

5.3. Rendering

In the following, renders comparing the BRDFs with reference results using generated meshes are provided. For this, a modified Cornell box is introduced in Figure 5.4, with its bottom surface as a placeholder, illustrating where the generated reference mesh or the flat BRDF surface are placed. Subsequently, front view (Figure 5.5) and top down (Figure 5.6) render comparisons are provided. Finally, render results at different viewing angles are compared using a skybox procured from Bitterli and MrChimp2313 [36] (Figure 5.7), followed by an impression of the application of coated panel to the roof of a house (Figure 5.8), made available under CC0 license [36]

The discussed BRDFs were implement and rendered using the PBRT-v4 [33] framework. For the back-scattering corrected implementation, path length was limited to three reflections.

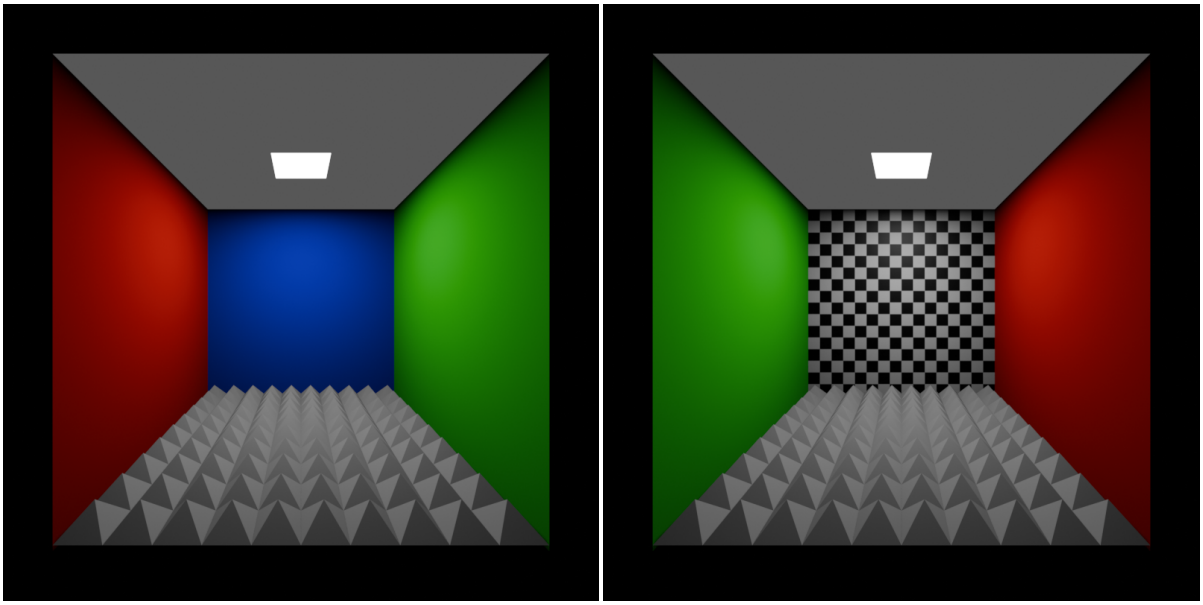

(a) Front view. **(b)** Rear view.

Figure 5.4: Custom Cornell box, with bottom surface as a placeholder for the generated reference mesh and the BRDF surface.

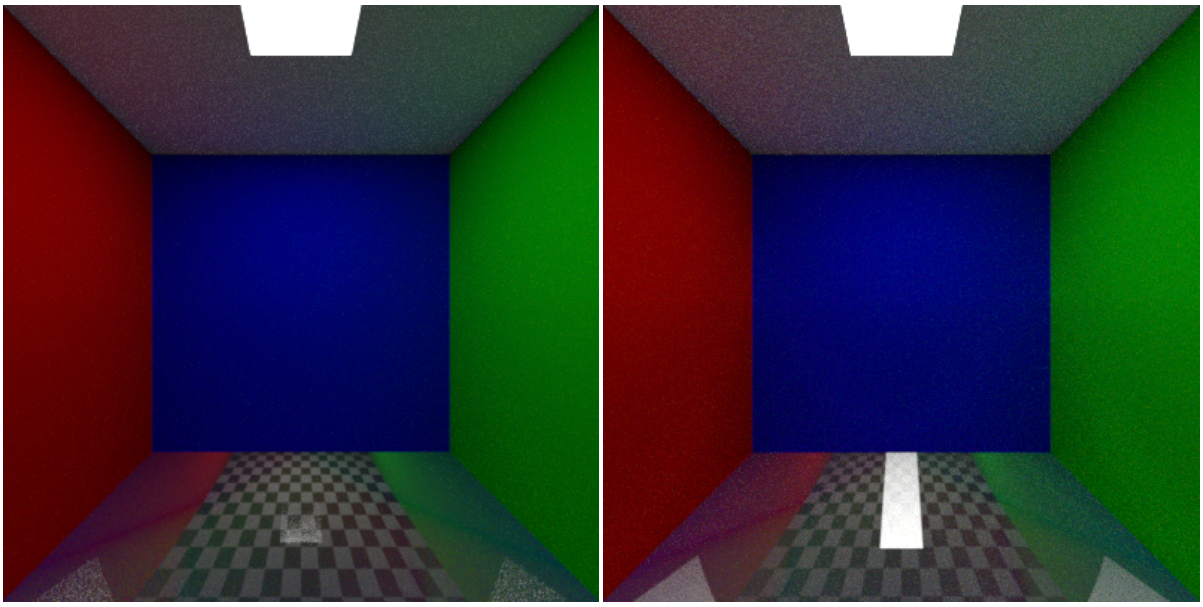

(a) Generated pyramidic mesh, with constant pyramid height, pyramid density $\rho = 0.6$ peaks$/\mu m^2$ and slant angle $\alpha = 54.7°$.

(b) Back-scattering corrected pyramidic BRDF, with path length limited to three reflections.

Figure 5.5: Custom Cornell box front view renders.

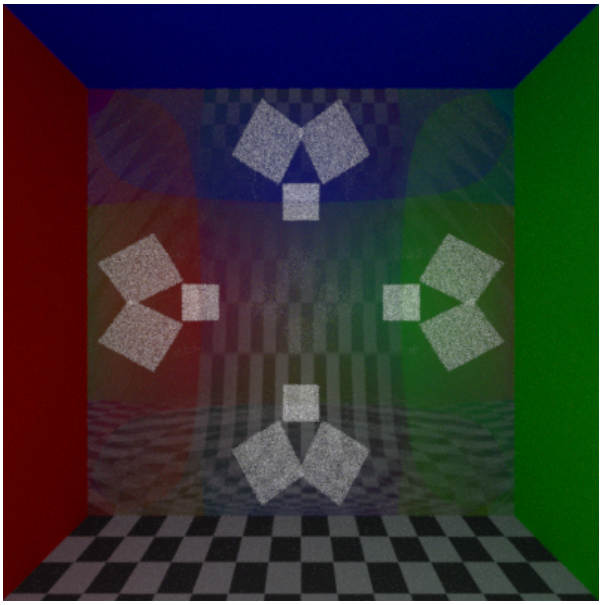

(a) Generated pyramidic mesh, with constant pyramid height, pyramid density $\rho = 0.6$ peaks$/\mu m^2$ and slant angle $\alpha = 54.7^{\circ}$.

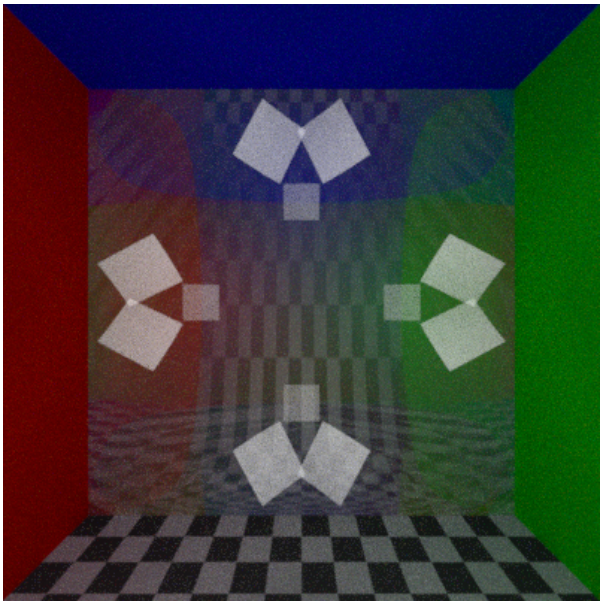

(b) Back-scattering uncorrected pyramidic BRDF.

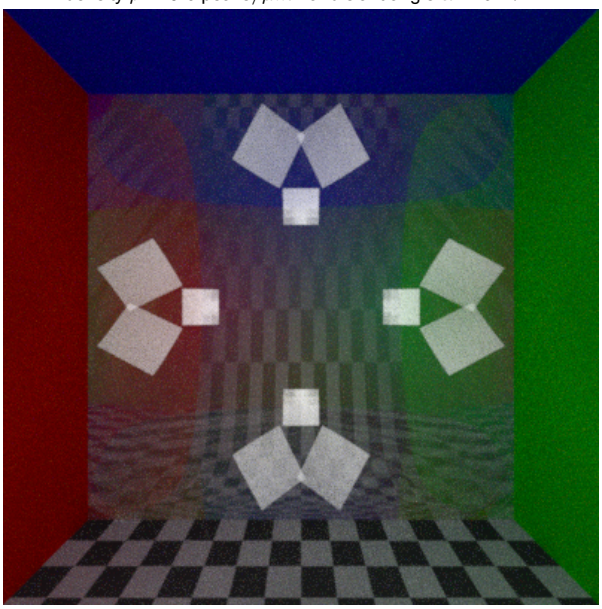

(c) Back-scattering corrected pyramidic BRDF.

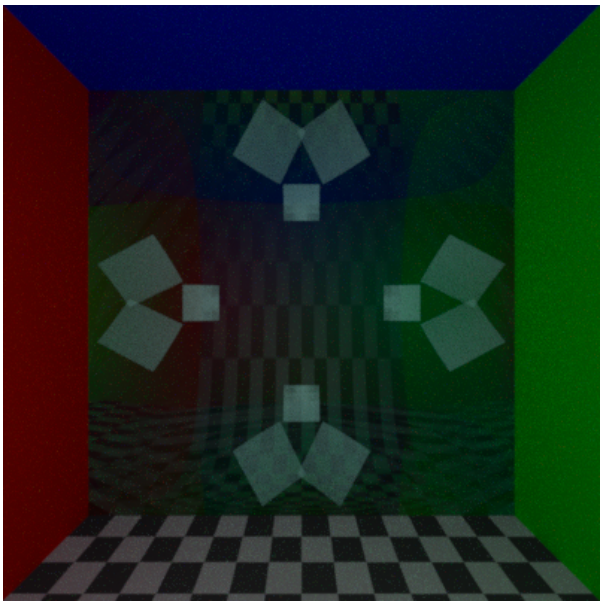

(d) Coated back-scattering corrected pyramidic BRDF.

Figure 5.6: Custom Cornell box top-down view renders. The path length of BRDF renders were limited to three reflections.

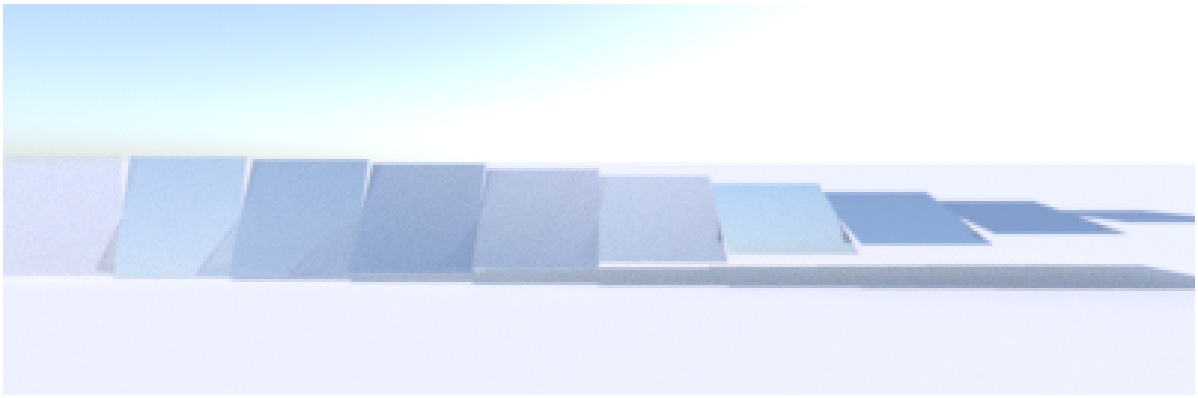

(a) Generated pyramidic mesh, with constant pyramid height, pyramid density $\rho = 0.6$ peaks/μm^2 and slant angle $\alpha = 54.7°$.

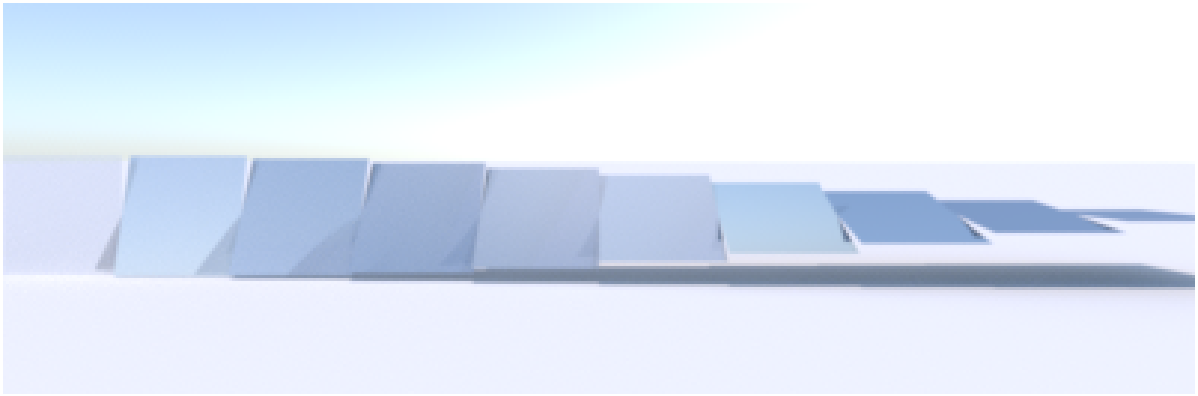

(b) Textured surface BRDF.

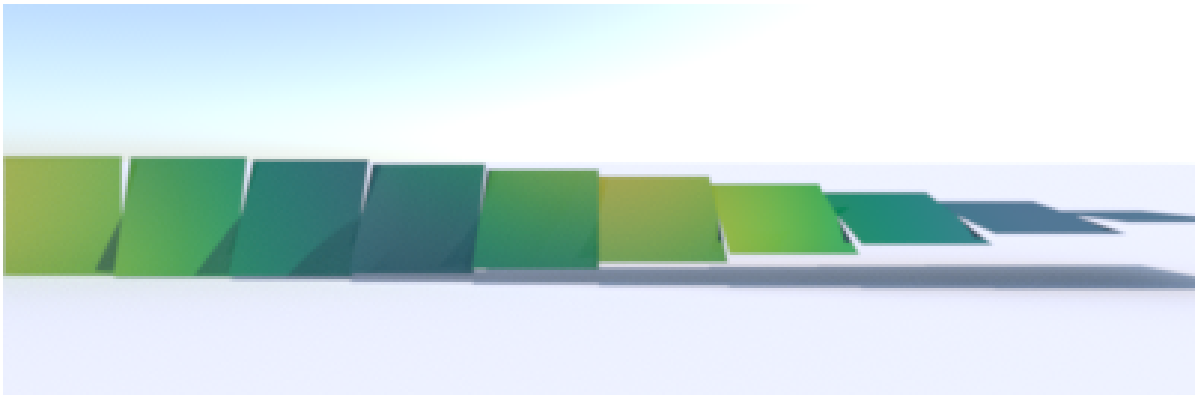

(c) Coated textured surface BRDF.

Figure 5.7: Renders comparing coated, uncoated back-scattering corrected pyramidic BRDFs and generated (uncoated) reference mesh at differing angles using a skybox, procured from Bitterli and MrChimp2313 [36]. BRDF path lengths were limited to three reflections.

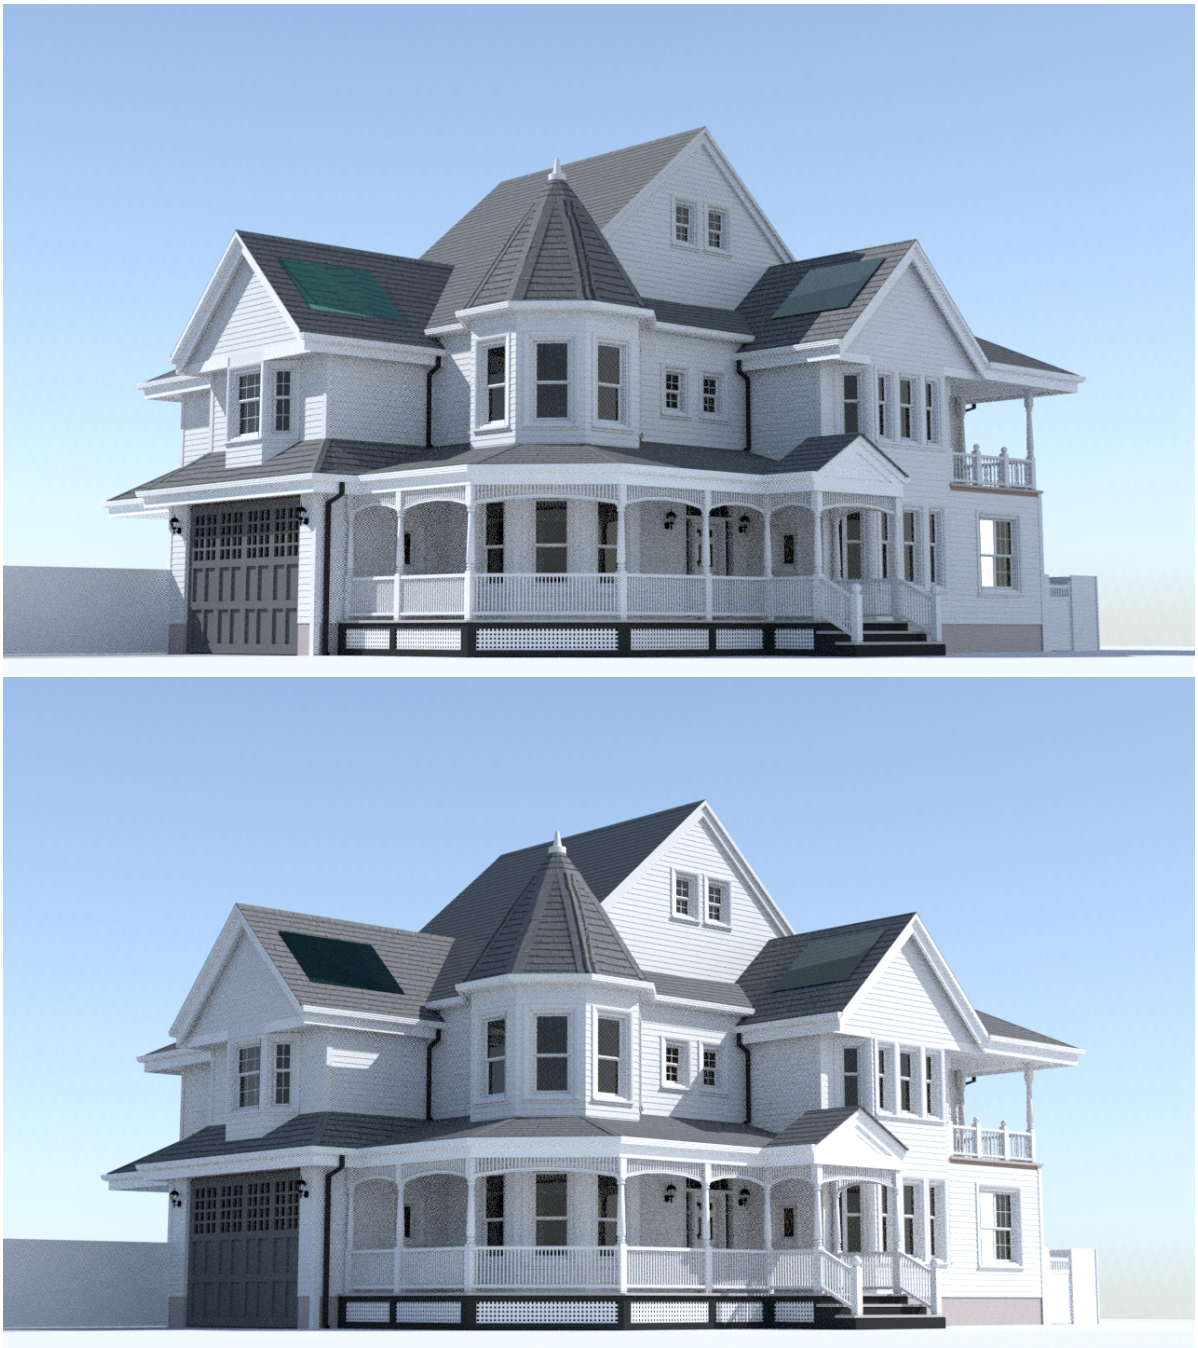

Figure 5.8: Coated back-scattering corrected pyramidic BRDF applied to panels on the roof of a house at different times of day.

5.4. Discussion

As can be seen from Section 5.1, the derived shadowing function accurately follows the reference shadowing function measurements. As discussed in Subsection 4.2.2, the derived shadowing function is also equivalent to the general shadowing function discussed by Heitz [3] and Ashikmin *et al.* [34], within the context of surfaces with uncorrelated normals/shadowing. Although likely correct for the constant height pyramidic surface model in general (Chapter 3), this has not been confirmed using different pyramid densities and slant angles.

The BRDF comparisons provided indicate the use of the derived shadowing function results in better BRDF accuracy than usage of the shadowing function derived by de Haas [4] (Section 5.2), reducing the average error metric from 0.277257 to 0.180245. Additionally, the back-scattering correction proves

effective, reducing the error metric at steep viewing angles, with the average error metric decreasing further to 0.148517. The cause of the error present at medium viewing angles has not been identified, while those at narrow viewing angles are likely due to self-burying (Section 4.6). Although not included in the figure, it is worthwhile to note the use of longer path lengths does not change error metrics significantly. This is expected, as the throughput at longer path lengths is relatively small (on the orders of 1-4%).

It is unclear why the total throughput of ray-traced reference results are lower than the throughput resulting from the BRDFs (Figure 5.3), especially as BRDF results reach their maximum at path lengths of four, while ray-traced reference results remain lower at higher path lengths.

Analysis of the difference in reflections in Figure 5.5 (Figure 5.9) indicate these are likely caused by self-burying. The viewing angle is narrow, with a probability distribution that offsets throughput to a longer path.

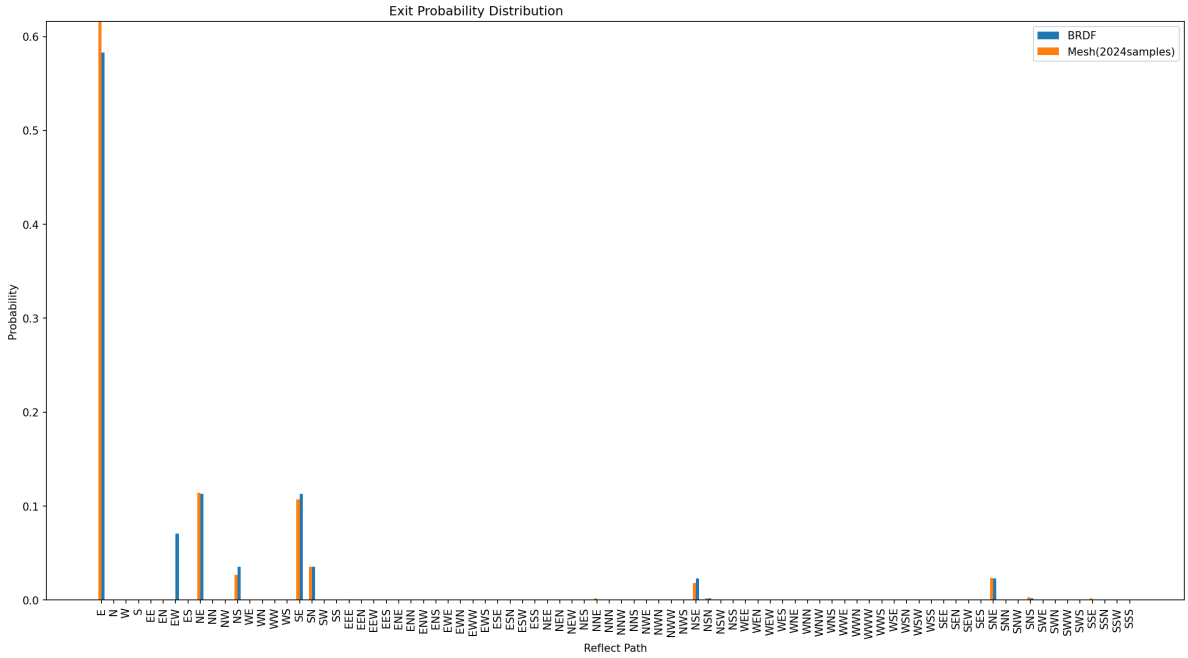

Figure 5.9: Exit-path histogram of disparity in Figure 5.5a and Figure 5.5b, with $\theta = 63°$ and $\phi = 0$.

Top down render results exemplify the effect of back-scattering correction, with highlights better resembling the reference render (Figure 5.6). The difference in brightness is likely due to the fact BRDF's throughput is slightly higher (Figure 5.3), resulting in noticeable differences within the partially-enclosed space. Similar throughput differences do not result in noticeable differences in brightness in the unenclosed skybox renders (Figure 5.7).

It is important to stress the fact the mesh based simulation used as a reference does not perfectly represent physical surfaces, as these are non-ideal, potentially containing (minor) height differences or beveled/rounded pyramid edges. It is unclear what effects these factors would have on the BRDF. Physical effects, such as those related to wave optics, are also unaccounted for. These start playing a more prevalent role at (sub)micron scales, in the case of visible light, close to the scale of the pyramids in question.

6 Conclusion

In this work, a multiple-scattering pyramidic BRDF has been analytically derived and improved with the addition of back-scattering correction. It is characterized solely by peak density ρ and slant angle α, making it applicable to constant height pyramidic surfaces in general, and is simple to reason about, with clear avenues of improvement. The derived shadowing functions has, furthermore, been shown to be accurate, and has turned out to be identical to more general results derived in previous work [3, 34]. Rendering results exhibit a close resemblance to reference measurements at steep angles, although missing corrections at shallow angles and a higher overall throughput result in visible differences.

The use of global microfacet averages for statistical properties such as the shadowing and back-scattering probability have been proven effective, despite their naive weighting. The symmetry has, meanwhile, allowed for efficient tabulation. Finally, despite the involved complexity, correction terms are shown to be effective at improving accuracy, with tabulation making their application relatively inexpensive. Although the tabulation of measurements as a whole is technically feasible as well, parametric models, like the one presented in this work, are applicable more generically. In predictive rendering, the discussed pyramidic BRDF can thus function as a building block in the creation of more complete solar cell BRDFs. As a starting point, an colored optical filter, although simplistic, has been shown to easily incorporate into the model.

6.1. Future Work

There are many avenues for improvement. The most pressing is finding corrections for self-burying behavior at shallow viewing angles. The deviations at intermediate viewing angles and the higher overall throughput also pose topics of interest. Other improvements include the introduction of additional layers in an attempt to model an assembled solar cell, and the generalization of the optical filter coating beyond a fitted model. Additional improvements include material behavior, such as the addition of Fresnel and an examination of transmission and absorption, at both the pyramids' as coating's interfaces.

As shortly mentioned in Section 4.4, several optimizations are still possible. This includes tabulation of path probabilities, the subsequent incorporation of the back-scattering table, the elimination of redundant paths and table regions, and the addition of approximations to speed up calculation.

Additional avenues of further research include the examination of pyramidic surfaces with non-constant height distributions, potentially modeled with multiple instances of the current model, each with a different height, but also includes examining the consequences of edge beveling/rounding on the BRDF, potentially introducing diffuse terms. Future work would especially benefit from the use of physical references, although this may not always be feasible.

Lastly it even may be possible to apply the method described in the work to other surfaces that exhibit comparable uniformly random dispersion of (convex) shapes. This is especially interesting with respect to the back-scattering correction term, as it provides an analytical way of improving comparable multiple-scattering BRDF models.

References

[1] B. Bitterli, *Rendering resources*, https://benedikt-bitterli.me/resources/, 2016.

[2] E. Veach, *Robust Monte Carlo methods for light transport simulation*. Stanford University, 1998.

[3] E. Heitz, "Understanding the masking-shadowing function in microfacet-based brdfs," *Journal of Computer Graphics Techniques*, vol. 3, no. 2, pp. 32–91, 2014.

[4] L. de Haas, "Pyramidic microfacets," M.S. thesis, Delft University of Technology, 2024. [Online]. Available: https://resolver.tudelft.nl/uuid:d27dac74-0c47-433e-97fd-6ab2378ce563.

[5] E. Commission, *Solar energy*, 2025. [Online]. Available: https://energy.ec.europa.eu/topics/renewable-energy/solar-energy_en.

[6] I. E. Agency, *Global energy review 2025*, Paris, Mar. 2025. [Online]. Available: https://www.iea.org/reports/global-energy-review-2025.

[7] T. V. Small, S. D. Butler, and M. A. Marciniak, "Solar cell brdf measurement and modeling with out-of-plane data," *Opt. Express*, vol. 29, no. 22, pp. 35 501–35 515, Oct. 2021. doi: 10.1364/OE.440190. [Online]. Available: https://opg.optica.org/oe/abstract.cfm?URI=oe-29-22-35501.

[8] M. Compean, T. Small, M. Hyde, and M. Marciniak, "Wave optics approach to solar cell brdf modeling with experimental results," *Opt. Express*, vol. 31, no. 16, pp. 26 289–26 300, Jul. 2023. doi: 10.1364/OE.494284. [Online]. Available: https://opg.optica.org/oe/abstract.cfm?URI=oe-31-16-26289.

[9] S. Han, M. Chu, D. P. Pham, S. K. Dhungel, and J. Yi, "Comparison of different approaches to texturing monocrystalline silicon wafers for solar cell applications," *Surface Science*, vol. 748, p. 122 540, 2024.

[10] A. Alasfour, J. Y. Zhengshan, W. Weigand, D. Quispe, and Z. C. Holman, "Sub-micrometer random-pyramid texturing of silicon solar wafers with excellent surface passivation and low reflectance," *Solar Energy Materials and Solar Cells*, vol. 218, p. 110 761, 2020.

[11] M. Moreno *et al.*, "A comparative study of wet and dry texturing processes of c-si wafers for the fabrication of solar cells," *Solar Energy*, vol. 101, pp. 182–191, 2014.

[12] J. Wang *et al.*, "Influence of the textured pyramid size on the performance of silicon heterojunction solar cell," *Solar Energy*, vol. 221, pp. 114–119, 2021.

[13] P. K. Basu, A. Khanna, and Z. Hameiri, "The effect of front pyramid heights on the efficiency of homogeneously textured inline-diffused screen-printed monocrystalline silicon wafer solar cells," *Renewable energy*, vol. 78, pp. 590–598, 2015.

[14] R. Santbergen *et al.*, "Ray-optics study of gentle non-conformal texture morphologies for perovskite/silicon tandems," *Optics Express*, vol. 30, no. 4, pp. 5608–5617, 2022.

[15] P. Procel *et al.*, "Opto-electrical modelling and optimization study of a novel ibc c-si solar cell," *Progress in Photovoltaics: Research and Applications*, vol. 25, no. 6, pp. 452–469, 2017.

[16] Z. Xu, X. Xu, and C. Cui, "Optical functional film with triangular pyramidal texture for crystalline silicon solar cells," *Solar Energy*, vol. 201, pp. 45–54, 2020.

[17] X. Chen *et al.*, "Maskless fabrication of quasi-omnidirectional v-groove solar cells using an alkaline solution-based method," *Chinese Physics B*, vol. 33, no. 1, p. 018 801, 2023.

[18] S. A. Boden and D. M. Bagnall, "Optimization of moth-eye antireflection schemes for silicon solar cells," *Progress in Photovoltaics: Research and Applications*, vol. 18, no. 3, pp. 195–203, 2010.

[19] B. G. Krishna and S. Tiwari, "Bioinspired solar cells: Contribution of biology to light harvesting systems," in *Sustainable Material Solutions for Solar Energy Technologies*, Elsevier, 2021, pp. 593–632.

[20] C.-H. Chang *et al.*, “Combined micro-and nano-scale surface textures for enhanced near-infrared light harvestingin silicon photovoltaics,” *Nanotechnology*, vol. 22, no. 9, p. 095 201, 2011.

[21] R. D. Hussein and D. Ismael, “Review on surface texturing method for solar cell efficiency enhancement,” *Journal of Physical Chemistry and Functional Materials*, vol. 5, no. 1, pp. 1–8, 2022.

[22] A. Soman and A. Antony, “Colored solar cells with spectrally selective photonic crystal reflectors for application in building integrated photovoltaics,” *Solar Energy*, vol. 181, pp. 1–8, 2019.

[23] A. B. Block *et al.*, “Colouring solutions for building integrated photovoltaic modules: A review,” *Energy and Buildings*, vol. 314, p. 114 253, 2024.

[24] J. C. Ortiz Lizcano *et al.*, “Colored optic filters on c-si ibc solar cells for building integrated photovoltaic applications,” *Progress in Photovoltaics: research and applications*, vol. 30, no. 4, pp. 401–435, 2022.

[25] B. Walter, S. R. Marschner, H. Li, and K. E. Torrance, “Microfacet models for refraction through rough surfaces.,” *Rendering techniques*, vol. 2007, 18th, 2007.

[26] E. Heitz, “Sampling the ggx distribution of visible normals,” *Journal of Computer Graphics Techniques (JCGT)*, vol. 7, no. 4, pp. 1–13, 2018.

[27] J. H. Lee, A. Jarabo, D. S. Jeon, D. Gutierrez, and M. H. Kim, “Practical multiple scattering for rough surfaces,” *ACM Transactions on Graphics (TOG)*, vol. 37, no. 6, pp. 1–12, 2018.

[28] Y. Hao, J. Duan, J. Liu, J. Yang, and J. Zhan, “Multiple scattering properties based on modified microsurface pbrdf model,” *IEEE Photonics Journal*, vol. 15, no. 5, pp. 1–8, 2023.

[29] F. Xie and P. Hanrahan, “Multiple scattering from distributions of specular v-grooves,” *ACM Transactions on Graphics (TOG)*, vol. 37, no. 6, pp. 1–14, 2018.

[30] K. E. Torrance and E. M. Sparrow, “Theory for off-specular reflection from roughened surfaces,” *Journal of the Optical society of America*, vol. 57, no. 9, pp. 1105–1114, 1967.

[31] T. Maruyama, J. Bandai, and S. Osako, “Reflection at transparent v-grooved surface,” *Solar energy materials and solar cells*, vol. 64, no. 3, pp. 261–268, 2000.

[32] R. Montes and C. Ureña, “An overview of brdf models,” *University of Grenada, Technical Report LSI-2012-001*, 2012.

[33] M. Pharr, W. Jakob, and G. Humphreys, *Physically Based Rendering: From Theory To Implementation*, 4th ed. The MIT Press, 2023, isbn: 978-0262048026. [Online]. Available: `https://pbr-book.org/4ed/contents` (visited on 04/25/2025).

[34] M. Ashikmin, S. Premože, and P. Shirley, “A microfacet-based brdf generator,” in *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 2000, pp. 65–74.

[35] Maplesoft, a division of Waterloo Maple Inc., *Maple*, version 2024, Waterloo, Ontario. [Online]. Available: `https://www.maplesoft.com/products/Maple/` (visited on 04/25/2025).

[36] B. Bitterli and MrChimp2313, *Victorian style house*, 2016. [Online]. Available: `https://benedikt-bitterli.me/resources/`.

A

Indeterminate Limit Derivation

The limit seen in Section 3.2 (Equation 3.3) is derived below. Note k and m are constants.

$$\lim_{x\to\infty}\left(1+\frac{k}{x}\right)^{mx} = \lim_{x\to\infty} f(x)^{g(x)} = 1^{\infty} \tag{A.1}$$

As the naïve approach results in the indeterminate form 1^{∞}, the function can be rewritten, such that l'Hôpital's rule can be applied to solve for the limit.

$$\lim_{x\to\infty} f(x)^{g(x)} = \lim_{x\to\infty} \exp\left(\frac{\ln(f(x))}{1/g(x)}\right) = \exp\left(\lim_{x\to\infty}\frac{\ln(f(x))}{1/g(x)}\right) \tag{A.2}$$

$$\lim_{x\to\infty}\left(1+\frac{k}{x}\right)^{mx} = \exp\left(\lim_{x\to\infty}\frac{\ln(1+k/x)}{1/mx}\right) \tag{A.3}$$

Note this uses the fact the exponential function is continuous, and for any continuous function $f(y)$, the limit of the composition $f(g(y))$ can be moved inside $f(y)$, assuming the limits for $f(y)$ and $g(y)$ exist:

$$\begin{aligned}\lim_{y\to a} f(g(y)) &= f(\lim_{y\to a} g(y))\\ \lim_{y\to a} \exp(g(y)) &= \exp(\lim_{y\to a} g(y))\end{aligned} \tag{A.4}$$

Since the resulting limit is of the indeterminate form $0/0$, both $\ln(f(x))$ and $1/g(x)$ are differential, and $1/g(x)$ does not equal 0 for $x \neq \pm\infty$, l'Hôpital's rule can be used, as given by Equation A.5, assuming the right-hand-side limit exists.

$$\lim_{y\to a}\frac{f(y)}{g(y)} = \lim_{y\to a}\frac{f'(y)}{g'(y)} \tag{A.5}$$

Evaluating the limit of the fractional part of Equation A.3 by substituting $f(y)$ and $g(y)$ with their respective expressions then yields the limit (Equation A.6).

$$f(y) = \ln\left(1+\frac{k}{y}\right) \qquad g(y) = \frac{1}{my}$$

$$f'(y) = \frac{y}{y+k} * \frac{-k}{y^2} = \frac{-k}{y^2+yk} \qquad g'(y) = \frac{-1}{my^2}$$

$$\lim_{y\to\infty}\frac{\ln(1+k/y)}{1/my} = \lim_{y\to\infty}\frac{kmy^2}{y^2+yk} = \lim_{y\to\infty}\frac{km}{1+k/y} = km \tag{A.6}$$

Filling this into Equation A.3 then yields the final result as given previously (Equation 3.3).

$$\lim_{x\to\infty}\left(1+\frac{k}{x}\right)^{mx} = \exp\left(\lim_{x\to\infty}\frac{\ln(1+k/x)}{1/mx}\right) = e^{km} \tag{A.7}$$

B

Miscellaneous expressions

Note the z-axis in this appendix is defined upwards, opposite to the z-axis used in the rest of this work.

Local Shadowing Function

Below the derivation shown in Equation 4.2.1 is shortly elaborated upon.

$$\hat{\omega} = \begin{bmatrix} \sin(\theta)\cos(\phi) \\ \sin(\theta)\sin(\phi) \\ \cos(\theta) \end{bmatrix} \tag{B.1}$$

$$\begin{aligned} \hat{n}_E &= \begin{bmatrix} \sin(\alpha) \\ 0 \\ \cos(\alpha) \end{bmatrix} \quad \hat{n}_W = \begin{bmatrix} -\sin(\alpha) \\ 0 \\ \cos(\alpha) \end{bmatrix} \\ \hat{n}_N &= \begin{bmatrix} 0 \\ \sin(\alpha) \\ \cos(\alpha) \end{bmatrix} \quad \hat{n}_S = \begin{bmatrix} 0 \\ -\sin(\alpha) \\ \cos(\alpha) \end{bmatrix} \end{aligned} \tag{B.2}$$

$$\begin{aligned} D_E = \hat{\omega} \cdot \hat{n}_E &= \quad \sin(\alpha)\sin(\theta)\cos(\phi) + \cos(\alpha)\cos(\theta) \\ D_W = \hat{\omega} \cdot \hat{n}_W &= -\sin(\alpha)\sin(\theta)\cos(\phi) + \cos(\alpha)\cos(\theta) \\ D_N = \hat{\omega} \cdot \hat{n}_N &= \quad \sin(\alpha)\sin(\theta)\sin(\phi) + \cos(\alpha)\cos(\theta) \\ D_S = \hat{\omega} \cdot \hat{n}_S &= -\sin(\alpha)\sin(\theta)\sin(\phi) + \cos(\alpha)\cos(\theta) \end{aligned} \tag{B.3}$$

$$D = \max(0, D_E) + \max(0, D_W) + \max(0, D_N) + \max(0, D_S) \tag{B.4}$$

Opposing Face Scattering

The opposing-face scattering visualized in Figure 4.9 can be described using matrices M_1 and M_2.

$$\hat{c}_0 = \hat{\omega} \tag{B.5}$$

$$\begin{aligned} \hat{c}_1 &= \underline{M}_1 \hat{c}_0 & \hat{c}_2 &= \underline{M}_2 \hat{c}_1 = \underline{M} \hat{c}_0 \\ \underline{M}_1 &= \begin{bmatrix} \cos(2\alpha) & 0 & -sin(2\alpha) \\ 0 & 1 & 0 \\ -\sin(2\alpha) & 0 & -\cos(2\alpha) \end{bmatrix} & \underline{M}_2 &= \begin{bmatrix} \cos(2\alpha) & 0 & sin(2\alpha) \\ 0 & 1 & 0 \\ \sin(2\alpha) & 0 & -\cos(2\alpha) \end{bmatrix} \\ \underline{M} &= \begin{bmatrix} \cos(4\alpha) & 0 & -sin(4\alpha) \\ 0 & 1 & 0 \\ \sin(4\alpha) & 0 & -\cos(4\alpha) \end{bmatrix} \end{aligned} \tag{B.6}$$

C

Maple

Due to the involved complexity, Maple was used to find the values of certain expressions, resulting from derivations in Subsection 4.5.2 depended upon Maple to find expressions for certain values. Of these, noteworthy results are provided below.[1]

$$\begin{aligned} Y_{mid} &:= simplify(solve(d_1(y,z) = d_2(y,z), y))\ assuming\ real \\ Y_{mid} &:= \frac{2(\cos^2(\alpha) - 1/2)c_{0y}\csc(\alpha)z}{-4\sin(\alpha)\cos^2(\alpha)c_{0z} + 4\cos^3(\alpha)c_{0x} + \sin(\alpha)c_{0z} - 3\cos(\alpha)c_{0x}} \end{aligned} \tag{C.1}$$

Simplified substitution for α, ρ and $\hat{c}_0$ then yields the simple numeric expression Y_{midn}:

$$\begin{aligned} Y_{midn} &:= simplify(subs([...], Y_{mid})) \\ Y_{midn} &:= \frac{z}{0.6731323870 + 2.36049614\cot(\phi)} \end{aligned} \tag{C.2}$$

In a similar vein, numeric expressions for the exit path-distances are provided below, split up into parts due to their size.

$$\begin{aligned} d_{1n} &:= \frac{d_{1num}}{d_{1den}} \\ d_{1num} &:= (-7.0671883z + 9.98134798y)\cos(\phi)\sin(\theta) + (0.43224778z - 0.61048541y)\cos(\theta) \\ d_{1den} &:= (3.04505424\cos^2(\phi) + 1.81997208\cos(\phi)\sin(\phi))\sin^2(\theta) \\ &\quad + (-9.82233945\cos(\phi) - 6.38906164\sin(\phi))\cos(\theta)\sin(\theta) - 3.04505424\cos^2(\theta) \end{aligned} \tag{C.3}$$

$$\begin{aligned} d_{2n} &:= \frac{d_{2num}}{d_{2den}} \\ d_{2num} &:= ((-7.0671883z - 9.98134798y)\cos(\phi)\sin(\theta) + (0.43224778z + 0.61048541y)\cos(\theta)) \\ d_{2den} &:= ((3.04505424\cos^2(\phi) - 1.81997208\cos(\phi)\sin(\phi))\sin^2(\theta)2 \\ &\quad + (-9.82233945\cos(\phi) + 6.38906164\sin(\phi))\cos(\theta)\sin(\theta) - 3.04505424\cos^2(\theta)) \end{aligned} \tag{C.4}$$

The expression T_{dn} for the two integral parts in the numerator of Equation 4.51 are provided below, with d representing d_1 or d_2 respectively, with ξ representing dz_F/dz (Equation 3.15). As these are not directly required for implementation, explicit expressions for P_{BSd}, P_{En} and ξ_n are omitted.

$$\begin{aligned} T_{dn} &:= simplify(P_{BSd} * P_{En} * \xi_n) \\ T_{dn} &:= -1.225283594e^{-1.203167728z^2} * (e^{T_e} - 1) \\ T_e &:= -0.2852077020d(\sin(\theta)\cos(\phi) - 3.510527281\cos(\theta)) * (0.4716113289d\sin(\theta)\cos(\phi) - 0.1660805661d\cos(\theta) + z) \end{aligned} \tag{C.5}$$

[1] Maple is a trademark of Waterloo Maple Inc.