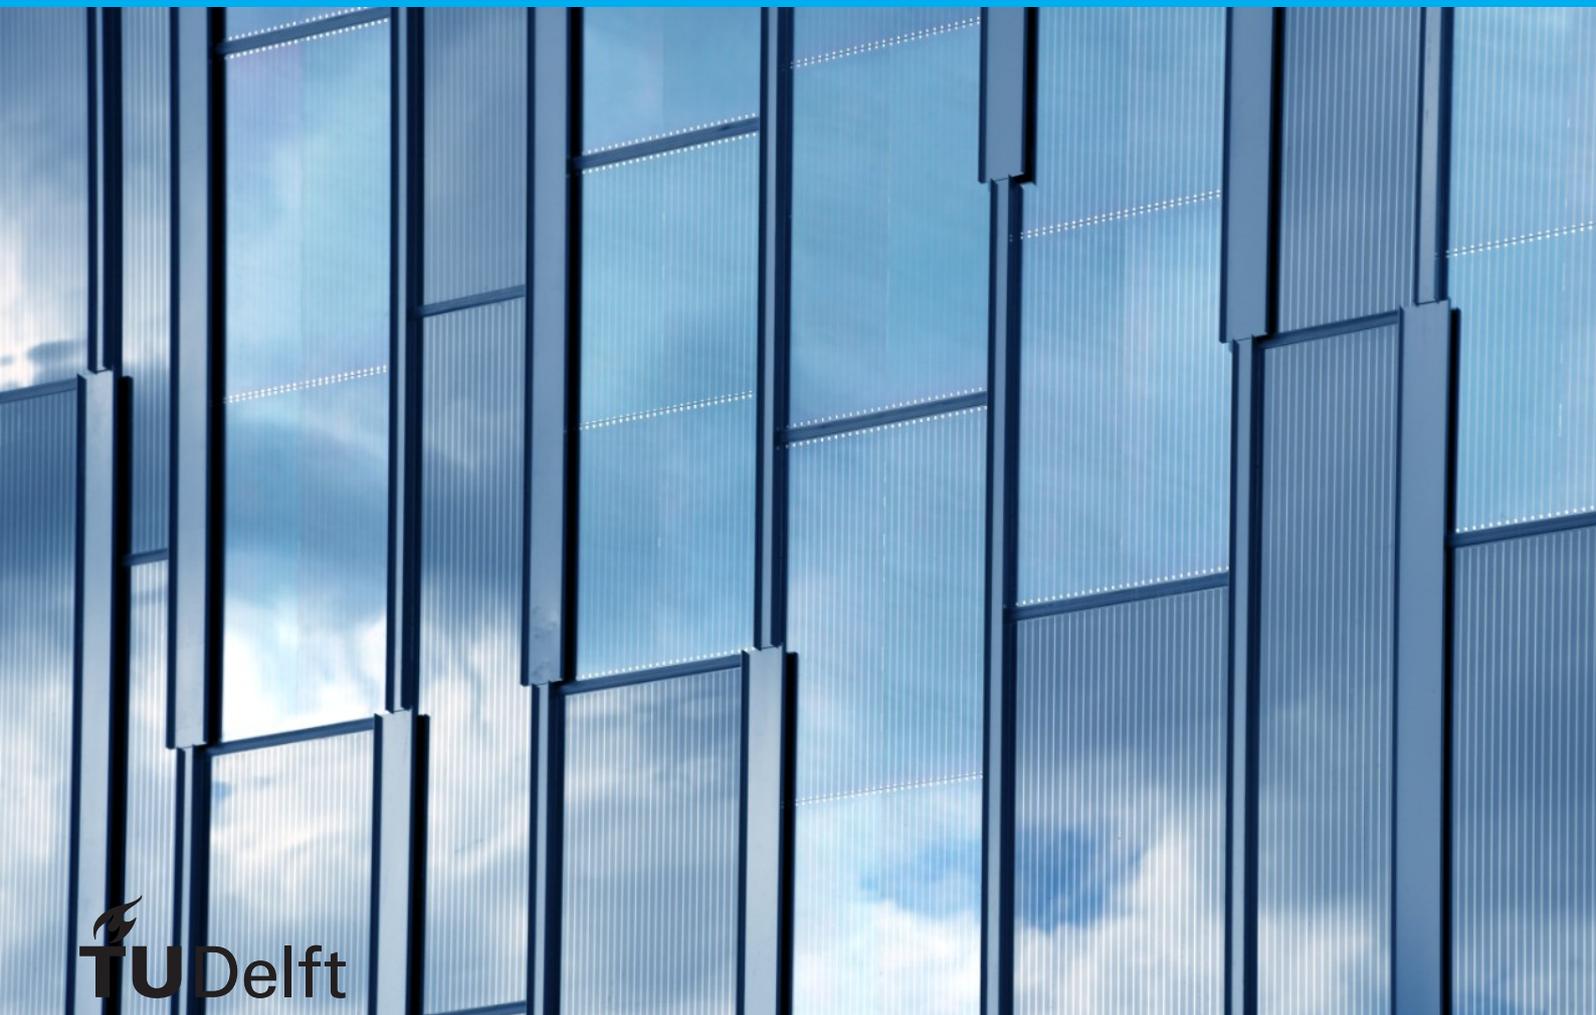


Pyramidic microfacets

Lyanne de Haas

May 2024



Pyramidal microfacets

Fitting a microfacet model to a pyramid-textured solar cell

by

Lyanne de Haas

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Thursday May 30, 2024 at 14:00.

Student number: 4361660
Project duration: November 2022 – May 2024
Thesis committee: Prof. dr. E. Eisemann, TU Delft, supervisor
Dr. R. Marroquim, TU Delft, supervisor
Dr. ir. R. Santbergen, TU Delft

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Cover image source: [6]

Acknowledgements

First and foremost, I would like to thank Ricardo Marroquim and Mark van de Ruit for their patience and for always being willing to help think of new approaches.

I want to thank Elmar Eisemann for providing feedback and ideas at the key moments during this thesis.

I am also grateful to Rudi Santbergen for introducing us to the world of photovoltaics, for explaining its intricacies and for being flexible when the project turned out to be more complicated than expected.

Thank you, Yang Chen, for letting me make use of your work and for taking the time to explain it to me.

And finally, I want to thank my family, for always being willing to listen when things were difficult, and my friends from Paal Centraal, for keeping me sane throughout all of this.

*Lyanne de Haas
Delft, May 2024*

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Scope and contribution of the project	1
1.3	Organisation	2
2	Background	3
2.1	Structure of a solar cell	3
2.2	Microfacet modelling	4
2.2.1	Microfacet BSDF	4
2.2.2	Microfacet distribution function	5
2.2.3	Shadowing-masking function	6
2.2.4	Fresnel	6
2.3	Measuring the BRDF	7
3	Derivation Of Statistical Properties	9
3.1	Normal distribution function	9
3.2	Shadowing-masking function	10
3.2.1	2D, or the V-groove problem	10
3.2.2	3D with fixed azimuth	12
3.2.3	Varying the azimuth	13
3.3	Multiple interactions	14
3.3.1	Exit direction	14
3.3.2	Probability of a sequence	15
4	Modelling And Measured Results	17
4.1	Recreating the surface	17
4.1.1	Assumptions and limitations	18
4.2	Normal distribution function	18
4.3	Shadowing-masking function	19
4.3.1	Validity for all face normals	20
4.4	Multiple interactions	21
4.4.1	Incoming ray at normal incidence	22
4.4.2	Incoming ray with 45 degree elevation	23
4.4.3	Incoming ray with 45 degree elevation, 25 degree azimuth	24
5	Conclusion And Future Work	25
5.1	Conclusion	25
5.2	Future work	25
5.2.1	Implementing the microfacet model	25
5.2.2	Adding layers	25
5.2.3	Exploring other textures	25

Nomenclature

Table 1 summarizes the notations used in this thesis.

\mathbf{g}	Macrosurface normal
\mathbf{f}	Microsurface normal
\mathbf{f}_e	Eastern face normal
\mathbf{f}_n	Northern face normal
\mathbf{f}_w	Western face normal
\mathbf{f}_s	Southern face normal
\mathbf{i}	Incoming ray direction
\mathbf{o}	Outgoing ray direction
\mathbf{r}	Vector with $\theta = \theta_o$ and $\phi = 0$
\mathbf{h}_r	Halfway vector between incoming and outgoing rays for reflection
\mathbf{h}_t	Halfway vector between incoming and outgoing rays for transmission
θ_o	Angle between macrosurface normal and outgoing direction
η_i	Index of refraction of the material on the incident side
η_o	Index of refraction of the material on the transmitted side
$D(\mathbf{v})$	Normal distribution function
$G(\mathbf{i}, \mathbf{o}, \mathbf{v})$	Shadowing-masking function
$G_1(\mathbf{o}, \mathbf{v})$	Smith shadowing-masking function
$\langle a, b \rangle$	Clamped dot product; 0 if $a \cdot b < 0$

Table 1: Definition of symbols

Introduction

1.1. Motivation

The importance of reducing humanity's impact on the environment is an increasingly discussed topic, with both organisations and individuals accepting more and more that they need to make changes. The EU aims to reduce greenhouse gas emissions by 55% in 2030, and wants to move towards being climate neutral in 2050 [3]. As one strategy in working towards this goal, clean energy techniques are being developed. One such technique is the deployment of Building Integrated Photovoltaic (BIPV) systems. These systems replace materials in a building with solar cells, allowing them to be seamlessly integrated into walls or roofs. Though this does not change the shape of the building, its appearance is altered.

At the moment, the options for the appearance of these BIPV systems are limited. Giving architects more options in this regard would better enable their use in a variety of scenarios.

Different options to improve colour versatility have been explored, involving techniques such as dyeing the protective glass covering the solar cells or colouring the cells themselves. Another option, which can be applied both to the front glass or to the solar cells themselves, is the use of an optical colour filter (OF). Broadly speaking, these optical filters selectively transmit a wavelength-dependent percentage of light. By adjusting the OF's composition, these percentages can be changed, thus altering their colour.

In 2021, the TU Delft Photovoltaic Materials and Devices (PVMD) group proposed a method for the use of coloured optical filters on crystalline silicon (c-Si) solar cells [10], those being one of the more popular types of cells in use today. These filters colour the cells in any desired shade and make them more aesthetically attractive for use in a city environment, as well as easier to blend into the environment, but do not severely impact the performance of the cells. To convince architects to make use of this technology, it is important they are able to get a good impression of what it will look like in use. To this end, the PVMD group would like to have access to a method to realistically render these solar cells.

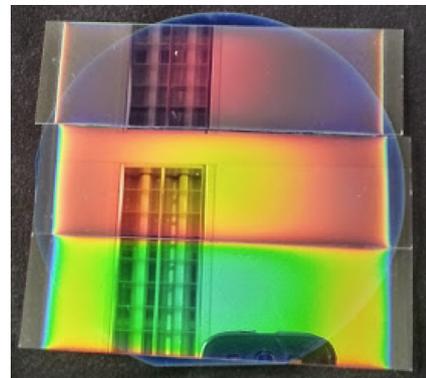


Figure 1.1: A solar cell overlaid with samples of the optical filters designed by Ortiz Lizcano et al. Image source: [10]

1.2. Scope and contribution of the project

The c-Si solar cells the filter is to be applied to has a microscopic texture, consisting of many overlapping pyramids, which forces light to interact with the surface multiple times. It is desirable to avoid recreating this microsurface in detail as this would be computationally expensive. The pyramidal surface of the c-Si cells will be analysed, and statistical functions that describe its geometric properties will be derived. These functions can be directly input into a so-called microfacet model that can approximate the way

the microsurface scatters light, allowing the user to render its appearance without the computational demands of exact recreation. The functions to be derived describe the orientation of the facets of the surface, as well as the visibility of the facets from any direction.

The microfacet model that will be adapted explicitly ignores multiple interactions, which are an integral part of the function of the microsurface and should be included. The effect of multiple interactions will therefore also be analysed and included into the model.

The resulting functions will be verified through measurements created using a digital gonioreflectometer and simulated pyramidal surfaces.

1.3. Organisation

Chapter 2 discusses previous work related to this thesis. In chapter 3, the functions needed to create the microfacet model are mathematically derived. Finally, chapter 4 compares the functions found in the previous chapter to those measured using a digital gonioreflectometer and discusses the accuracy of the model.

2

Background

This chapter will explore previous work related to the subject of this thesis.

2.1. Structure of a solar cell

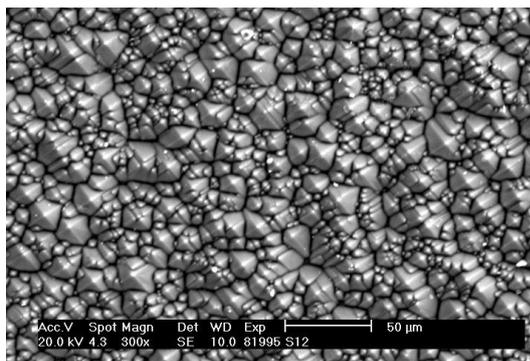


Figure 2.1: c-Si solar cell surface as captured by electron microscope. Image source: [11]

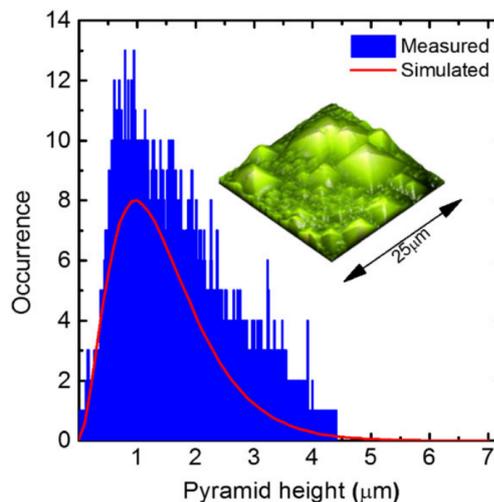


Figure 2.2: Measured (blue) and fitted (red) distribution of pyramid heights on a solar cell surface (inset). Image source: [14]

The base cell described by Ortiz Lizcano et al. [10] consists of a crystalline silicon (c-Si) solar cell. This is the most commonly used type of cell due to its low production cost and high conversion efficiency [14], meaning, compared to other types of solar cells, they produce more electricity from the same amount of light.

A pyramid texture is added to the front side of these C-Si solar cells, an example of which is pictured in Figure 2.1. This is done through a so-called wet-alkaline etching process [12], where, as the name implies, liquid alkaline chemicals are applied to the surface. These chemicals react with the silicon, eating away at it. The rate at which this process occurs varies locally, resulting in the distinctive pyramid texture. The silicon has several crystal orientations, which etch down at very different speeds. The angle between the orientation of the silicon wafer itself and the orientation that is slowest to etch is 54.7° [15]. The pyramid sides follow this slowest orientation and have a fixed angle of 54.7° [14]. These pyramids are randomly distributed on the surface, 'growing' downwards from a seed point at their tips. The pyramids vary in height, following a distribution measured by Procel et al. [14]. The distribution as found in this paper can be seen in Figure 2.2. The data plotted in this figure was obtained from P. Procel for use in this thesis.

Adding a texture to the cells reduces the amount of light reflected and thus increases the conversion

efficiency by increasing the light absorbed into the cell [14]. This effect is demonstrated in Figure 2.3. Each interaction allows a percentage of the light to be transmitted into the surface. When a second bounce occurs, the same percentage of the light that was not transmitted on the first bounce, and would have left the surface if no texture was present, will enter the material and be converted into electricity. The purpose of using this specific pyramidic texture is to always force at least two interactions, but more can occur, each time reducing the reflected energy by a percentage.

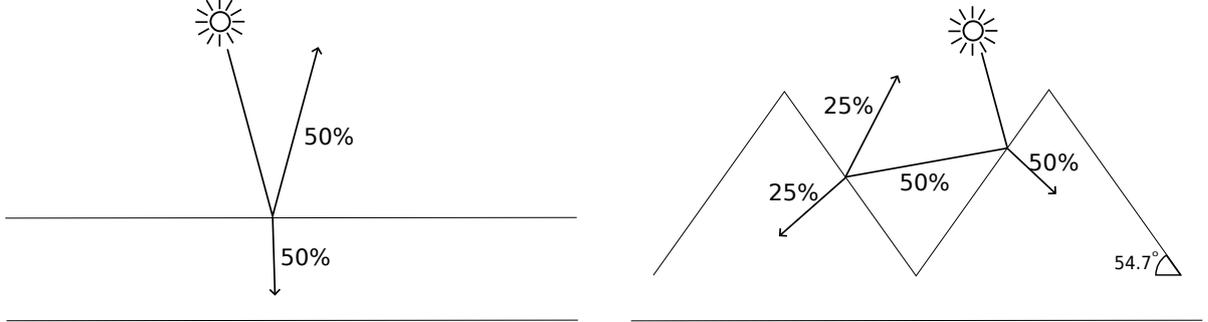


Figure 2.3: Demonstration of pyramidic texture forcing multiple interactions with a surface with a transmittance of 0.5. On a flat surface (left), 50% of the light is reflected. On a surface with pyramids (right) the light is reflected again and only 25% escapes.

2.2. Microfacet modelling

Perfectly smooth surfaces are rare. Usually, a surface has some sort of texture, which could vary from something as small as scratches on a desk to something as rough as cat fur. When these textures are too small to be seen directly, they are referred to as microspheres. When trying to realistically render an image, it is important to capture these imperfections, as they influence the appearance of a material. Explicitly incorporating the microsphere is computationally expensive, as it would require a very high-definition mesh. The microfacet model developed by Torrance and Sparrow [17] allows the bidirectional scattering distribution function (BSDF) of microfaceted surfaces to be statistically modelled, thus speeding up computation considerably by allowing the renderer to sample the model BSDF instead of the exact microsphere. A limitation of the model is the assumption that any ray of light hits the surface exactly once. Since our surface is designed to force multiple bounces, this will need to be included separately. The Torrance-Sparrow model assumes a microsphere can be fully described by two statistical measures, a microfacet distribution function D (2.2.2) and a shadowing-masking function G (2.2.3). It also uses a Fresnel term F (2.2.4).

2.2.1. Microfacet BSDF

The microfacet model is used to obtain a bidirectional scattering distribution function (BSDF). The BSDF is a function that takes an incoming direction \mathbf{i} , outgoing direction \mathbf{o} and macrosurface normal \mathbf{g} as input, and outputs the ratio of energy transferred from the incoming to the outgoing direction. It is found by summing its two components, the bidirectional reflectance distribution function (BRDF) and bidirectional transmittance distribution function (BTDF), denoting the reflection and transmission respectively [18].

$$f_s(\mathbf{i}, \mathbf{o}, \mathbf{g}) = f_r(\mathbf{i}, \mathbf{o}, \mathbf{g}) + f_t(\mathbf{i}, \mathbf{o}, \mathbf{g}). \quad (2.1)$$

The BRDF and BTDF both depend on the statistical functions D and G , and on the Fresnel term F . For the BTDF, we also need the index of refraction of the materials on the incident (η_i) and transmitted (η_o) side. They are defined as follows [18]:

$$f_r(\mathbf{i}, \mathbf{o}, \mathbf{g}) = \frac{F(\mathbf{i}, \mathbf{h}_r)G(\mathbf{i}, \mathbf{o}, \mathbf{h}_r)D(\mathbf{h}_r)}{4|\mathbf{i} \cdot \mathbf{g}||\mathbf{o} \cdot \mathbf{g}|}, \quad (2.2)$$

$$f_t(\mathbf{i}, \mathbf{o}, \mathbf{g}) = \frac{|\mathbf{i} \cdot \mathbf{h}_t||\mathbf{o} \cdot \mathbf{h}_t|}{|\mathbf{i} \cdot \mathbf{g}||\mathbf{o} \cdot \mathbf{f}_n|} \frac{\eta_o^2(1 - F(\mathbf{i}, \mathbf{h}_t))G(\mathbf{i}, \mathbf{o}, \mathbf{h}_t)D(\mathbf{h}_t)}{(\eta_i(\mathbf{i} \cdot \mathbf{h}_t) + \eta_o(\mathbf{o} \cdot \mathbf{h}_t))^2}. \quad (2.3)$$

2.2.2. Microfacet distribution function

The microfacet distribution function $D(\mathbf{v})$, also known as the normal distribution function (NDF), describes the distribution of microfacet normals on the surface. For any normal \mathbf{v} and given an infinitesimal solid angle $d\omega_v$ centred on \mathbf{v} and an infinitesimal macrosurface area dA , $D(\mathbf{v})d\omega_m dA$ is the total area of the microsurface corresponding to dA that has a normal within $d\omega_v$ [18]. Simply said, D is a density function that for any given vector, describes the probability of any infinitesimal area of the microsurface having that vector as a normal.

There are some properties that should hold for any valid NDF.

- D is a density function and, therefore, must be non-negative:

$$0 \leq D(\mathbf{v}) \leq \infty \quad (2.4)$$

- Over the whole surface, the area of the microsurface must be equal to or greater than the corresponding area of the macrosurface.

$$1 \leq \int D(\mathbf{v})d\omega_v \quad (2.5)$$

This can be intuited in 2D as the macrosurface spanning the most direct route from A to B, where the microsurface deviates from this path and, therefore, has a greater length. This concept is visualised in Figure 2.4.

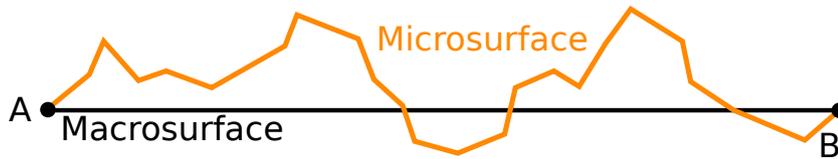


Figure 2.4: Illustration of a micro- and macrosurface, showing that the macrosurface can be viewed as the most direct path from point A to point B, while the microsurface is a longer path between the points.

- The projected area of the microsurface is the same as the projected area of the macrosurface for any direction \mathbf{o} .

$$(\mathbf{o} \cdot \mathbf{g}) = \int D(\mathbf{v})(\mathbf{o} \cdot \mathbf{v})d\omega_m. \quad (2.6)$$

The left hand side of equation 2.6 is the area of the macrosurface, the right hand side is that of the microsurface. This concept is visualised in Figure 2.5.

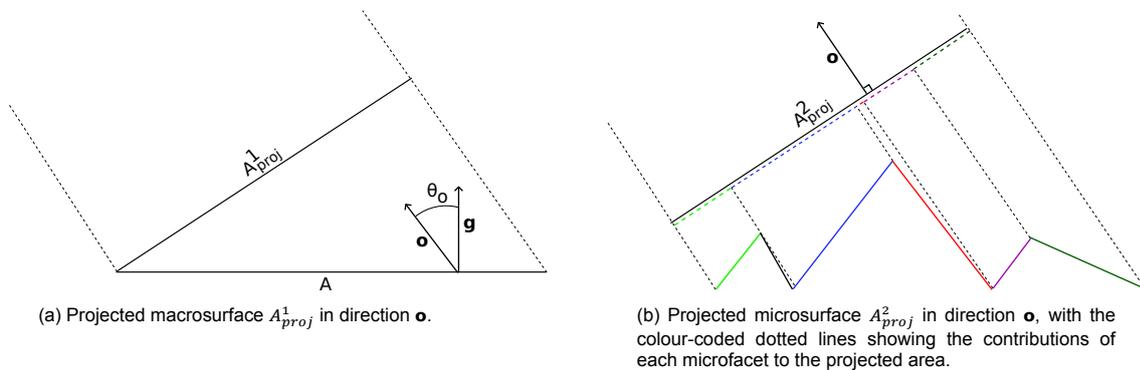


Figure 2.5: Projections of a macro- and microsurface in the same direction \mathbf{o} . The projections of both surfaces have the same area.

2.2.3. Shadowing-masking function

The second function, $G(\mathbf{i}, \mathbf{o}, \mathbf{v})$, describes, for the part of the surface that has normal \mathbf{v} , what percentage is simultaneously visible in directions \mathbf{i} and \mathbf{o} [18]. This function shows when peaks in the surface occlude other parts of the surface. Typically, one would expect the most occlusion to occur at shallow angles.

A valid function for G needs to have the following properties:

- As G denotes a percentage, it should be a fraction between zero and one:

$$0 \leq G(\mathbf{i}, \mathbf{o}, \mathbf{v}) \leq 1 \quad (2.7)$$

- G describes the simultaneous visibility in directions \mathbf{i} and \mathbf{o} , so these vectors should be interchangeable:

$$G(\mathbf{i}, \mathbf{o}, \mathbf{v}) = G(\mathbf{o}, \mathbf{i}, \mathbf{v}) \quad (2.8)$$

- When viewing the front of the macrosurface, we should never be able to see the back of any microsurface facet and vice versa; in these cases, G should be set to zero. The dot product of a vector with a surface normal is positive if and only if the vector is on the front side of the surface.

$$G(\mathbf{i}, \mathbf{o}, \mathbf{v}) = 0 \quad \text{if } (\mathbf{i} \cdot \mathbf{v})(\mathbf{i} \cdot \mathbf{g}) \leq 0 \quad (2.9)$$

$$\text{or } (\mathbf{o} \cdot \mathbf{v})(\mathbf{o} \cdot \mathbf{g}) \leq 0 \quad (2.10)$$

Smith approximation

In general, the exact function for $G(\mathbf{i}, \mathbf{o}, \mathbf{m})$ is difficult, if not impossible, to find. In 1967, the Smith shadowing-masking approximation method was proposed [16]. This approximation can be used in those cases where the spatial distribution of the microfacets is not correlated with their heights. Instead of directly finding one term with both the incident and scattering vectors, this method approximates it as the product of two shadowing terms G_1 with a single direction each,

$$G(\mathbf{i}, \mathbf{o}, \mathbf{m}) \approx G_1(\mathbf{i}, \mathbf{m})G_1(\mathbf{o}, \mathbf{m}), \quad (2.11)$$

for which G_1 can be extracted from the microfacet distribution.

V-grooves

Most microsurface models assume a surface with facets in many different orientations. A surface with only a few possible face orientations, such as the one discussed in this paper, is relatively rare. There is, however, one commonly discussed texture which is similarly structured and will be of use to us. This is the V-groove surface ([1, 4, 7]), which is illustrated in Figure 2.6. As suggested by the name, this surface consists of grooves running along one axis, meaning the normal distribution function only has two possible values. We are especially interested in the derivation of the shadowing-masking function for V-grooves, as a jumping-off point for a derivation for our surface. There are two works that have served as the main inspiration, providing different proofs for the same function. Blinn [1] uses a trigonometric proof, employing the law of sines and complementary angles. Heitz [7], meanwhile, builds upon the equality between the area of the macro- and microsurface as established in equation 2.6. This proof will be discussed in section 3.2.1.

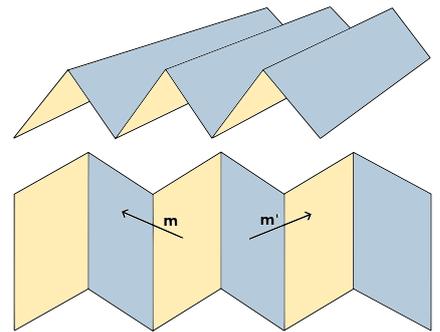


Figure 2.6: Illustration of a V-groove surface texture. The surface has two normals, \mathbf{m} and \mathbf{m}'

2.2.4. Fresnel

The final component needed for the BSDF is the Fresnel term F [5]. All light waves can be described as two linear polarisation components: s , perpendicular to the plane of incidence (the plane containing both the surface normal and the light direction vector), and p , parallel to it. These components will be transmitted into the surface or reflected from it in different ratios. The four Fresnel equations describe the ratio of transmitted and reflected light for each of the polarisation components [13]. They depend only on the refractive indices of the media at the surface, which are material-dependent, and on the angle between the incoming light and the surface normal. In general, more light will be transmitted at normal incidence and when nearing grazing angles, reflection very quickly rises to close to one. This effect is responsible for, for example, the reflections we might see on water in the distance, like in Figure 2.7.



Figure 2.7: An example of water being reflective when viewed at a grazing angle, due to the Fresnel effect. Image source: [19]

Under certain constraints, the formulation can be simplified. For example, dielectrics do not have a complex component to their refractive index thereby allowing a simpler expression, or the two polarization components can be combined to yield an approximation for unpolarized light. Because the Fresnel term does not depend on the geography of the microsurface, it is outside the scope of this thesis and will not be investigated further here.

2.3. Measuring the BRDF

To verify the model deduced for our microsurface, we can measure the normal distribution function and shadowing-masking function of simulated surfaces using the digital gonioreflectometer developed by Y. Chen and used in [2]. A physical gonioreflectometer (Figure 2.8) measures the BRDF of a surface by rotating a light source and sensor around it, and measuring the reflectance for different combinations of positions. Chen's digital version allows the user to input a height field into his software, and then tracks in which directions rays sent into the surface are scattered.

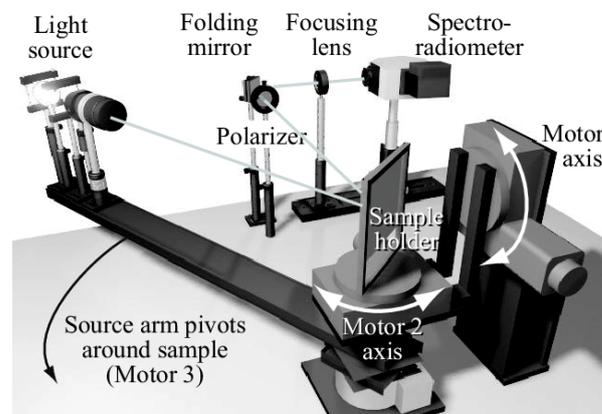


Figure 2.8: An example of a physical gonioreflectometer. The light source and sample can both pivot, to measure the reflectance in all directions. Image source: [9]

3

Derivation Of Statistical Properties

In this chapter, expressions for the normal distribution function and shadowing-masking function of the pyramidic microsurface will be derived, based on similar derivations for related surfaces. Additionally, an analysis of the effect of multiple interactions on the energy transfer will be performed.

3.1. Normal distribution function

Since the surface consists exclusively of pyramids which all have the same angle and orientation, we expect D to only contain a few nonzero values. A single pyramid has four faces, each of them the same size. Thus, if we were to consider a grid with adjacent pyramids of the same size aligned such that none overlap, out of the facets on the grid, 25% would be facing north, 25% east, 25% south, and 25% west. In other words, D will consist of four Dirac deltas for the four normals. If the surface were to not be fully covered in pyramids, but still had some empty space in between them, we would expect to see another Dirac delta at zero elevation and azimuth, corresponding to the areas with a flat surface. In practice, however, this is not something that would occur on a c-Si solar cell.

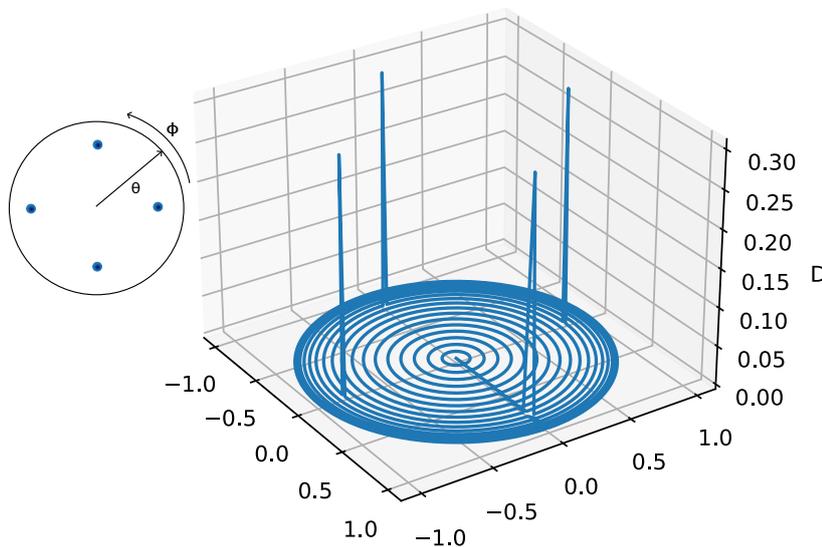


Figure 3.1: Theoretical normal distribution function D , plotted such that rotation around the origin denotes azimuth and distance from the origin denotes elevation. The plot shows four peaks, one for each of the possible face orientations.

We do not have such a neat surface. Instead, pyramids can have varying heights and can overlap one another. This overlap will cause part of the overlapping faces to be obscured, thus causing their peaks in D to be smaller. In our case, because the pyramids are uniformly randomly scattered across

the surface, we expect D to look the same as for the neatly aligned grid. The chance of a pyramid overlapping another on any one side is the same as it overlapping on any other side. Over a large area, then, the expectation is that these overlaps will even out and result in the same four equal Dirac deltas.

To find the value of the peaks in D , we refer to the work by Heitz [7]. In this paper, the following value is derived for D for a V-groove surface with normals \mathbf{m} and \mathbf{m}' :

$$D(\mathbf{v}) = \frac{1}{2} \frac{\delta_{\mathbf{m}}(\mathbf{v})}{\mathbf{m} \cdot \mathbf{g}} + \frac{1}{2} \frac{\delta_{\mathbf{m}'}(\mathbf{v})}{\mathbf{m}' \cdot \mathbf{g}}. \quad (3.1)$$

The Dirac delta functions ensure D is only nonzero at two points, corresponding to the two surface normals. The division by $\mathbf{m} \cdot \mathbf{g}$ is needed to ensure equation 2.6 (the equality between the projected micro- and macrosurface area) holds. The microsurface has a larger total area than the macrosurface, and if D did not include this denominator, its integral over the surface would equate to a number larger than 1.

In an almost identical manner, we can show the value of D for a pyramidal surface is

$$D(\mathbf{v}) = \sum_{\mathbf{i} \in \{\mathbf{f}_n, \mathbf{f}_e, \mathbf{f}_s, \mathbf{f}_w\}} \frac{1}{4} \frac{\delta_{\mathbf{i}}(\mathbf{v})}{\mathbf{i} \cdot \mathbf{g}}. \quad (3.2)$$

To show this is the correct formula for D , we use the property of D that $\int_{\Omega} D(\mathbf{v}) \langle \mathbf{g}, \mathbf{v} \rangle d\mathbf{v} = 1$, as given in [18]. We can show that this equality holds when the value in Equation 3.2 is inserted.

$$\begin{aligned} \int_{\Omega} D(\mathbf{v}) \langle \mathbf{g}, \mathbf{v} \rangle d\mathbf{v} &= \int_{\Omega} \langle \mathbf{g}, \mathbf{v} \rangle \sum_{\mathbf{i} \in \{\mathbf{f}_n, \mathbf{f}_e, \mathbf{f}_s, \mathbf{f}_w\}} \frac{1}{4} \frac{\delta_{\mathbf{i}}(\mathbf{v})}{\mathbf{i} \cdot \mathbf{g}} d\mathbf{v} \\ &= \frac{1}{4} \frac{\mathbf{f}_e \cdot \mathbf{g}}{\mathbf{f}_e \cdot \mathbf{g}} + \frac{1}{4} \frac{\mathbf{f}_w \cdot \mathbf{g}}{\mathbf{f}_w \cdot \mathbf{g}} + \frac{1}{4} \frac{\mathbf{f}_n \cdot \mathbf{g}}{\mathbf{f}_n \cdot \mathbf{g}} + \frac{1}{4} \frac{\mathbf{f}_s \cdot \mathbf{g}}{\mathbf{f}_s \cdot \mathbf{g}} \\ &= \frac{1}{4} + \frac{1}{4} + \frac{1}{4} + \frac{1}{4} = 1. \end{aligned}$$

3.2. Shadowing-masking function

To derive the properties of the shadowing-masking function for our surface, we will progress through a number of cases, increasing in complexity. For each case we will analyse how varying the incident angle impacts the area of the pyramids that is masked.

3.2.1. 2D, or the V-groove problem

We start with a 2D scenario, leaving the azimuth out of consideration. We consider two pyramids placed next to each other, as in Figure 3.2. Since it is in 2D, this case is identical to that of a V-groove problem. This problem has been solved for V-grooves [1, 7], and the same approach can be used here.

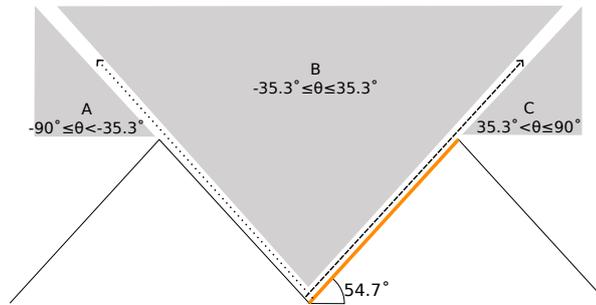


Figure 3.2: The different cases for G applied to the orange face in the case of 2D pyramids. In section A masking occurs, in section B there is no masking, and in section C the face is backfacing.

The facet marked in orange is the one for which we will analyse the shadowing. There are three distinct regions:

- If θ_o is larger than 35.3° , so in section C in Figure 3.2, the light is coming from behind the facet under consideration and so none of it will be visible. This means G_1 will have a value of 0 in that case.
- If θ_o is between 35.3° and -35.3° , corresponding with section B in the figure, the entire facet is visible and $G_1(\theta)$ will have a value of 1.
- The most complex section of the function is section A in Figure 3.2. Here, the left pyramid casts a shadow on our facet. The proportion of the visible surface is described by the following formula:

$$G_1(\mathbf{o}, \mathbf{v}) = 2 \frac{(\mathbf{v} \cdot \mathbf{g})(\mathbf{g} \cdot \mathbf{o})}{\langle \mathbf{o}, \mathbf{v} \rangle}. \quad (3.3)$$

Function 3.3 can be derived in different ways. It can be shown trigonometrically, as done by Blinn [1], but also by using the fact that the projected area of the microsurface is the same as that of the macro-surface. This proof was provided by Heitz [7].

The projected area A_{proj}^1 of a flat surface is equal to the surface area A multiplied by the cosine of the angle θ_o between the surface normal \mathbf{g} and the projection direction \mathbf{o} , which is the same as the dot product of \mathbf{g} and \mathbf{o} .

$$A_{proj}^1 = A \cos(\theta_o) \quad (3.4)$$

To simplify the derivation, the dimensions are usually chosen such that $A = 1$.

This A_{proj}^1 is equal to the projected area of the visible microsurface in the same direction, which is given by

$$A_{proj}^2 = \int_{\Omega} D(\mathbf{v}) \langle \mathbf{o}, \mathbf{v} \rangle G_1(\mathbf{o}, \mathbf{v}) d\mathbf{v}. \quad (3.5)$$

We also know that in the 2D case there are only two face orientations, and so $D(\mathbf{v}) = \sum_{i=1}^2 \frac{1}{2} \frac{\delta_i(\mathbf{v})}{\mathbf{i} \cdot \mathbf{g}}$.

Setting Equations 3.4 and 3.5 to be equal to each other allows us to derive the formula for G_1 .

$$\begin{aligned} \cos(\theta_o) &= \int_{\Omega} D(\mathbf{v}) \langle \mathbf{o}, \mathbf{v} \rangle G_1(\mathbf{o}, \mathbf{v}) d\mathbf{v} \\ &= \frac{1}{2} G_1(\mathbf{o}, \mathbf{f}_e) \frac{\langle \mathbf{o}, \mathbf{f}_e \rangle}{\mathbf{f}_e \cdot \mathbf{g}} + \frac{1}{2} G_1(\mathbf{o}, \mathbf{f}_w) \frac{\langle \mathbf{o}, \mathbf{f}_w \rangle}{\mathbf{f}_w \cdot \mathbf{g}} \end{aligned}$$

For an angle θ_o in section A in Figure 3.2, the area for which we want to derive the function, the eastern faces (\mathbf{f}_e) are backfacing and $G_1(\mathbf{o}, \mathbf{f}_e) = 0$. Then

$$\begin{aligned} \cos(\theta_o) &= \frac{1}{2} G_1(\mathbf{o}, \mathbf{f}_w) \frac{\langle \mathbf{o}, \mathbf{f}_w \rangle}{\mathbf{f}_w \cdot \mathbf{g}} \\ G_1(\mathbf{o}, \mathbf{f}_w) &= 2 \frac{\cos(\theta_o)(\mathbf{f}_w \cdot \mathbf{g})}{\langle \mathbf{o}, \mathbf{f}_w \rangle} \\ &= 2 \frac{(\mathbf{o} \cdot \mathbf{g})(\mathbf{f}_w \cdot \mathbf{g})}{\langle \mathbf{o}, \mathbf{f}_w \rangle} \end{aligned}$$

which, when generalised to a random face normal \mathbf{v} , is the function given in Equation 3.3. The full function is plotted in Figure 3.3. The slope seen around -35.3° is the result of the step size in the calculation; in reality, we would see a Heaviside step function here, creating an immediate drop.

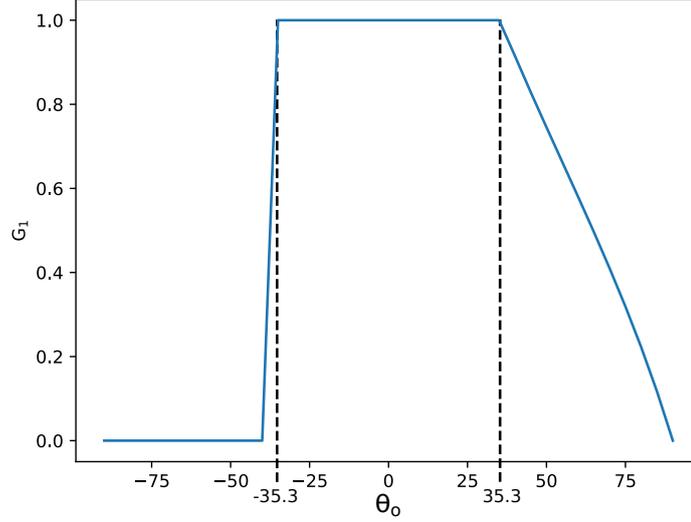


Figure 3.3: Plot of the value of G_1 against θ_o for the 2D case, showing the three sections of G_1 . The dotted lines illustrate the transitions between sections in Figure 3.2, at -35.3° and 35.3° .

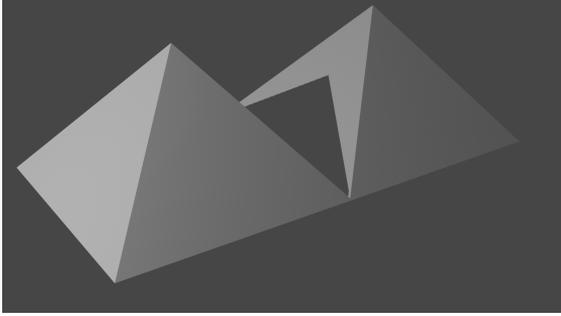


Figure 3.4: Shadowing in 3D from azimuth aligned with pyramid axis, made in Blender. Note the triangular shape of the shadow.

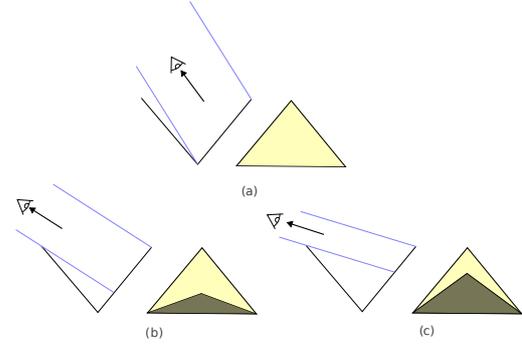


Figure 3.5: Side and front view of the masked area of a face at different view angles, showing how the masked area will change as the view angle is varied.

3.2.2. 3D with fixed azimuth

We now complicate our scenario slightly by considering a 3D case, where there may be any amount of randomly distributed pyramids of the same height, but the light is only coming from a single azimuth aligned with the grid. Each pyramid will throw a shadow in the shape of an isosceles triangle. This is illustrated in Figures 3.4 and 3.5. Referencing back to Figure 3.2, we are once again interested in deriving a formula for $G_1(\theta)$ in the case where \bullet falls in section A. The value of $G_1(\theta)$ in sections B and C is the same as in the 2D scenario. As we know from Heitz [7], the projected area of the macrosurface is equal to that of the microsurface. Setting these to be equal, where the first term is the projected area of the macrosurface and the second is that of the microsurface, we have,

$$\begin{aligned} \langle \mathbf{o}, \mathbf{g} \rangle &= \int_{\Omega} D(\mathbf{v}) \langle \mathbf{o}, \mathbf{v} \rangle G_1(\mathbf{o}, \mathbf{v}) d\mathbf{v} \\ &= \sum_{\mathbf{i} \in \{\mathbf{f}_n, \mathbf{f}_e, \mathbf{f}_s, \mathbf{f}_w\}} D(\mathbf{i}) \langle \mathbf{o}, \mathbf{i} \rangle G_1(\mathbf{o}, \mathbf{i}). \end{aligned}$$

Let's say we are considering the light direction facing the western face (\mathbf{f}_w). Then we see \mathbf{f}_n and \mathbf{f}_s from the side. We can assume the terms for both side faces are equal, as for an azimuth aligned with one face, the sides are exposed symmetrically as you increase the elevation.

$$\begin{aligned} \langle \mathbf{o}, \mathbf{g} \rangle = & D(\mathbf{f}_w) \langle \mathbf{o}, \mathbf{f}_w \rangle G_1(\mathbf{o}, \mathbf{f}_w) \\ & + 2 \cdot D(\mathbf{f}_s) \langle \mathbf{o}, \mathbf{f}_s \rangle G_1(\mathbf{o}, \mathbf{f}_s) \\ & + D(\mathbf{f}_e) \langle \mathbf{o}, \mathbf{f}_e \rangle G_1(\mathbf{o}, \mathbf{f}_e) \end{aligned}$$

In the range for \mathbf{o} under consideration, section A in Figure 3.2, $\langle \mathbf{o}, \mathbf{f}_e \rangle$ is 0, as the eastern face is backfacing from this angle. This means we can remove this term from the equation. For a random grid, $G_1(\mathbf{o}, \mathbf{f}_w) = G_1(\mathbf{o}, \mathbf{f}_s)$, as found in the measurements provided in the next chapter.

$$\begin{aligned} \langle \mathbf{o}, \mathbf{g} \rangle = & D(\mathbf{f}_w) \langle \mathbf{o}, \mathbf{f}_w \rangle G_1(\mathbf{o}, \mathbf{f}_w) \\ & + 2 \cdot D(\mathbf{f}_s) \langle \mathbf{o}, \mathbf{f}_s \rangle G_1(\mathbf{o}, \mathbf{f}_w) \\ = & (D(\mathbf{f}_w) \langle \mathbf{o}, \mathbf{f}_w \rangle + 2 \cdot D(\mathbf{f}_s) \langle \mathbf{o}, \mathbf{f}_s \rangle) G_1(\mathbf{o}, \mathbf{f}_w) \end{aligned}$$

Solving this for G_1 :

$$G_1(\mathbf{o}, \mathbf{f}_w) = \frac{\langle \mathbf{o}, \mathbf{g} \rangle}{D(\mathbf{f}_w) \langle \mathbf{o}, \mathbf{f}_w \rangle + 2 \cdot D(\mathbf{f}_s) \langle \mathbf{o}, \mathbf{f}_s \rangle}$$

To arrive at a valid formula for any elevation, we use the fact that G_1 is capped at 1 for any angle in section B in Figure 3.2 and add a minimum to the function. We set the function to zero when $\mathbf{o} \cdot \mathbf{f}_w \leq 0$ to ensure we never divide by zero.

$$G_1(\mathbf{o}, \mathbf{f}_w) = \begin{cases} 0 & \text{if } \mathbf{o} \cdot \mathbf{f}_w \leq 0 \\ \min\left(1, \frac{\langle \mathbf{o}, \mathbf{g} \rangle}{D(\mathbf{f}_w) \langle \mathbf{o}, \mathbf{f}_w \rangle + 2 \langle \mathbf{o}, \mathbf{f}_s \rangle}\right) & \text{otherwise} \end{cases} \quad (3.6)$$

Figure 3.6 shows plots of G_1 for the 2D and 3D case together. We note that the descent of the plot on the right side is slower for the 3D case, indicating a greater percentage of faces with a given normal is visible at a given elevation.

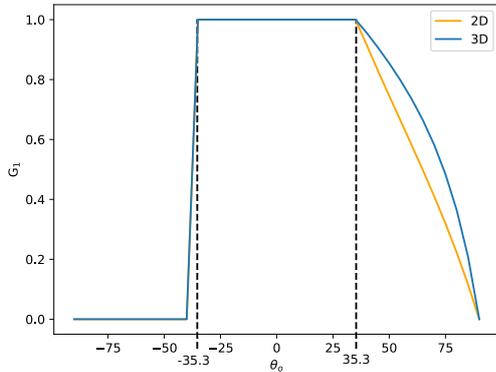


Figure 3.6: Plot of the value of G_1 against θ_0 for the 2D and 3D cases, showing the three sections of G_1 . The dotted lines illustrate the transitions between sections in Figure 3.2, at -35.3° and 35.3° .

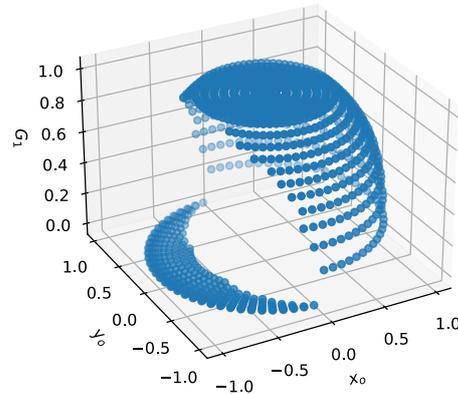


Figure 3.7: 3D plot of the value of G_1 for all \mathbf{o} , given a single \mathbf{v} . x_0 and y_0 are the x- and y-component of \mathbf{o} , respectively.

3.2.3. Varying the azimuth

So far, we have only considered a light direction aligned exactly with the orientation of our pyramids, varying only the elevation and not the azimuth. We want a formulation for G_1 that will work for any \mathbf{o} . The expectation is that G_1 will look the same when varying the azimuth, except for the point at which the face becomes backfacing. Because the pyramids are randomly distributed, the probability of another

pyramid shadowing the face under consideration is always the same. There is always an equal chance of another pyramid being 'in the way'. The simplest way, then, to find the correct value of G_1 is to rotate \mathbf{o} so that its azimuth is 0, then calculate G_1 using this rotated vector. Let \mathbf{r} be the vector with $\theta = \theta_o$ and $\phi = 0$, then the formulation for G_1 is given in equation 3.7, and a 3D plot of the function is shown in Figure 3.7.

$$G_1(\mathbf{o}, \mathbf{v}) = \begin{cases} 0 & \text{if } \mathbf{o} \cdot \mathbf{v} \leq 0 \\ \min\left(1, \frac{\langle \mathbf{r}, \mathbf{g} \rangle}{D(\mathbf{f}_e)(\langle \mathbf{r}, \mathbf{f}_e \rangle + 2\langle \mathbf{r}, \mathbf{f}_s \rangle)}\right) & \text{otherwise} \end{cases} \quad (3.7)$$

3.3. Multiple interactions

Normally, a microfacet model [18] explicitly assumes the light only interacts with the surface once. The surface being analysed here is designed to force rays to interact with it multiple times. Since energy is absorbed by the surface on subsequent bounces after the first one, both the direction and energy of the ray when it finally exits the surface will be different. This means there is a need to account for multiple bounces separately and add this to the model. Our first goal will be to derive the possible exit directions given an incoming light direction and the number of bounces we want to consider. We can then make a choice to use this first model directly, or calculate the energy throughput and chance of escaping. This second method incorporates the shadowing-masking function into the multiple interaction model.

3.3.1. Exit direction

To calculate the reflection direction after a single bounce, we use the following formula,

$$\mathbf{o} = \mathbf{i} - 2(\mathbf{i} \cdot \mathbf{f})\mathbf{f} \quad (3.8)$$

where \mathbf{i} is the incoming ray, \mathbf{o} is the reflection, and \mathbf{f} is the surface normal. For any bounce j after the first, \mathbf{i} simply becomes the result \mathbf{o} from the previous bounce, meaning we have a recursive formula for the reflection direction after j bounces, with $\mathbf{o}_1 = \mathbf{i}$, and \mathbf{f}_j the normal of the face hit in bounce j :

$$\mathbf{o}_j = \mathbf{o}_{j-1} - 2(\mathbf{o}_{j-1} \cdot \mathbf{f}_j)\mathbf{f}_j \quad (3.9)$$

Using this formula, we can find the resulting reflection direction for any combination of faces we could bounce off of. The number of possible results increases exponentially with the number of bounces taken into account; there are 4^b possibilities for b bounces. However, not all combinations can actually occur. For example, it is impossible to hit two faces with the same orientation in a row, which we can prove by contradiction.

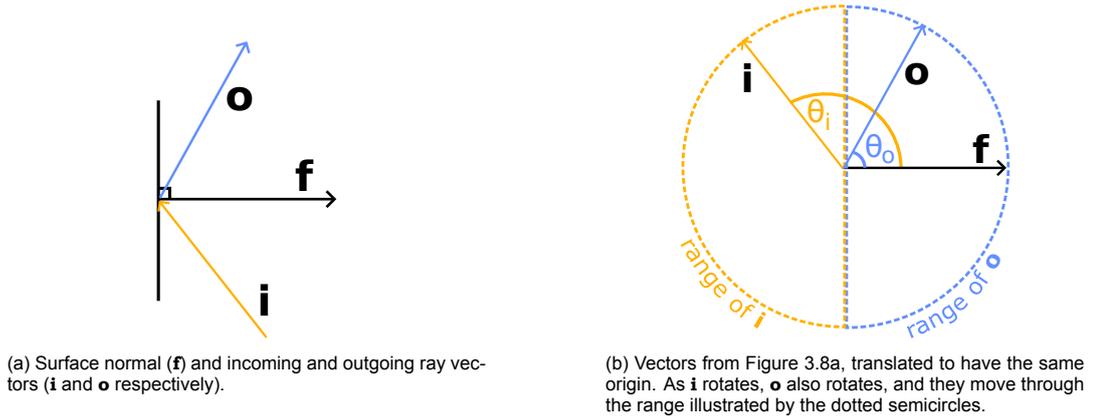


Figure 3.8: The range of movement for \mathbf{i} and \mathbf{o} limit the angles between the vectors and the surface normal to $\theta_i > 90^\circ$ and $\theta_o < 90^\circ$ for any valid reflection.

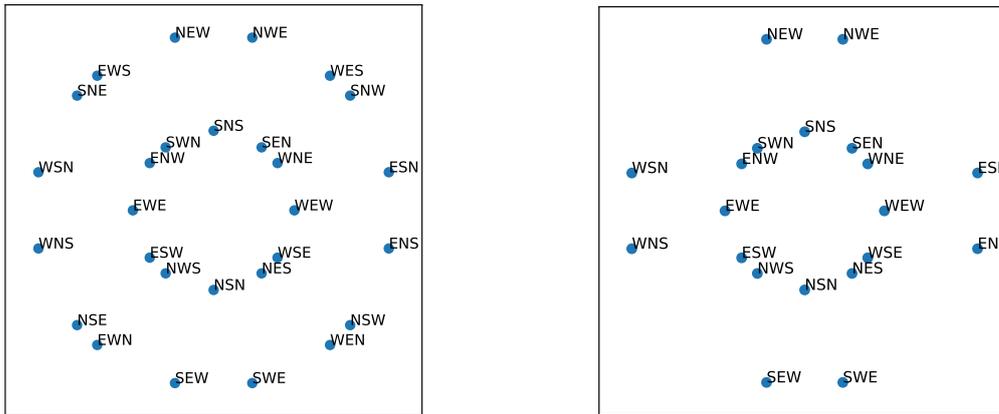
As illustrated in Figure 3.8, for a valid reflection, the incoming vector will be at an angle θ_i of over 90° with the surface normal, and the outgoing vector will be at an angle θ_o of under 90° . The dot product of

two vectors is positive at angles below 90° and negative at angles over 90° . This means a ray pointing into a surface can bounce on that surface when the dot product between the ray vector and surface normal is negative, and the ray will hit the backface if the dot product is larger than 0. We therefore know that for a valid bounce on a face with normal \mathbf{f} , $\mathbf{i} \cdot \mathbf{f} < 0$ and want to show that if a first bounce on a face has occurred, the resulting outgoing ray could not reflect off the same face normal, so $\mathbf{o} \cdot \mathbf{f} \geq 0$. We do so by assuming $\mathbf{o} \cdot \mathbf{f} < 0$ and showing this leads to a contradiction.

$\mathbf{o} \cdot \mathbf{f} < 0$	
$(\mathbf{i} - 2(\mathbf{i} \cdot \mathbf{f})\mathbf{f}) \cdot \mathbf{f} < 0$	substituting equation 3.8 for \mathbf{o}
$\mathbf{i} \cdot \mathbf{f} - 2(\mathbf{i} \cdot \mathbf{f})\mathbf{f} \cdot \mathbf{f} < 0$	by distributive property of the dot product
$\mathbf{i} \cdot \mathbf{f} - 2(\mathbf{i} \cdot \mathbf{f})(\mathbf{f} \cdot \mathbf{f}) < 0$	by scalar multiplication property of the dot product
$\mathbf{i} \cdot \mathbf{f} - 2(\mathbf{i} \cdot \mathbf{f}) < 0$	by $\mathbf{f} \cdot \mathbf{f} = \ \mathbf{f}\ ^2 = 1$
$\mathbf{i} \cdot \mathbf{f} > 0.$	

This cannot be true, as we previously noted that to not hit the backface, $\mathbf{i} \cdot \mathbf{f} < 0$ must hold for the incoming ray. It must therefore be true that $\mathbf{o} \cdot \mathbf{f} \geq 0$ and therefore that a ray can never bounce off two faces with the same orientation in a row.

Eliminating bounce sequences hitting the same face twice in a row reduces the possibilities to $4 \cdot 3^{b-1}$ for b bounces. We can reduce this number further by realizing that more of these combinations involve hitting a backface. We can check for this by only keeping reflections where the dot product between the face normal and the light direction is non-positive. Figure 3.9 shows the results of this check for 3 bounces. In this case, 20 options remain out of the original 64.



(a) Before removing sequences containing a backface reflection (b) After removing sequences containing a backface reflection

Figure 3.9: Possible reflection directions of an incoming ray from straight above after 3 bounces on a pyramidal surface, projected using the Lambert equal-area mapping. Only vectors in the upper hemisphere are plotted here. Labels denote the sequence of bounces: e.g. ESW means east-south-west bounces in that order.

3.3.2. Probability of a sequence

We are now left only with sequences of bounces that should, in theory, be possible. However, even if a sequence is not impossible, the chance of it occurring may be very low. This depends on the angle at which the ray is approaching the face.

It is important to take this into consideration in the model to increase its accuracy. We can determine the probability of each face being hit by projecting the area of each face onto the incoming ray direction with the cosine of the angle θ_i between the face normal and the incoming ray. This projection represents how much of the surface is 'visible' from the incoming direction. If we restrict ourselves to using unit

vectors, this cosine is the same as the dot product used to check for backfaces in the previous section. If the cosine is below 0, it is impossible to hit the face from this direction. If the ray is coming at a face with a grazing angle, the cosine will be small. We multiply the cosine at each bounce with those of the previous bounces in the sequence, and colour the points in Figure 3.9b to match the product of the cosines in their sequence of interactions. The result is shown in Figure 3.10.

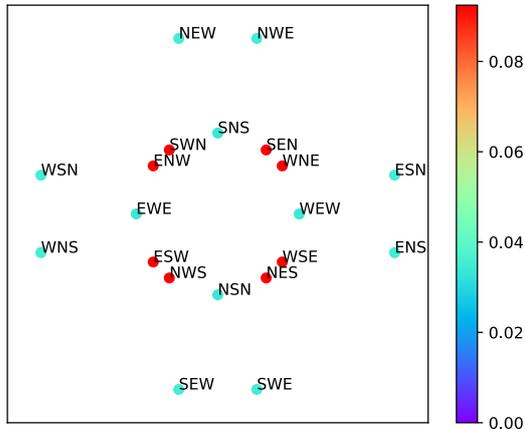


Figure 3.10: Cosine product for any possible sequence of 3 bounces. Here, the coordinates denote the projection of the exit direction onto the x-y plane, while the colour denotes the cosine product.

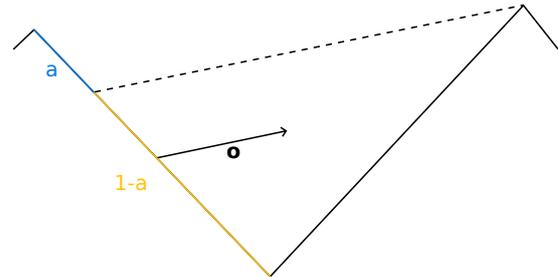


Figure 3.11: Illustration of the probability a of a ray in direction \bullet escaping the surface, and the probability $1 - a$ of it reflecting again.

Instead of just calculating the probability of each face being hit, we can move toward computing the energy transfer for each sequence of bounces by more accurately modelling the probability of actually escaping the surface after an interaction. To do so, we need to not only take into account the projection of a face onto the incoming direction, but also the chance of the ray escaping after this bounce. As a ray reflects from the last face in the sequence, it might escape the surface entirely with a probability a , but also has a chance of $1 - a$ of another pyramid being in the way, as illustrated in Figure 3.11.

This probability a is one that this paper has discussed before: it is equal to the value of G_1 . We expect points that lie towards the edge of the figure, so rays that bounce at a shallow angle, to have a bigger chance of encountering another pyramid on their way. Since G_1 also becomes small at shallow angles, this should be exactly what we see in the results. We can incorporate G_1 into the model by multiplying the energy transferred with the value of G_1 for the final bounce, which yields the result shown in figure 3.12. Comparing this to figure 3.10, we see that points further toward the edge of the image seem to have a smaller chance of escaping the surface, as was the expectation.

The multiple interactions model has been laid out step by step in this section. When implementing the microfacet model, it will depend on the implementation strategy whether only the exit directions or probability of hitting a face should be included, or whether the value of energy transferred is a better fit.

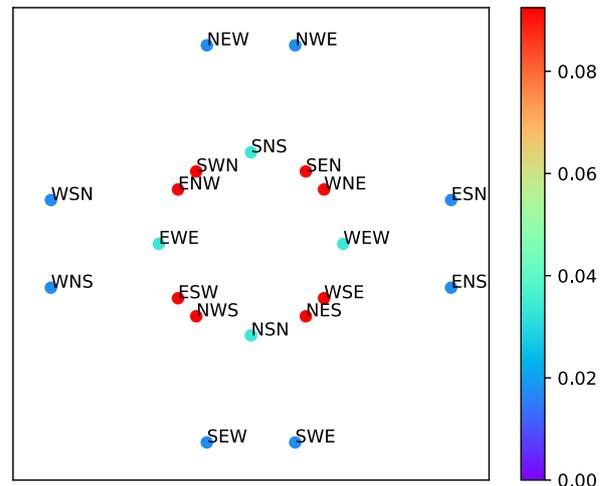


Figure 3.12: Probability of occurrence and escape for any possible sequence of 3 bounces. Here, the coordinates denote the projection of the exit direction onto the x-y plane, while the colour denotes the energy transferred.

4

Modelling And Measured Results

In this chapter, pyramidal microsurfaces will be simulated and their properties will be measured using a digital gonioreflectometer. The resulting measurements will then be compared to the expected results derived in Chapter 3.

4.1. Recreating the surface

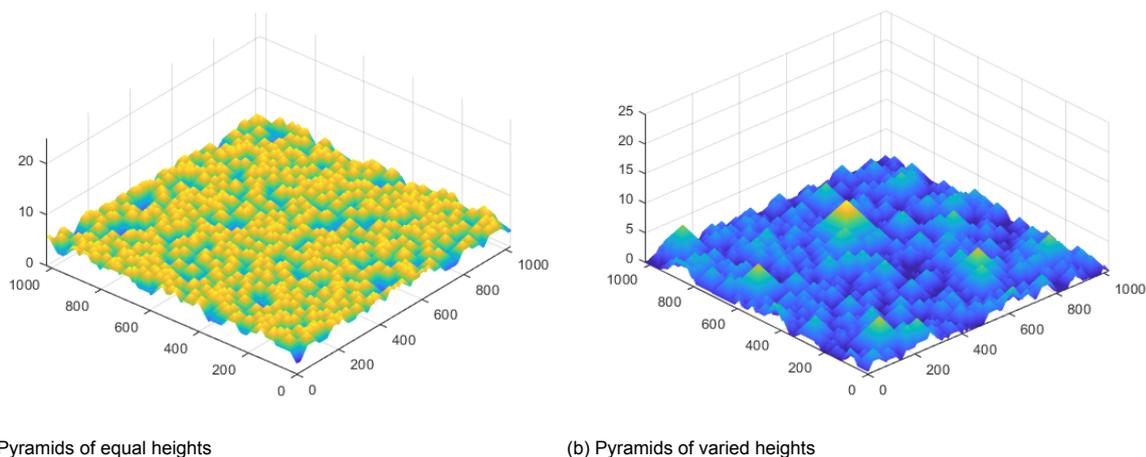


Figure 4.1: Examples of randomly distributed 50x50 micron surfaces

A height map of the surfaces to be tested was needed as input for the gonioreflectometer. These surfaces were created in Python using PyOpenGL. Because all pyramids have the same side angle, a single pyramid object could be created and copied across the grid, rescaling it if necessary. This allowed for easy manipulation of the positions of the pyramids, making it possible to test many different (random) configurations. Meshes were created by generating a list of seed points, and spawning a pyramid at each of these seeds on top of a square mesh serving as the floor. Pyramid density could be controlled by varying the total amount of pyramids added.

To generate a height map, an orthogonal projection from straight above the mesh was used. The depth buffer was then extracted and saved in a plain text format, which could be used as input for the gonioreflectometer.

There are two main types of grid that were used for testing. The first is a grid with randomly distributed pyramids that are all the same height. This is also the grid for which Section 3 derives the expected functions. The second grid is more realistic to the surface of a real solar cell. Here, we not

only generate a list of seeds, but also a list of heights for each of the pyramids. These heights are drawn from the height distribution given by Procel et al. [14], exact data for which was obtained from the author. Both types of surfaces will be measured and compared in this section.

4.1.1. Assumptions and limitations

Height field resolution

The digital gonireflectometer is resource-intensive, and there are limits to the resolution of the height field it can handle. As Figure 4.2 demonstrates, the height field measures the height of the surface at equally spaced points, and linearly interpolates between them. In the cases where an edge does not exactly align with a grid point on the height field, that edge will be flattened slightly. This means in those areas, we will find a variety of different normals depending on how far the grid points lie from the edge. This effect becomes smaller as the mesh resolution is increased, but there is a limit to how detailed this resolution can be. This needs to be balanced with the size of the mesh, as it also needs to contain enough pyramids to average out any bias in their random placement.

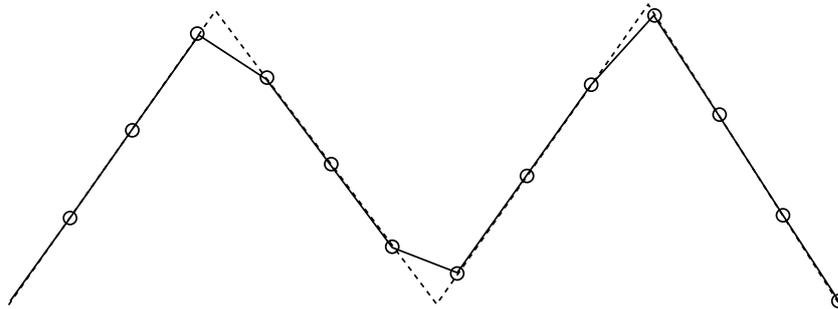


Figure 4.2: Rendered microsurface (dotted line) versus height field (solid line), linearly interpolated between the grid points (circles) extracted from the depth buffer.

Bevelled edges

In these grids, we assume the pyramids' shape consists of four perfect triangles. When looking at the microscope image in Figure 2.1, it seems this may not actually be the case. It appears the edges of the pyramids may be slightly bevelled. However, the impact of these bevelled edges was judged to presumably be negligible. It also seems plausible that any changes caused by including bevelled edges will be drowned out by the grid resolution issues discussed above.

4.2. Normal distribution function

We use the digital gonireflectometer to measure the NDF, for both equal and random height pyramids. For a small sample surface, the heights of the peaks in the NDF graph will not be equal, as pyramids may overlap each other more on one side than the other. We need to use a sample surface large enough to allow this random overlap to even out on all sides. We also need to use a surface with a high enough density. If the density is too low, not all floor space will be covered by pyramids and some of the surface will have a normal going straight up. This will result in a fifth peak in the middle of the graph. This configuration is not true to the surfaces this model is based on. As the goal of the pyramids is to trap more light in the surface, allowing some of the area to be bare would allow light to escape unnecessarily. This configuration will thus not be investigated further. As is shown in Figure 4.3a, the NDF consists of four peaks, as derived in Section 3.1.

Though the measured results mostly match those that were expected, one key difference is that the peaks are not quite narrow enough to be Dirac deltas. Some faces have a normal that is close to, but not quite equal to the expected four normals. This is caused by two factors, the first of which, noticeable in Figure 4.3a, is simply the step size of the measurements. The plots linearly interpolate between data points, resulting in the peaks not being perfect Dirac deltas. The second issue is the limited resolution of the height field, as discussed in Section 4.1.1. Figure 4.3b shows what happens when the mesh resolution is too small for the mesh size, so when each pyramid spans too few grid

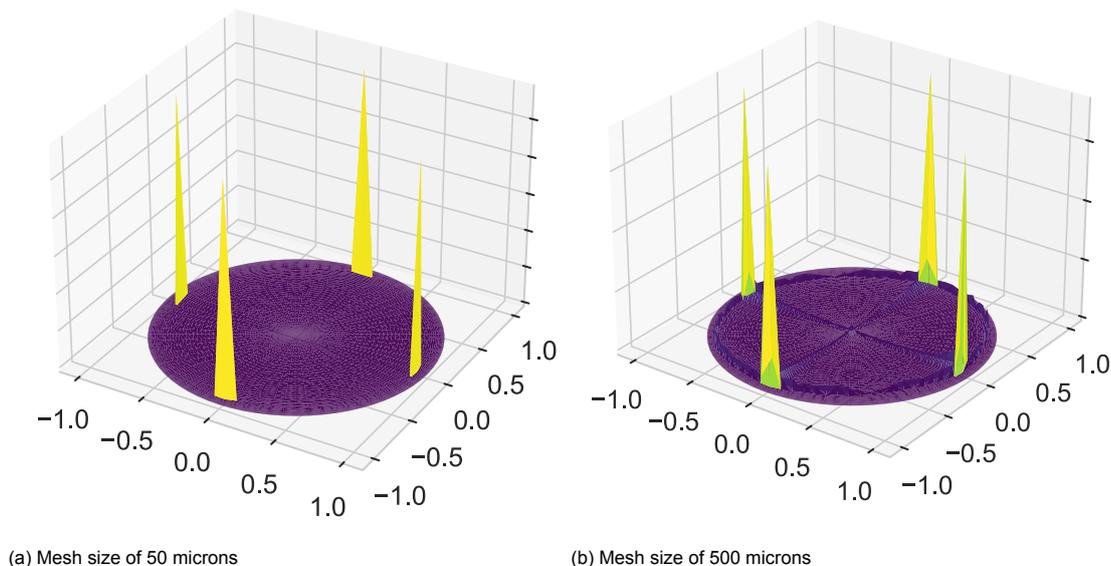


Figure 4.3: Measured normal distribution function with a mesh resolution of 1024x1024

points. The NDF now contains small amounts of many more face normals than the four we would expect to see, though the four peaks are still clearly defined. The deviation explained in Section 4.1.1 will only occur on pyramid edges, so the errant normals make up only a relatively small portion of the total NDF.

4.3. Shadowing-masking function

To measure the shadowing-masking function, we once again input meshes into the digital gonireflectometer. The results are given per section of microsurface normals, for all incident angles. This gives us a 3D graph depicting G_1 for a single facet orientation, where rotation around the origin denotes an increasing azimuth and distance from the origin denotes an increasing elevation. This graph is shown in Figure 4.4. We can also consider a 2D version of this graph, made by taking a profile of the 3D function along any desired azimuth. In the case of Figure 4.5, the profile was taken along the x-axis of Figure 4.4, which is the same azimuth as the light direction. The results are as expected. In Figure 4.4 we see that the function is 0 whenever the face is backfacing the view direction being evaluated, and 1 whenever the view direction is high enough over top of the surface that no shadowing occurs at all. When neither of those conditions hold, the function drops from one to zero at the same speed on all sides, creating a bell shape.

We will compare Figure 4.5 to the theoretical results by plotting them together. Figure 4.6 shows the theoretical and measured functions for G_1 plotted on top of each other, confirming the results are very similar. However, though the functions lie close to one another, they do not match completely. This will always be the case: the measured function will always decline slightly slower than the theoretical function. This difference is caused by the use of a finite grid for measurements. Pyramids on the edge of the grid closest to the light direction will never be shadowed by another pyramid, but always remain fully visible. As the grid is made larger, the pyramids on the edge of the grid make up a smaller percentage of the total and the effect on G_1 becomes smaller, but the graphs will never fully match for a finite grid.

We see the shape of the function conforms to the expectations formed in Section 3.2. At an elevation smaller than 35.3° , the full surface is visible and G_1 is one. There is a fast decline to 0 as the angle of incidence moves to be behind the face. In the rest of the figure, as the face is shadowed by pyramids next to it, G_1 curves back down towards 0.

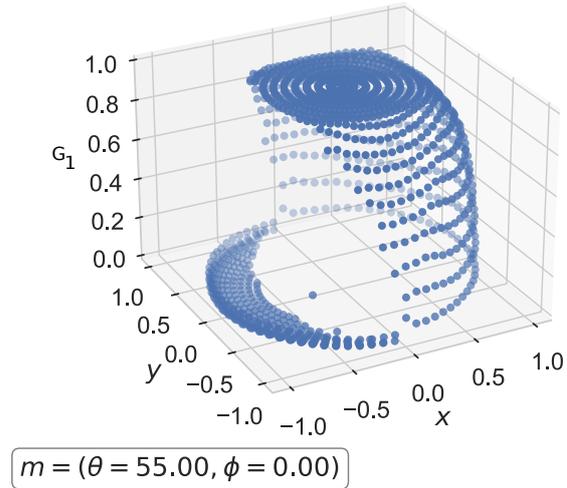


Figure 4.4: 3D plot of measured G_1 , for all possible ray directions \mathbf{o} , given a microsurface normal \mathbf{m} with an elevation of 52.5° to 57.5° and azimuth of -2.5° to 2.5° . x and y are the x - and y -component of \mathbf{o} , respectively.

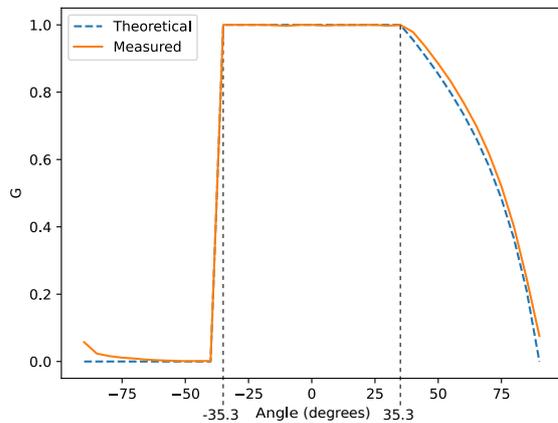


Figure 4.6: Measured (solid orange, Figure 4.4) versus theoretical (dotted blue) function for G_1 , with the black dotted lines marking the rotation over the apex of a pyramid as illustrated in Figure 4.7

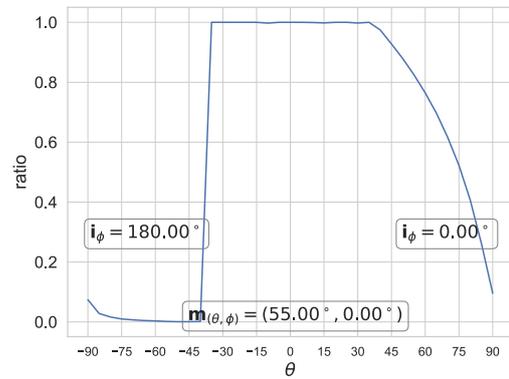


Figure 4.5: Profile of Figure 4.4 along x -axis. The left half of the plot has azimuth 180° , the right half has azimuth 0° , matching the azimuth of the microfacet normal under consideration.

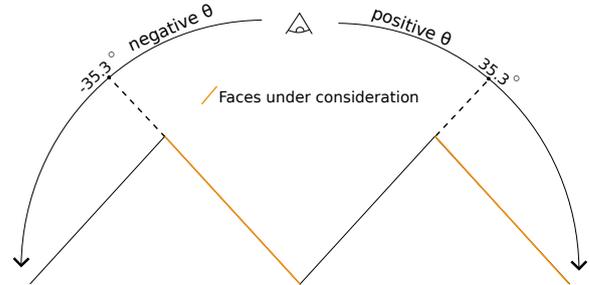


Figure 4.7: Illustration of how Figures 4.5 and 4.6 should be interpreted, showing the faces under consideration in orange. The value on the x -axis of the plots changes with the rotation of \mathbf{o} across the surface.

4.3.1. Validity for all face normals

Figure 4.6 shows the measurement matches the derived function in Equation 3.6 for the microsurface normal \mathbf{m} with azimuth $\phi_m = 0^\circ$. We need to verify this function holds for the other three face normals as well. Figure 4.8 demonstrates that all four functions are equal to each other. Each of the four subplots shows G_1 for one face normal minus G_1 for the normal with a 90° greater azimuth. We see this evaluates to zero almost everywhere. In each plot, there are a few points which are nonzero. These points lie in the section where the function jumps from zero to one as the ray direction \mathbf{o} transitions from being behind the face to in front of it (sections B and C in Figure 3.2). Which exact points fall in the transition varies slightly per measurement, resulting in the non-zero values in Figure 4.8. Since the measurements are equal for all four normals, the derived G_1 will hold for all of them.

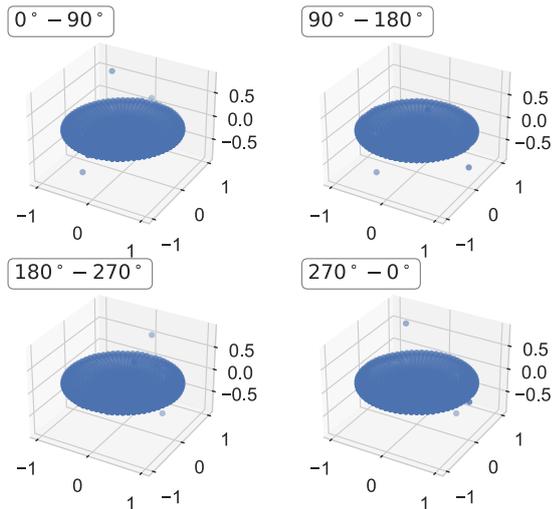


Figure 4.8: 3D plots of the differences between the measured G_1 for each face normal and the G_1 for the counterclockwise adjacent face normal. The labels denote the azimuths of the normals that have been subtracted from each other in each plot.

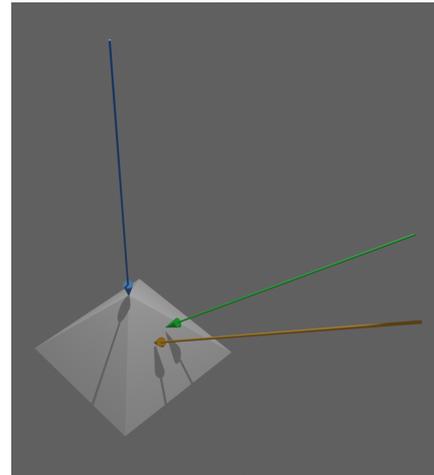


Figure 4.9: 3D illustration of the incoming ray directions as discussed in Sections 4.4.1 (blue), 4.4.2 (orange), and 4.4.3 (green). Zero azimuth is aligned with the orange arrow, pointing to the bottom right of the image.

4.4. Multiple interactions

Finally, the multiple interaction model is verified. Once again, this is done using the digital gonioreflectometer, by measuring the energy throughput for all outgoing directions after a given incoming direction and number of bounces. Only rays escaping the surface after the given number of bounces are considered. The measurements were taken for 20 random mesh configurations, and then averaged to eliminate any noise due to the surface configuration. The results are given as values at the outgoing light direction on a hemisphere, mapped onto the x-y plane using the Lambert azimuthal equal-area projection [8]. This is the same plotting method that was used in Section 3.3, which makes the resulting images easy to compare to the theoretical results. We compare the first three bounces, as for any number above that, the energy transfer was found to be minimal, meaning only a very small percentage of light interacts with the surface more than three times.

Figure 4.9 illustrates the incoming ray directions to be evaluated in the following subsections. We will first compare results for an incoming ray from straight above in Section 4.4.1. We then vary the elevation of the ray in Section 4.4.2, and finally, in Section 4.4.3, set both the elevation and the azimuth of the ray to be nonzero.

4.4.1. Incoming ray at normal incidence

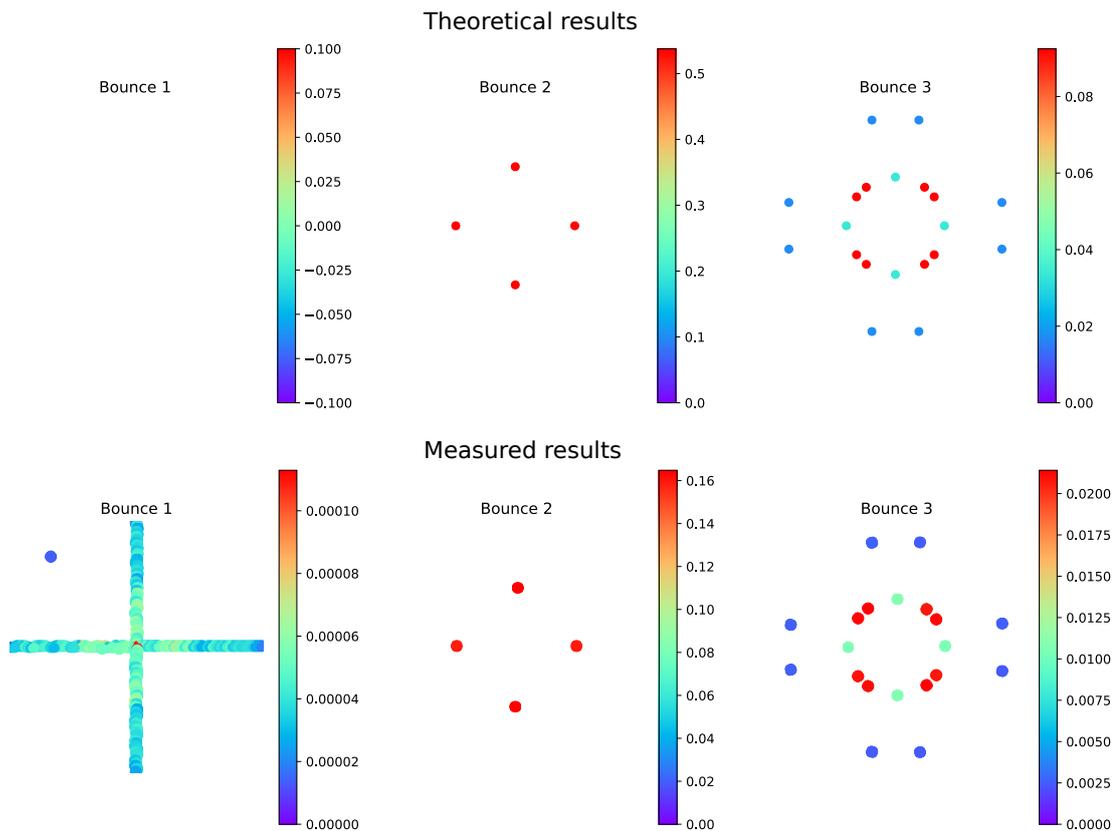


Figure 4.10: Theoretical versus measured results for the first three bounces for an incoming ray at normal incidence. The plots show the energy transmission for rays escaping the surface in each direction after one, two and three bounces, respectively.

We will first compare the results for a ray originating straight above the surface in Figure 4.10. The theoretical results show no non-zero values for the first bounce, indicating rays never escape the surface after the first bounce. The measured results appear to show very different results for the first bounce. However, upon closer consideration, the values of the points are very close to 0. The range of non-zero values in the measured results can be explained by the imperfections in the height field that were previously discussed in Section 4.2 and illustrated in Figure 4.2, and have a small enough throughput that they can be safely ignored.

In the cases of both the two bounce and three bounce plots, the theoretical results match the measured results exactly, both in location and colour.

When looking at the legends of both sets of plots, we see that these do not match in general, though the proportional energy throughput between points in one plot does match. This happens because the value in the measured results denotes the percentage of total energy being diverted in any one direction after a given number of bounces. The theoretical model does not incorporate this full measurement of energy throughput, but instead only takes into account the probability of a sequence of interactions occurring, and the chance of a ray escaping after its last interaction.

4.4.2. Incoming ray with 45 degree elevation

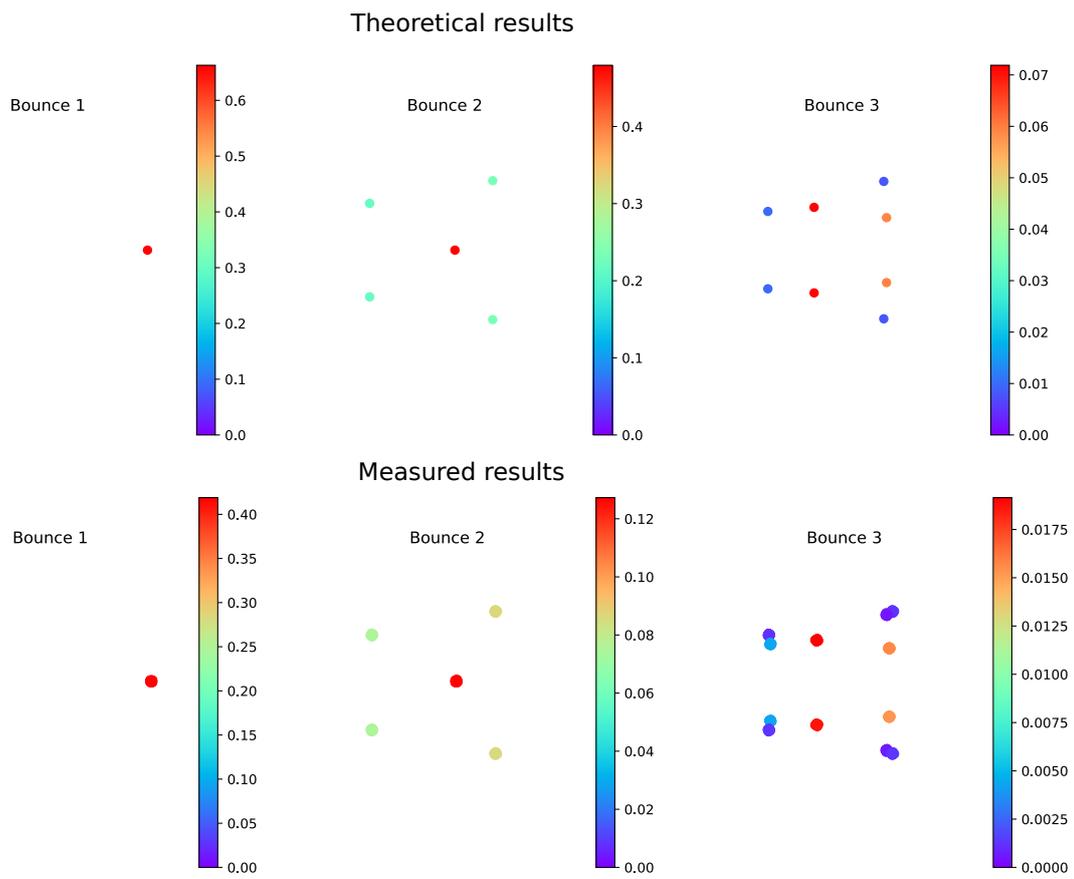


Figure 4.11: Theoretical versus measured results for the first three bounces for an incoming ray at a 45° elevation. The plots show the energy transmission for rays escaping the surface in each direction after one, two and three bounces, respectively.

Figure 4.11 shows the theoretical and measured results when considering an incoming ray at a 45° angle. For all measurements both the configuration and values of the points are nearly identical to the expected results. In the measured three bounce results we see points that do not appear in the theoretical results. Like in the case of one bounce, these points have a very low energy transfer value. As they also appear close to other valid points, they may be the result of the inaccuracies in the height field described in Section 4.1.1.

4.4.3. Incoming ray with 45 degree elevation, 25 degree azimuth

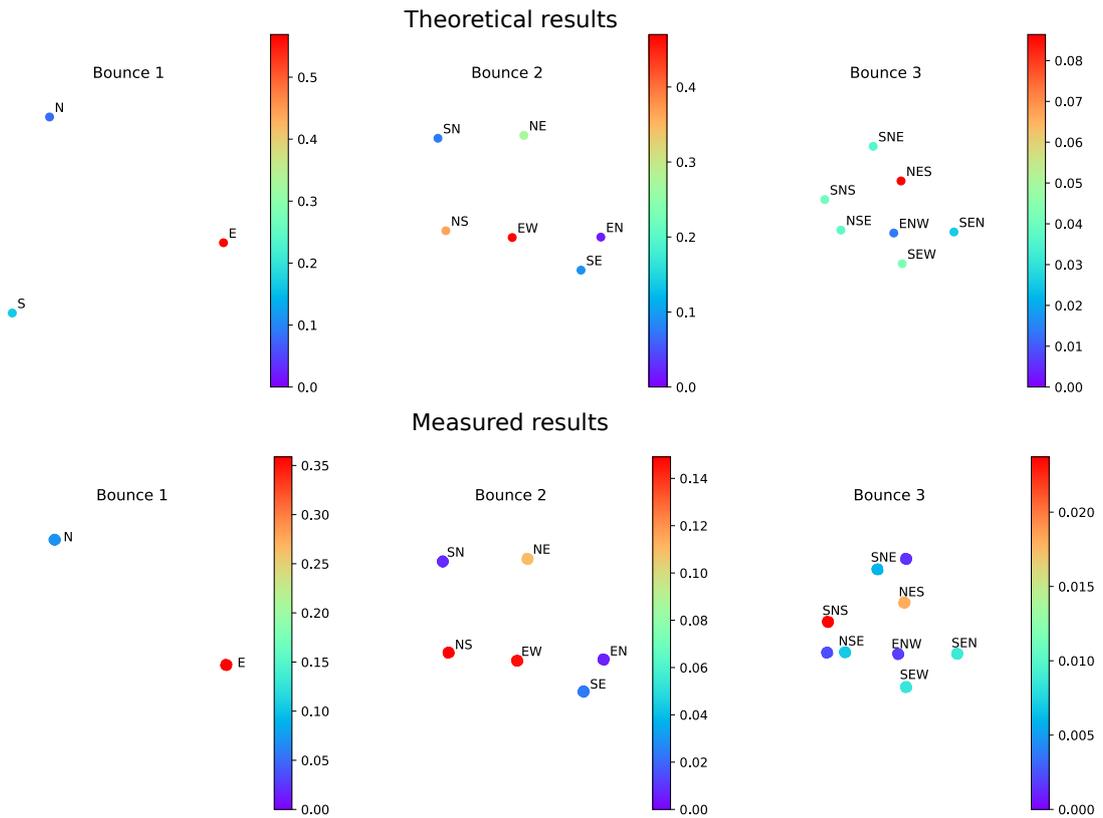


Figure 4.12: Theoretical versus measured results for the first three bounces for an incoming ray at a 45° elevation and 25° azimuth. The plots show the energy transmission for rays escaping the surface in each direction after one, two and three bounces, respectively.

The previous examples were both based on an incoming ray with zero azimuth, meaning the incoming ray is aligned with the pyramid orientation. We also need to check whether the model holds for nonzero azimuth. The last example we compare is one where the azimuth of the incoming ray is set to 25°. This angle was chosen for its lack of alignment to either the faces or their edges. The result, as is shown in Figure 4.12, lacks the symmetry we saw in previous examples. To make comparison easier, the points in the theoretical results have been labelled with the sequence of faces interactions have taken place with, and the same labels have been placed on the measured results. Though the locations of corresponding points match exactly, we do see that a point in the theoretical results is missing in the measured ones and vice versa. We also see that the energy transfer values, displayed as the colour of the points, are not equal.

For the first bounce, the ratio of energy transfer between the north and east faces matches in both results. However, we see that though the theoretical model indicates there should be some interaction with the south face, this is not found to be the case in the measurements. A possible explanation for this fact is that, though it is theoretically possible to hit the south face and immediately escape, the probability is small enough that this did not occur in the simulations.

For the third bounce we see the opposite phenomenon: here, the measured results show two points that are not visible at all in the theoretical model. All theoretical exit directions were checked, including those resulting in zero energy transfer, and none were found to match these two points. They are therefore attributed to noise in the meshes, probably caused by the edge of the pyramids closest to the incoming ray being bevelled in the height field.

5

Conclusion And Future Work

In this final chapter, we will recap the work performed in this thesis, and suggest avenues for future research building upon the steps made here.

5.1. Conclusion

The objective of this thesis was to analyse the spatial properties of the pyramidal surface of a c-Si solar cell and formulate statistical expressions of its properties for use in a microfacet model.

To achieve this goal, a theoretical analysis of the surface was performed. A derivation for V-grooves was adapted to find expressions for the normal distribution function and shadowing-masking function. Additionally, a method to determine possible exit directions and the probability of escape in each exit direction was developed.

To verify the accuracy of the models, the pyramidal surface was recreated digitally and the normal distribution function, shadowing-masking function and energy transfer after n bounces were measured using a digital gonireflectometer. These measurements were then compared to the theoretically derived functions and confirmed them.

The functions derived in this thesis fully statistically describe the pyramidal microsurface under analysis, allowing it to be simulated in rendering using a microfacet model.

5.2. Future work

5.2.1. Implementing the microfacet model

Though this work supplies all components needed for a microfacet model, it does not yet implement them. The next step in this process should be to implement the model in a suitable renderer and evaluate the appearance of the resulting images against samples of textured silicon.

5.2.2. Adding layers

The original request from the PVMD group was for images of solar cells with the optical filters they developed. This thesis only constitutes the first step in working towards rendering these solar cells fully, as they consist of many layers. To realistically render the appearance of Lizcano et al.'s OF solar cells, it will be necessary to first implement the layers in a regular c-Si solar cell, and then expand this to include the multi-layered optical filter. As this filter gives the cell an iridescent effect, a spectral renderer will be needed to accurately capture this phenomenon.

5.2.3. Exploring other textures

Though the pyramidal microsurface explored in this work is commonly used in c-Si cells, it is not the only option available. Owing to the fact that the texture of these cells is dependent on the crystal structure of the silicon, other textures are often semi-structured in a way comparable to the pyramidal surface modelled in this work. This thesis could constitute a jumping-off point for the exploration of microfacet models representing other existing surface textures, to eventually allow a visual comparison between the use of different solar cells in building-integrated applications.

Bibliography

- [1] James F. Blinn. "Models of Light Reflection for Computer Synthesized Pictures". In: *Proceedings of the 4th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '77. San Jose, California: Association for Computing Machinery, 1977, pp. 192–198. ISBN: 9781450373555. DOI: 10.1145/563858.563893. URL: <https://doi.org/10.1145/563858.563893>.
- [2] Olaf Clausen et al. "Investigation and simulation of diffraction on rough surfaces". In: *Computer Graphics Forum* 42.1 (Oct. 2022), pp. 245–260. DOI: 10.1111/cgf.14717. URL: <https://doi.org/10.1111/cgf.14717>.
- [3] European Commission. *2030 Climate Target Plan*. 2023.
- [4] Robert Cook and Kenneth Torrance. "A Reflectance Model for Computer Graphics". In: *ACM Trans. Graph.* 1 (Jan. 1982), pp. 7–24. DOI: 10.1145/965161.806819.
- [5] Augustin Jean Fresnel. *Mémoire sur la loi des modifications que la réflexion imprime à la lumière polarisée*. Jan. 1823.
- [6] Guardian Glass. *Building-integrated photovoltaics (BIPV)*. 2024. URL: <https://www.guardianglass.com/la/en/why-glass/build-with-glass/applications-of-glass/glass-for-facades/bipv>.
- [7] Eric Heitz. "Understanding the Masking-Shadowing Function in Microfacet-Based BRDFs". In: *Journal of Computer Graphics Techniques (JCGT)* 3 (June 2014).
- [8] Martin Lambers. "Mappings between Sphere, Disc, and Square". In: *Journal of Computer Graphics Techniques (JCGT)* 5.2 (Apr. 2016), pp. 1–21. ISSN: 2331-7418. URL: <http://jcgt.org/published/0005/02/01/>.
- [9] Hongsong Li et al. "Automated three-axis gonioreflectometer for computer graphics applications". In: *Optical Engineering* 45 (Apr. 2006). DOI: 10.1117/12.617589.
- [10] Juan Camilo Ortiz Lizcano et al. "Colored optic filters on c-Si IBC solar cells for building integrated photovoltaic applications". In: *Progress in Photovoltaics: Research and Applications* 30 (4 Apr. 2022). Introduction of solar cell layer data, pp. 401–435. ISSN: 1062-7995. DOI: 10.1002/pip.3504. URL: <https://onlinelibrary.wiley.com/doi/10.1002/pip.3504>.
- [11] Robert Macku, Pavel Koktavý, and Jiří Šicner. "Comprehensive Study of Solar Cell Structure Defects by Means of Noise and Light Emission Analysis". In: *Advances in Electrical and Electronic Engineering* 10 (June 2012). DOI: 10.15598/aeer.v10i2.620.
- [12] Arend Jasper Nijdam. "Anisotropic wet-chemical etching of silicon pits, peaks, principles, pyramids and particles". PhD thesis. University of Twente, 2001.
- [13] Rüdiger Paschotta. *Fresnel equations*. Apr. 2024. URL: <https://doi.org/10.61835/tq1>.
- [14] Paul Procel et al. "Opto-electrical modelling and optimization study of a novel IBC c-Si solar cell". In: *Progress in Photovoltaics: Research and Applications* 25.6 (2017), pp. 452–469. DOI: <https://doi.org/10.1002/pip.2874>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/pip.2874>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/pip.2874>.
- [15] Silson. *Silicon Etching*. <https://www.silson.com/product/silicon-etching/>. Accessed: 2024-03-11. 2022.
- [16] B. Smith. "Geometrical shadowing of a random rough surface". In: *IEEE Transactions on Antennas and Propagation* 15.5 (1967), pp. 668–671. DOI: 10.1109/TAP.1967.1138991.
- [17] K. E. Torrance and E. M. Sparrow. "Theory for Off-Specular Reflection From Roughened Surfaces*". In: *J. Opt. Soc. Am.* 57.9 (Sept. 1967), pp. 1105–1114. DOI: 10.1364/JOSA.57.001105. URL: <https://opg.optica.org/abstract.cfm?URI=josa-57-9-1105>.

-
- [18] Bruce Walter et al. *Microfacet Models for Refraction through Rough Surfaces*. 2007.
- [19] Ruston Youngblood. *Photo by Ruston Youngblood on unsplash*. July 2015. URL: <https://unsplash.com/photos/reflection-of-sky-on-water-U1POHKqXYKs>.