

Fabricating BRDFs at High Spatial Resolution Using Wave Optics

Anat Levin¹ Daniel Glasner¹ Ying Xiong² Frédo Durand³ William Freeman³ Wojciech Matusik³ Todd Zickler²
¹Weizmann Institute ²Harvard ³MIT CSAIL




Figure 1: First column: a wafer fabricated using photolithography displaying spatially varying BRDFs at 220dpi. Second column: designed pattern, color coded according to the 5 reflectance functions used. Dithering is exaggerated for better visualization. Rightmost columns: fabricated pattern as imaged under two different illumination directions. The designed pattern includes anisotropic reflectance functions at two opposite orientations. Hence, the image is inverted when light moves from the horizontal to the vertical directions.

Abstract

Recent attempts to fabricate surfaces with custom reflectance functions boast impressive angular resolution, yet their spatial resolution is limited. In this paper we present a method to construct spatially varying reflectance at a high resolution of up to 220dpi, orders of magnitude greater than previous attempts, albeit with a lower angular resolution. The resolution of previous approaches is limited by the machining, but more fundamentally, by the geometric optics model on which they are built. Beyond a certain scale geometric optics models break down and wave effects must be taken into account. We present an analysis of incoherent reflectance based on wave optics and gain important insights into reflectance design. We further suggest and demonstrate a practical method, which takes into account the limitations of existing micro-fabrication techniques such as photolithography to design and fabricate a range of reflection effects, based on wave interference.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Viewing and display algorithms;

Keywords: Fabrication, BRDF design, wave optics.

Links: DL PDF WEB

1 Introduction

The physical construction of surfaces with controlled appearance and reflectance properties is important for many industrial appli-

cations, including printing, product design, luminaire design, security markers visible under certain illumination conditions, and many others. The topic has been gaining much research interest in computer graphics [Weyrich et al. 2009; Finckh et al. 2010; Papas et al. 2011; Kiser et al. 2012; Dong et al. 2010; Hašan et al. 2010; Matusik et al. 2009; Patow and Pueyo 2005; Patow et al. 2007; Weyrich et al. 2007; Malzbender et al. 2012]. In computer vision, photometric stereo algorithms can be improved if the surface reflectance properties can be precisely controlled [Johnson et al. 2011]. Custom designed BRDFs can also help in appearance acquisition tasks such as a BRDF chart [Ren et al. 2011] and a planar light probe [Alldrin and Kriegman. 2006].

Recent attempts to fabricate surfaces with custom reflectance functions boast impressive angular resolution, yet their spatial resolution is limited. For example the authors of [Weyrich et al. 2009] generate a single dot of controlled reflectance properties, with size 3×3 cm. In this paper we present a method to construct spatially varying reflectance at a high resolution of up to 220dpi (dots per inch), albeit with a lower angular resolution. Figure 1 shows a prototype wafer fabricated using photolithography with spatially varying BRDFs, designed according to our method.

Bidirectional Reflectance Distribution Functions (BRDFs) are usually explained using the micro facets theory [Torrance and Sparrow 1967] and similar geometric optics extensions [Westin et al. 1992; Ashikhmin et al. 2000; Pont and Koenderink 2005; Oren and Nayar 1994; Koenderink et al. 1999; Wolff et al. 1998]. According to this model the surface is a collection of small, randomly scattered facets, each facet is assumed to have a simple reflectance function, often an ideal mirror. The distribution of facet normals inside a surface patch determines how light is reflected to different directions.

Recent approaches to appearance fabrication rely on geometric optics appearance models. For example, Weyrich et al. [2009] use the micro-facet model. Given a user specified BRDF, they compute a corresponding spatial arrangement of facets. Subsequent approaches seek continuous surfaces with desired redirection of light [Finckh et al. 2010; Papas et al. 2011]. The resolution of these designs is limited by the capabilities of the surface machining methods (e.g., CNC mills, engravers). However, it is important to note that even with better manufacturing methods the ability to scale down these designs is limited, since the reflectance will be dominated by diffraction. To illustrate this, we show in Figure 2 a target

ACM Reference Format

Levin, A., Glasner, D., Xiong, Y., Durand, F., Freeman, W., Matusik, W., Zickler, T. 2013. Fabricating BRDFs at High Spatial Resolution Using Wave Optics. *ACM Trans. Graph.* 32, 4, Article 144 (July 2013), 13 pages. DOI = 10.1145/2461912.2461981 <http://doi.acm.org/10.1145/2461912.2461981>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Copyright © ACM 0730-0301/13/07-ART144 \$15.00.
DOI: <http://doi.acm.org/10.1145/2461912.2461981>

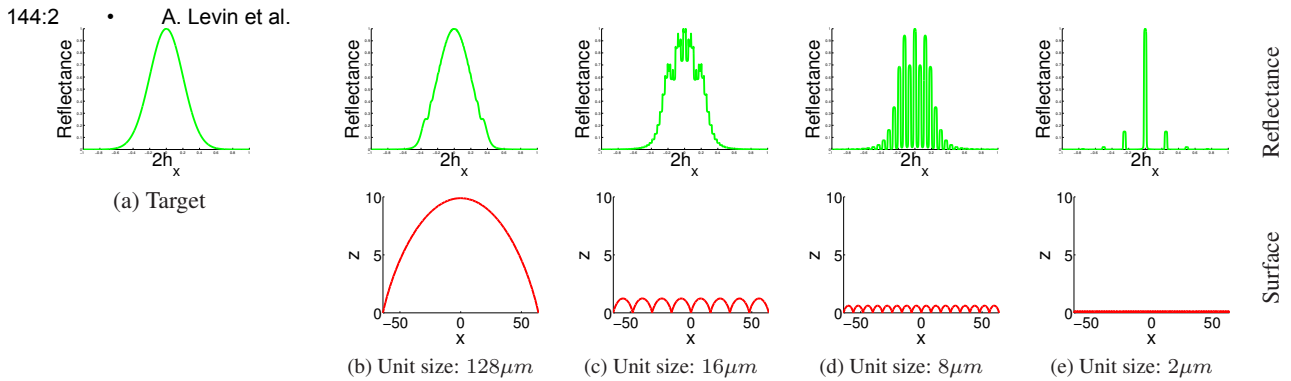


Figure 2: Limits of the geometric optics model. (a) Target reflectance. (b-e) The reflectance (top row) produced when we scale the size of the basic surface unit (second row). All surfaces have equal gradient histograms and should produce equal reflectance from a geometric optics viewpoint. Yet, for small structures wave effects play a major role and the actual reflectance is very different than the geometric prediction. When the basic unit size is as low as $2\mu\text{m}$ the reflectance is just a set of impulses. The simulation in this figure assumes the surface is viewed and illuminated from above, with the light source area occupying a subtended angle of 2° , and the surface area imaged by a dot is larger than $128\mu\text{m}$.

reflectance function and a surface with a corresponding normals histogram. From pure geometric optics considerations this surface should produce the exact target BRDF. Figure 2 (b-e) illustrates the same basic surface shape repeated at different scales. Since the gradient histogram is fixed at all scales, these should produce the same reflectance from a geometric optics viewpoint. However at micro-scales wave effects play a major role and the reflectance function is very different from the prediction of the geometric optics model.

This paper builds on previous work on deriving reflectance models based on wave optics [He et al. 1991; Nayar et al. 1991; Stam 1999; Cuypers et al. 2012] but focuses on practical fabrication procedures, which take into account the limitations of existing micro-fabrication techniques such as photolithography.

Using a photolithography process, it is significantly easier to fabricate surfaces which are composed of a small number of piecewise-flat layers (a 1D example is shown in Figure 4(a)), than continuous depth surfaces. The implication is that it is not possible to generate a variety of normal orientations – normals over the entire surface point at the same direction. The geometric optics model predicts that such piecewise-flat surfaces behave as mirrors, however as we demonstrate in this paper we can achieve a range of reflection effects with such a construction based solely on light interference.

Our method offers a more restricted family of reflectance functions compared, for example, with the method of Weyrich et al. [2009], its major advantage is high spatial resolution. We fabricate spatially varying BRDFs with a dot size of 0.112mm , 220 dpi. This is orders of magnitude improvement compared to the few centimeters wide dot units in previous approaches.

The complexity of the photolithography process increases with the number of different depth values since multiple etching passes are required. Thus, it is important to understand what reflectance functions are achievable as a function of the number of depth values. In particular our analysis and experiments show that a single etching pass (which produces only two different depth values), allows control over the BRDF up to a narrow specular spike at the mirror direction. Moreover, as few as three etching steps can provide full control and eliminate this specular spike. The prototype wafer which we experimented with in this paper was manufactured with a single etching pass, allowing for limited control over the specular spike.

While the price of lithography fabrication of the master prototype is high, it can be used as a template for stamping many successive copies quite inexpensively. This is similar to white light rainbow holograms mass-manufactured on credit cards [Iwata and Tsujiuchi 1974].

In this paper, we focus on photolithography, but in a companion technical report [Levin et al. 2013], we also explore an inexpensive alternative to micro-fabrication, in the form of metallic powders. Such powders are readily available at a variety of particle sizes and morphologies. Our analysis provides guidelines for the relation between the particles' shape and size and the reflectance functions they can produce.

1.1 Related Work

In addition to the work on reflectance design from the computer graphics community mentioned above this work is related to the following research areas.

Digital microoptics refers to the manipulation of light using surfaces equipped with piecewise flat microstructures at a discrete set of heights. This includes diffraction gratings used in spectrometers, and computer-generated holograms (CGHs) that are used, for example, to create laser point-pattern projectors, three-dimensional visualizations, and special-purpose diffusers [Kress and Meyrueis 2009]. Our work is related to these because we also seek patterns of microstructures that are optimized to provide a desired transformation (via reflection in our case) of an incident electromagnetic field, but while these designs are optimized to work under specific illumination conditions such as plane waves, we seek patterns which will work under general illumination conditions.

Holography [Yaroslavsky 2004; Benton and Bove 2008] is based on recording an interference pattern in a holographic medium between the field emitted from a 3D scene illuminated by a coherent wave and a coherent reference beam. When illuminating the hologram with the conjugate plane wave, the interference pattern preserved in the hologram recreates the wavefront corresponding to the original scene. The interference patterns can arise due to variation in absorbance or transmissivity, variations in the refractive index, and variations in the thickness of the medium. In rainbow holography the field is passed through a 1D horizontal slit, sacrificing the vertical parallax of the hologram but allowing the hologram to be viewed under natural white illumination. Computational holography works such as [Lucente and Galyean 1995; Matsushima 2005; Ziegler et al. 2007] attempt to compute the desired holographic plane pattern without illuminating a physical scene. When creating digital holograms, one typically designs for particular illumination conditions, whereas BRDF design aims for a specified behavior under general viewing and lighting geometries.

Micro-fabrication and Photolithography: The prototypes in this work were fabricated using photolithography. A good introduction to the subject is provided in [Sinzinger and Jahns 2006]. Photolithography uses light to transfer a geometric pattern from a photomask to a light-sensitive chemical photoresist on a substrate. A series of chemical treatments then either engraves (mostly using etching) the exposure pattern into, or enables deposition of a new material in the desired pattern upon, the material underneath the photo resist. Photolithography can create relatively small patterns but it still limits our minimal feature size to $2 - 3\mu m$. Its main disadvantages are that it requires a flat substrate to start with, and it is not very effective at creating shapes that are not flat. Multiple depth layers are usually obtained with successive etching iterations, but alignment between successive layers is challenging. Proper planning of the etching process allows k different depth values to be achieved using $\log_2 k$ etching passes [Walker and Jahns 1990]. Alternatively, gray scale lithography can partially expose the photoresist layer to multiple etching depths, but the complexity of calibrating the system increases with the number of different depth values. In both cases the complexity of the process grows drastically with the required number of depth layers.

Electron Beam Lithography uses an electron beam to write the features and thus can achieve a significantly higher spatial resolution at the price of a drastically slower throughput. It can also achieve better alignment between successive layers since the beam resolution is higher.

1.2 Outline

We first show how reflectance can be computed as a function of surface structure using a Fourier transform, and how it depends on the incoherence of the light source (Sec. 2). We then focus on the design of piecewise-flat surfaces that can be achieved using a photolithographic process and we derive the class of reflectance functions that can be obtained (Sec. 3). We demonstrate fabrication results with a variety of spatially varying reflectance effects (Sec. 4). Derivation details and proofs can be found in the Appendices.

2 Reflectance Under Extended Sources

To understand the formation of reflectance at micro scale one needs to take into account the wave nature of light rather than the geometric optics approximation. In this section we explain how the reflectance can be derived as a function of the surface shape. We start by summarizing the coherent illumination setting. Next, we analyze the effects which come into play in the more realistic setting of incoherent illumination.

We restrict the discussion to single material surfaces; therefore, the albedo is constant and the reflectance depends only on the surface shape. We consider metallic or silicone surfaces which reflect most of the incoming light, and the transmission, sub-scattering and Fresnel effects can be neglected. Our analysis ignores shadowing and masking. In practice the height differences in the surfaces we consider are very small, making occlusions negligible except for extreme grazing angles.

To simplify notation we consider a 2D case, in which all vectors are denoted on the x, z plane, omitting their y coordinate.

2.1 Point sources (coherent lighting)

The scattering of light from surfaces is derived in multiple textbooks [Beckmann and Spizzichino 1963; Goodman 1968] and was introduced to the graphics community by [He et al. 1991; Stam 1999]. We review the derivation in Appendix A. The important property, formally stated below, is that the reflectance is equal to the intensity of the Fourier transform of a signal whose phase is proportional to the depth of the surface.

We consider a roughly planar surface $\mathbf{S} = (x, z(x))$. \mathbf{S} is illuminated by a coherent point light source at direction $\mathbf{l} = (\mathbf{l}_x, \mathbf{l}_z)$ and viewed from direction $\mathbf{v} = (\mathbf{v}_x, \mathbf{v}_z)$. In our notation \mathbf{l} and \mathbf{v} are unit vectors. We denote by $\mathbf{h} = (\mathbf{h}_x, \mathbf{h}_z) = (\mathbf{l} + \mathbf{v})/2$ the half vector (note that unlike the standard convention [Rusinkiewicz 1998] \mathbf{h} is not a unit vector).

Let $a(x)$ denote the surface modulation function, that is, a signal whose phase is proportional to the surface height:

$$a(x) = e^{-ik2\mathbf{h}_z z(x)} \quad (1)$$

where $k = 2\pi/\lambda$ is the wave number and λ the wavelength of light.

The standard derivation states that under coherent illumination a camera, or an eye, observes the intensity (squared amplitude) of $\mathfrak{A}(\omega_x)$, the Fourier transform (Fraunhofer diffraction) of the surface modulation $a(x)$:

$$R_c(\mathbf{v}, \mathbf{l}) = \left| \gamma(\mathbf{v}, \mathbf{l}) \cdot \mathfrak{A} \left(\frac{2\mathbf{h}_x}{\lambda} \right) \right|^2. \quad (2)$$

The scalar function $\gamma(\mathbf{v}, \mathbf{l})$ encodes the foreshortening between \mathbf{v} and \mathbf{l} – when light arrives at an oblique angle, the same energy is spread over a wider surface area. The exact form of γ can be found in Appendix A Eq. (23).

The Fourier transform \mathfrak{A} is evaluated over a finite domain Δ_d corresponding to the physical dot area imaged onto a single pixel, that is:

$$\mathfrak{A}(\omega_x) \propto \int_{x \in \Delta_d} e^{-i2\pi\omega_x x} a(x) dx. \quad (3)$$

As a sanity check, consider the case of a perfectly flat surface, the reflectance is, as expected, a mirror. To see this, express the surface as $z(x) = s \cdot x$, and the normal $\mathbf{n} \propto (-s, 1)$. Thus $a(x)$ is a complex sinusoid:

$$a(x) = e^{-ik2\mathbf{h}_z s x} \quad (4)$$

and its Fourier transform \mathfrak{A} is an impulse at frequency $\omega_x = -2\mathbf{h}_z s/\lambda$. Using Eq. (2), the reflectance is an impulse at \mathbf{h} such that $\mathbf{h}_x = -\mathbf{h}_z s$, which is exactly the normal direction.

If we denote by $|E|^2$ the intensity of the input wave, we can relate reflectance to the conventional BRDF definition as in [Stam 1999]:

$$\text{BRDF}(\mathbf{v}, \mathbf{l}) = \frac{R_c(\mathbf{v}, \mathbf{l})}{\Delta_d |E|^2 \langle \mathbf{n}, \mathbf{l} \rangle \langle \mathbf{n}, \mathbf{v} \rangle}. \quad (5)$$

Keeping (5) in mind, we continue the analysis in the rest of this paper in terms of R and not in terms of the BRDF.

Except for the γ scale, the reflectance R depends only on the half vector. Our analysis omits γ and treats the reflectance as a function of \mathbf{h} alone. It can be shown that γ is approximately constant for small incident angles.

2.2 Extended sources (incoherent lighting)

In a real-world setting, lighting is created by extended sources. The electromagnetic fields these generate have limited spatial coherence, meaning that the cross-correlation of any such field at two well-separated surface points is small. For an incoherent incident field, the phrase *coherence area* is used to describe the range of separations between pairs of surface points for which significant cross-correlation (and therefore significant interference) exists [Goodman 1968]. We expect this coherence area to be inversely proportional to the angular extent of the source. An infinitesimal point source will produce a plane wave that is coherent everywhere (the previous section), and the coherence area will decrease as the source is

extended over a larger solid angle.¹

To analyze reflectance under extended sources we must define two parameters, shown in Figure 3(a), related to the size of the surface and the viewing geometry. The *dot pitch* Δ_d is the length of the surface that projects to a single pixel in the observer at the highest allowable viewing resolution. Put another way, it is the limit of how much an observer will be able to “zoom in” without discerning some of the microstructure that we design to aggregate to induce our reflectance. The second parameter is the *source area* Δ_a , which describes the extent of the (far-field) light source—the solid angle that it subtends relative to any point on surface S.

The core result of this section is a relatively simple analytical expression for the reflectance, at wavelength λ , when the dot pitch Δ_d is sufficiently large relative to the *coherence area* Δ_c , which is in turn equal to wavelength over source area: $\Delta_c = \lambda/\Delta_a$. This analysis provides two things. First, it gives an upper bound on the spatial reflectance resolution that can be achieved for a given lighting environment (i.e., a given source area); and second, it provides a means for designing and fabricating custom-reflectance at that spatial resolution.

As mentioned above, light is spatially incoherent because the light source is not limited to a point, but has an area and can be seen as composed of multiple *independent* point sources. The illumination reaching the surface arises from a small subtended angle Δ_a of the source area². The perceived incoherent reflectance is the averaged coherent *intensity* over all illumination directions in the source [Goodman 1968]:

$$\begin{aligned} R_{ic}(\mathbf{h}) &= \int_{u \in \Delta_a} \left| \mathfrak{A} \left(\frac{\mathbf{v}_x + (\mathbf{1}_x + u)}{\lambda} \right) \right|^2 du \\ &= \int_{u \in \Delta_a} \left| \mathfrak{A} \left(\frac{2\mathbf{h}_x + u}{\lambda} \right) \right|^2 du. \end{aligned} \quad (6)$$

Eq. (6) tells us that (up to a global scaling) the reflectance is the norm of the signal $\mathfrak{A}(\omega_x)$ multiplied by a rect function centered at $\omega_x = 2\mathbf{h}_x/\lambda$, with width Δ_a/λ . Formally,

$$R_{ic}(\mathbf{h}) = \|\mathfrak{A} \cdot \Pi_{2\mathbf{h}_x/\lambda, \Delta_a/\lambda}\|^2 \quad (7)$$

where,

$$\Pi_{\mu, c}(x) = \begin{cases} 1 & |(x - \mu)/c| < 0.5 \\ 0 & |(x - \mu)/c| > 0.5 \end{cases}$$

Alternatively, we can say that R_{ic} is the rect-filtered power spectrum:

$$R_{ic}(\mathbf{h}) = \mathfrak{J} \left(\frac{2\mathbf{h}_x}{\lambda} \right) \quad (8)$$

with

$$\mathfrak{J}(\omega_x) = |\mathfrak{A}(\omega_x)|^2 \otimes \Pi_{0, \Delta_a/\lambda}. \quad (9)$$

Note that similarly, the viewing direction has a finite size which is defined by the aperture. As a result the power spectrum is further blurred by the aperture area. While we omit this for simplicity, adjusting the derivation is straightforward.

Eq. (7) states that the reflected intensity is the norm of the Fourier transform \mathfrak{A} multiplied by a rect function. Using the convolution theorem and Parseval theorem, the norm of this signal equals (up to a global scaling) the norm of the spatial modulation signal a con-

¹Note that we are interested only in spatial coherence, since temporal coherence does not have a significant effect on the reflectance.

²We parameterize Δ_a using the x axis projection of unit length vectors.




Figure 3: Resolution parameters in incoherent illumination. (a) Angular resolution: illumination from an area light source covering a subtended angle Δ_a . Spatial resolution: The dot size Δ_d . (b) Incoherent reflectance is formed by applying coherent reflectance rules (Fourier transform) over coherent windows of size $\Delta_c = \lambda\Delta_a^{-1}$. The reflectance from all coherent areas inside a dot area Δ_d is averaged to form the incoherent reflectance.

involved with a phase shifted sinc:

$$R_{ic}(\mathbf{h}) = \|\mathbf{a} \otimes W\|^2 \quad (10)$$

with

$$W(x) = e^{-ik(2\mathbf{h}_x x)} W_0(x), \quad W_0(x) = \text{sinc}(x/\Delta_c). \quad (11)$$

$W(x)$ is the Fourier transform of the frequency window $\Pi_{2\mathbf{h}_x/\lambda, \Delta_a/\lambda}$. To express the spatial window size we reparameterize the angular resolution and define $\Delta_c = \lambda\Delta_a^{-1}$. Δ_c is the *spatial coherence area* mentioned above. It represents the spatial area over which the illumination wave interferes coherently with its shifted copies, and hence the coherent reflection rules of Sec. 2.1 apply.

Eqs. (7) and (10) highlight the known inverse relation between spatial resolution and angular resolution of the light source: when the angular resolution Δ_a increases the spatial resolution Δ_c decreases (the blur function W widens).

To better understand the formation of incoherent reflectance let us rewrite Eq. (10) as:

$$R_{ic}(\mathbf{h}) \propto \int_{x_0 \in \Delta_d} \left| \int \mathbf{a}(x - x_0) W_0(x) e^{-ik(2\mathbf{h}_x x)} dx \right|^2 dx_0. \quad (12)$$

According to Eq. (12), the incoherent reflectance is formed as follows: isolate an area of width Δ_c from a around x_0 , by multiplying with the window function W_0 (Figure 3(b)). Then apply the coherent illumination formula in that window. That is, compute a Fourier transform over the windowed area Δ_c by multiplying with the complex sinusoid $\exp(-ik(2\mathbf{h}_x x))$ of the corresponding viewing and illumination directions. The reflectance of different Δ_c sized windows inside Δ_d is averaged incoherently – only the power is averaged, without the phase.

Typical physical dimensions: To get some sense of typical sizes, an incandescent bulb in a room, which is a few centimeters wide and a few meters away from an observer usually occupies a subtended angle Δ_a of about $0.5^\circ - 2^\circ$. As an example a subtended angle of 1.8° , for which $\Delta_a = 1.8/180 \cdot \pi$, corresponds to a coherence area $\Delta_c = \lambda\Delta_a^{-1} = 16\mu\text{m}$, with a wavelength $\lambda = 0.5\mu\text{m}$. Lower coherence areas require larger area sources. At the extreme case making the coherence area as short as the wavelength requires a subtended angle of 180° . In our simulations we assumed that $\Delta_c = 16\mu\text{m}$ and targeted a dot pitch of $\Delta_d = 112\mu\text{m}$, which corresponds to a 220dpi resolution.




Figure 4: Parameterization and layout of a “staircase” piece-wise flat surface created by photolithography. Depths z_j are drawn from a discrete set of values.

3 Photolithographic Reflectance Fabrication

Having described the formation of reflectance our goal now is to design surfaces that give rise to reflectance functions of interest. Our analysis of the reflectance under incoherent illumination has important implications to the design of such surfaces. To achieve a desired reflectance one needs to design a surface $z(x)$ such that the rect-filtered power spectrum of its surface modulation function $a(x)$ produces the desired reflectance.

Moreover, when the desired dot pitch is sufficiently larger than the coherence area Δ_c our analysis shows that there is substantially more design freedom to be gained. In this case, we can combine within each dot Δ_d many distinct Δ_c -sized structures, thereby creating reflectance as their average. This is a useful result because even when only simple piecewise-planar microstructures are available on our substrate, we can still create interesting reflectance functions through averaging.

Our approach is as follows. We introduce a stochastic model for the placement of basic surface elements. The parameters of the model are designed such that when we sample repeatedly and independently, we obtain collections of Δ_c -sized structures whose reflectances *averaged* over Δ_d approximate the desired reflectance. Even if every particular sample is different, when $\Delta_d \gg \Delta_c$ we average multiple independent samples and by the law of large numbers, their mean approaches the expected value. In addition, by employing such a stochastic design we can avoid structural artifacts in the Fourier transform at scales comparable to or larger than Δ_d .

Given a set of fabricate-able microstructures and a target reflectance function $R_T(\mathbf{h})$, we will construct a surface in two steps. First, we define a stochastic sampling process for selecting microstructures that has the property that a sample will produce R_T in expectation. Then, we tile a surface with random i.i.d. samples from this process. In the remainder of this section we describe the set of microstructures that can be reliably created through photolithography, and then we describe sampling-based designs for two distinct families of reflectance functions.

Constrained surface model: As discussed in Sec. 1.1, fabricating arbitrary continuous height fields at micrometer scale is quite challenging. It is significantly easier to fabricate piecewise flat surfaces composed of a small number of layers. We take these considerations into account in our design.

We restrict our surface to be a combination of piecewise flat steps as illustrated in Figure 4(a). Formally, let x_j denote the center of the j 'th step, z_j its height and a_j its width. Taking photolithography limits into account we restrict step widths a_j to be larger than $2\mu\text{m}$ and the values of z_j to belong to a small discrete set (i.e. we allow only a small number of layers).




Figure 5: Constructing custom reflectance lobes: A window from an ensemble of piecewise flat depth surfaces (top row) and the resulting reflectance (second row). The reflection functions involve weak secondary lobe which become visible only at a contrast adjusted display (third row). (a-c) Isotropic Glossy reflectance at different widths. Note that due to the Fourier transform relation, wider features on the surface lead to narrower reflectance lobes. (d) Anisotropic reflectance achieved with rectangular surface features. We plot the reflected energy as a function of the x, y coordinates of the half vector, at the range $(\mathbf{h}_x, \mathbf{h}_y) \in [-0.2, 0.2] \times [-0.2, 0.2]$. This corresponds, e.g. to the case $\mathbf{l} = \mathbf{n} = (0, 0, 1)$, and $(\mathbf{v}_x, \mathbf{v}_y) \in [-0.4, 0.4] \times [-0.4, 0.4]$.

3.1 Custom Glossy Lobe

We start with a family of metallic reflectances. Given a user defined amount of gloss we seek a surface that reflects light over a lobe of directions around a mirror direction. Different lobe widths produce different amounts of gloss.

According to Eq. (8) the reflectance is proportional to the rect-filtered power spectrum of the flat surface. Thus, adjusting the step widths a_j allows us to inversely adjust the width of the Fourier lobe. Wider steps lead to a narrower Fourier transform (high gloss, mirror-like) and thinner steps lead to a wider Fourier lobe (low gloss, matte). Figure 5(a-c) shows a visualization of some step surfaces and the reflectance lobes they produce.

We use the following random process for tiling a segment with piecewise flat steps (see Figure 4(b)):

Sampling Process 1 Sample a sequence of lengths $\{a_j\}$ independently from a distribution p_a . The concatenation of these lengths determines the step centers x_j . To cover a two dimensional region sample two such length sequences and define rectangular regions as their outer product. Finally sample heights z_j for each rectangle segment independently from another distribution p_z .

As discussed above, to analyze the reflectance produced by this process it is enough to analyze the expected power spectrum.

Claim 1 Consider a surface modulation function $a(x)$, whose phase is a piecewise flat signal sampled according to Sampling Process 1. Denoting $\omega_x = 2\mathbf{h}_x/\lambda$, the expected reflectance is given by:

$$R_E(\mathbf{h}) = E_{p_a}[R_{a,E}(\mathbf{h})] \quad (13)$$

where $R_{a,E}$ denotes the expected spectrum from steps of width a :

$$R_{a,E}(\mathbf{h}) = (1 - |\tau|^2) \left[a^2 \cdot \text{sinc}^2 \left(\frac{\omega_x}{a-1} \right) \right] + |\tau|^2 a \Delta_c \Pi \left(\frac{\omega_x}{\Delta_c^{-1}} \right) \quad (14)$$

with $\tau = E_{p_z}[e^{-ik2\mathbf{h}_z z_j}]$, and E_{p_a}, E_{p_z} represent expectations with respect to the distributions p_a, p_z correspondingly.

The proof is in Appendix C. Intuitively, the first term in Eq. (14) comes from the fact that each flat step is box-shaped and thus has a squared-sinc power spectrum with width a^{-1} , inversely proportional to the step length. Since the step lengths a are random variables sampled from p_a , the resulting spectrum is the expectation of such squared-sinc functions with respect to p_a .

Note that if the mean of the height random variables p_z being averaged is not zero (i.e. $|\tau| > 0$) the expected reflectance includes a second “specular spike” term that is a very narrow impulse (of width Δ_c^{-1}) centered at the mirror-direction ($\mathbf{h}_x = 0$). We discuss this further in Sec. 3.2.

Surface tiling algorithm: For now, assume for simplicity that we can construct zero mean signals for which $\tau = 0$. We show how we use the above construction to generate a glossy lobe of interest.

We start with a target reflectance R_T . According to Claim 1, the expected reflectance R_E produced is the average, according to p_a , of the squared-sinc functions $R_{a,E}$. Thus we need to find a step width distribution p_a that will best approximate the target reflectance R_T . We do that by minimizing the squared norm:

$$\|E_{p_a}[R_{a,E}(\mathbf{h})] - R_T(\mathbf{h})\|^2 \quad (15)$$

which can be easily solved as a constrained least squares problem using a reasonable discretization on \mathbf{h} and a .

An example is shown in Figure 5 for the case of a Gaussian-shaped reflectance lobe,

$$R_T(\mathbf{h}) \propto e^{-\frac{\mathbf{h}_x^2}{2\sigma^2}}. \quad (16)$$

Our construction averages squared-sincs at different widths (corresponding to steps at different widths). By doing this we better approximate the smooth Gaussian fall-off and reduce the secondary lobes one would observe with a single squared-sinc (corresponding to equal width steps). This construction significantly decreases the secondary lobes which become visible only after a contrast adjustment (third row of Figure 5).

The horizontal and vertical step width distributions in our construction need not be the same, that is, the steps can be rectangles rather than squares. This allows us to design and fabricate anisotropic reflectance functions (Fig. 5(d)).

Physical step dimensions: To understand the connection between the frequency parameter ω_x and the lobe width, we recall that $\omega_x = 2\mathbf{h}_x/\lambda$. Let us denote by α the ratio between a step width a_0 and the wavelength $\alpha = a_0/\lambda$. The reflectance lobe in this case has the form $\text{sinc}(\mathbf{h}_x 2a_0/\lambda) = \text{sinc}(2\alpha\mathbf{h}_x)$, whose first zero occurs at $\mathbf{h}_x = 1/(2\alpha)$. Thus if $\mathbf{l} = \mathbf{n} = (0, 0, 1)$, the first sinc zero occurs at $2\mathbf{h}_x = \mathbf{v}_x = \alpha^{-1}$. That is, a surface composed of step widths a_0 reflects most light up to a maximal viewing angle of α^{-1} . As an example, with a feature size of $a_0 = 2\mu\text{m} \approx 4\lambda$, (for visible wavelengths $\lambda \approx 0.5\mu\text{m}$) the first sinc zero occurs at $2\mathbf{h}_x = \mathbf{v}_x = 0.25$ and the maximum width of the lobe is $15^\circ \times 2 = 30^\circ$. A fully diffuse lobe reflecting over a 180° angle requires feature sizes a_0 smaller than the wavelength λ .

Resolution tradeoffs: Before proceeding to the next family of reflectances we illustrate the behavior of the stochastic construction. As noted in Sec. 2 reflectance involves a tradeoff between spatial and directional resolution. To understand the main factors let us denote by $R_E(\mathbf{h})$ the theoretical expected reflectance predicted by Claim 1, and by $R_o(\mathbf{h}) = \mathcal{J}(2\mathbf{h}_x/\lambda)$ the power spectrum obtained in practice from a finite surface patch of size $\Delta_d \times \Delta_d$ sampled




Figure 6: Spatial and directional resolution tradeoffs. (a) The theoretical expected power spectrum. (b) Reflectance from a single region of $\Delta_d = \Delta_c = 16\mu\text{m}$. (c) Reflectance from a dot area of $\Delta_d = 0.16\text{mm}$, assuming a coherence area of $\Delta_c = 16\mu\text{m}$ produce a better approximation to the expected power spectrum. (d) Increasing the angular resolution and the coherence area such that $\Delta_c = 80\mu\text{m}$ while keeping the same dot area $\Delta_d = 0.16\text{mm}$ produces another a noisy reflectance.

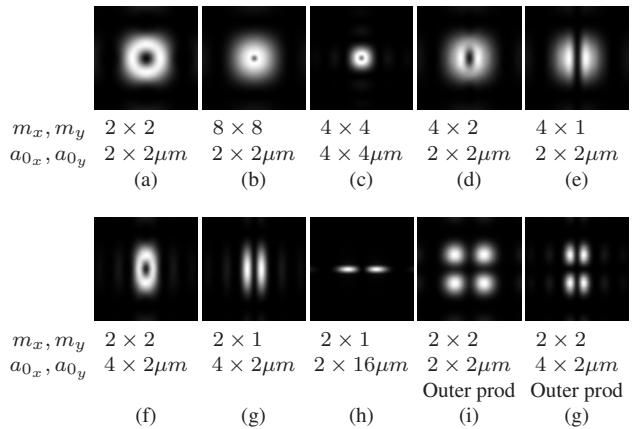


Figure 7: Anti-mirror reflectance functions with a Laplacian shape, achieved by selecting sequences of $m_x \times m_y$ fixed width rectangles of size $a_{0x} \times a_{0y}$ whose zero mean is enforced. As in Figure 5, we plot the reflected energy as a function of the x, y coordinates of the half vector, at the range $(\mathbf{h}_x, \mathbf{h}_y) \in [-0.2, 0.2] \times [-0.2, 0.2]$.

from the model. Let $n = (\Delta_d/\Delta_c)^2$ denote the number of coherent areas inside the dot pitch. How well does R_o approximate R_E ? Figure 6 provides a comparison. In Fig. 6(b-c) assume a coherent area of $\Delta_c = 16\mu\text{m}$ which corresponds to illumination angle resolution of about 2° . If $n = 1$ and the dot area is as low as $\Delta_d = \Delta_c$ (Fig. 6(b)), the reflectance function R_o we obtain is noisy, since it is generated by a single sample from the step signal model. In Fig. 6(c) we reduce spatial resolution such that $\Delta_d = 10\Delta_c$, and a dot area covers $n = 100$ coherent areas. In this setting, R_o provides a reasonable approximation to the expected reflectance R_E . In Fig. 6(d) we decrease the extent of the light source to 0.4° creating a 5 times wider coherence area $\Delta_c = 80\mu\text{m}$. Keeping the same dot area as in Fig. 6(c), now averages only $n = 4$ coherent areas. The reflectance R_o is again noisy since it involves less independent instantiations of the same model. In other words, to accommodate such a directional resolution we need to further reduce spatial resolution and average over a wider Δ_d .

When n is large, the law of large numbers takes effect, and any R_o sampled according to Sampling Process 1 will produce a good approximation to R_E . For modest n there is often a residual error from the random samples. One practical way to reduce this is to start from a random sample and optimize the design iteratively to reduce an objective function. In Appendix B we describe one such approach based on a genetic algorithm for minimizing $\|R_o - R_E\|$. In future work one could possibly construct better surfaces by exploring error diffusion or other half toning algorithms [Ulichney 1987].

3.2 Anti Mirror Reflectance

In this section we study another family of reflection functions, the anti-mirror reflectances. These are surfaces which reflect light in all directions *except the mirror direction*. Figure 7(a-c) visualizes some of these reflectance functions. This non-standard reflectance may be useful for camouflage or for creating reflective band-pass optical filters on the reflected scene, possibly enabling a form of a gradient camera [Tumblin et al. 2005].

To construct anti-mirror reflectance we select step heights such that the mean of the surface modulation function *equals* zero in every local region. This ensures that the DC component of the Fourier transform vanishes, and the reflectance has the shape of a wide ring with a dark “hole” in its center.

We do this by introducing a change to the step construction in Sec. 3.1.

Sampling Process 2 Consider a set of m possible depth values $\mathfrak{Z} = \{\zeta_1, \dots, \zeta_m\}$, with zero mean phase, $\tau = \frac{1}{m} \sum_{j=1}^m e^{-ik2h_z \zeta_j} = 0$. Tile the segment using sequences of m equal size steps of width a_0 , whose heights are sampled as a random permutation of the sequence $\zeta_1 \dots \zeta_m$.

Again, assuming $\Delta_d \gg \Delta_c$, the reflectance will approach the expected one $R_E(\mathbf{h}) = E[\mathfrak{J}(2\mathbf{h}_x/\lambda)]$. Thus, it is enough to compute the expected power spectrum.

Claim 2 The expected power spectrum of a surface modulation function sampled according to Sampling Process 2 is:

$$E[\mathfrak{J}(\omega_x)] \approx \left(1 - \text{sinc}\left(\frac{\tilde{\omega}_x}{(ma_0)^{-1}}\right)\right)^2 \text{sinc}\left(\frac{\omega_x}{a_0^{-1}}\right)^2 \quad (17)$$

with $\tilde{\omega}_x \equiv \omega_x \pmod{a_0^{-1}}$.

The proof is provided in Appendix C. The intuitive explanation is that the construction ensures that in every segment of length ma_0 we get exactly a zero mean signal. Note that the requirement $\tau = 0$ as used in the previous section implies that the DC component of the signal is zero mean, yet the variance of the DC component distribution could be large. In contrast, the current construction ensures that the DC component of *every* sample is zero, not only in expectation.

In practice the DC component is not a point but is supported over a sinc of finite width. This width is inversely proportional to the width of the primal support ma_0 . Therefore the power spectrum exhibits an angular “hole” of size $(ma_0)^{-1}$.

2D arrangements: Constructing anti-mirror reflectance in 2D provides a number of anisotropic choices. Each step is a rectangle of size $a_{0x} \times a_{0y}$ and the m constrained steps are arrangements of $m_x \times m_y$ units, satisfying $m_x \cdot m_y = m$.

The fully isotropic case $a_{0x} = a_{0y}$ and $m_x = m_y = \sqrt{m}$ is demonstrated in Fig. 7(a-c). By adjusting step size a_{0x}, a_{0y} we can control the width of the outer ring (the 2nd sinc in Eq. (17)), and by adjusting m_x, m_y we can control the width of the central zero hole (first sinc in Eq. (17)), whose width is $(ma_0)^{-1}$.

In Fig. 7(d-e) we use $a_{0x} = a_{0y}$ so the outer ring is isotropic. However $m_x \neq m_y$ results in an anisotropic inner hole. In Fig. 7(f-h) $a_{0x} \neq a_{0y}$ and the outer reflectance shape is anisotropic as well.

Finally in Fig. 7(i-g) the surface is constructed as an outer product of two one dimensional sequences. The horizontal dimension is a sequence of m_x length a_{0x} segments whose zero mean is enforced, and similarly the vertical sequence. As a result the zero mean is enforced at each axis independently, and the central black region takes a cross shape.




Figure 8: Reflectance under white illumination. The same step structures are demonstrated with two allowable depth values and with eight allowable depth values. For the 2 layers the zero mean property holds only for a single wavelength (in this case 550nm), and at other wavelengths a specular spike component is visible. The 8 layer design can achieve zero mean at all wavelengths.

To produce a larger family of reflectance functions, we can use any linear combination of the above anti-mirror functions, as well as the reflectance functions described in Sec. 3.1. This amounts to tiling the surface with samples from multiple desired models. This is essentially the strategy used in Sampling Process 1 and Eq. (15), which tiles the surface with step widths sampled from a distribution p_a instead of using a fixed a_0 width.

Zero mean and wavelength sensitivity: The reflectance functions predicted by Claims 1 and 2 produce pure glossy lobes, or pure anti-mirrors only if we can achieve a zero mean signal ($\tau = 0$). Otherwise, the Fourier transform includes an impulse around zero, which implies that the reflectance involves a specular spike, a narrow impulse response at the mirror direction (when $\mathbf{h} \approx 0$). Here we provide more detail about constructing zero mean surfaces. If the step heights $\{z_j\}$ can take arbitrary values achieving zero mean distributions is easy. However, to make photolithography fabrication practical the depth values z_j should belong to a small set of discrete values $\mathfrak{Z} = \{\zeta_1 \dots \zeta_k\}$. If the illumination is monochromatic at wavelength λ , and if we restrict our attention to small incident angles for which $\mathbf{h}_z \approx 1$, it is enough if z_j takes one of two discrete values $\mathfrak{Z} = \{\zeta_1 = 0, \zeta_2 = \lambda/4\}$, and then

$$e^{-ik2\zeta_1} = 1, \quad e^{-ik2\zeta_2} = e^{-\frac{i2\pi i}{\lambda} \frac{2\lambda}{4}} = -1.$$

This can be implemented with a single etching pass. For general white illumination, we need to find k depth values such that for every wavelength in the visible spectrum (i.e. every $\lambda \in [400, 700]nm$)

$$\frac{1}{k} \sum_{j=1}^k e^{-\frac{2\pi i}{\lambda} \mathbf{h}_z \zeta_j} \leq \epsilon \quad (18)$$

where ϵ is a small constant. Our simulations show that one can achieve good results with as low as 4–8 different depth values. Since photolithography processes can be designed such that k depth values can be created with only $\log_2 k$ etching passes [Walker and Jahns 1990], this implies that reasonable results under white light can be achieved using 2–3 passes.

In Figure 8 we show a simulation of the reflectance under white illumination. We compare an 8-layer construction with a 2-layer one. The former achieves zero mean at all wavelengths and as a result produces a relatively uniform white lobe. For the latter the mean vanishes only at a single wavelength. Thus, under white illumination the 2-layer construction shows a specular spike. In this example the zero mean was achieved at a green wavelength and the spike has a magenta color (since its green component vanishes). The Fourier transform scales as we change wavelength. As a result, the reflectance lobe of the red channel is a bit wider than that of the blue one and the periphery of the lobe is slightly reddish.

The prototype wafers we created for this paper were manufactured with a single etching pass, implying that at best, the specular spike vanishes at a single wavelength. In Sec. 4 we show the reflectance generated under white illumination as well as under monochromatic

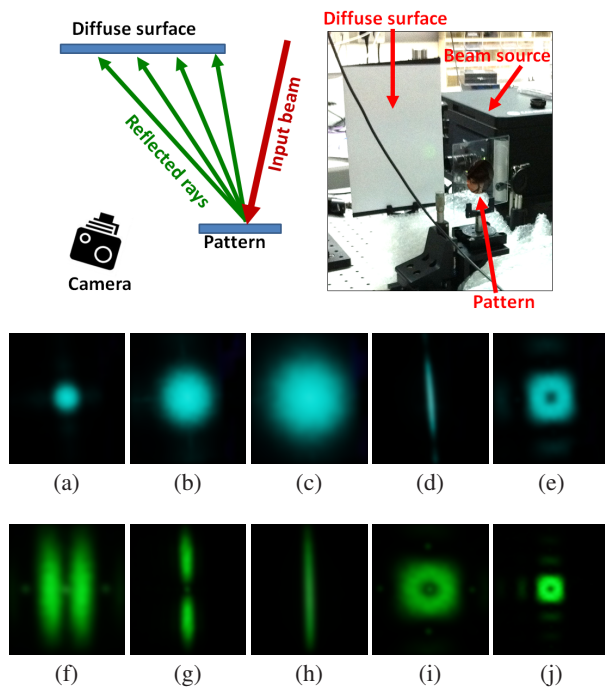


Figure 9: Our acquisition setup (top), and the reflectance of a few patterns fabricated according to our approach. (a-c) A Glossy lobe of 3 different widths. (d,h) Anisotropic reflectance. (e,i,j) Anti mirror reflectance, (f-g) Anisotropic anti-mirrors.

light and compare the results to the predictions of the simulation.

4 Results

Our prototype wafers were fabricated using photolithography as described in Sec. 1.1. Fabrication cost for a $10 \times 10\text{cm}$ wafer is on the order of \$4,000, produced by Photo-Science Inc. A wafer with a scanning electron microscope view inset is shown on the right of Fig. 1. Since we used a single etching process the generated surface has only two distinct depths. Following the discussion in Sec. 3.2 this implies that a specular spike is visible under white light but can be canceled at a certain wavelength.

Different regions on our prototype wafers are designed to generate a variety of different custom reflectance functions. In Figure 9 we show measurements of the reflectance from different points on the wafer corresponding to different reflections, including a glossy lobe at a variety of widths, an anisotropic reflectance, anti-mirrors, and anisotropic anti-mirrors. To capture these images we illuminate a point on the surface using a tunable laser monochromatic source with a narrow beam and image the reflectance on a diffuse surface.

For the examples shown in the first row the appropriate wavelength is $\lambda = 500\text{nm}$. Under illumination from such a monochromatic source the specular spike is not visible. The examples in the second row come from a wafer with more significant fabrication errors and thus some specular spike is visible at all wavelengths. The specular spike can be seen as the narrow impulse at the center of the reflectance images in the second row. The illumination in these images comes from a $\lambda = 550\text{nm}$ source.

In Figure 10 we show the reflected image of a planar dots pattern on our prototype wafer. The imaging setup is shown along with close-ups on parts of the wafer with different reflectance functions. The reflected dots are blurred by the surface reflectance function and the image resembles the shape of the designed reflectance function. We first took images of the wafer under white illumination, at which a specular spike is visible. As predicted by the simulation of Figure 8,

the specular spike has a magenta color since its green component is weak. Next we took images using a narrow band filter centered at a wavelength of $\lambda = 550\text{nm}$. As shown in the figure, the spike is significantly reduced in the monochromatic image.

To demonstrate the high spatial resolution which can be achieved with our approach, we fabricated surfaces with a dot (pitch) size of 0.112mm (220dpi). The fabrication process allows us to imprint a different design at each dot, giving rise to a different reflectance type at each dot. The physical dimensions of the entire pattern are only $4 \times 4\text{cm}$ which demonstrates a drastic resolution improvement compared to the 3cm dot units of prior geometric optics approaches [Weyrich et al. 2009]. An animation of the patterns under varying illumination directions is included in the supplementary video. Figures 1 and 11 show a few examples. In Figure 1 and in the first two columns of Figure 11 the patterns are composed of anisotropic reflectances with opposite orientations, and as a result the reflectance is inverted when the lighting direction changes from horizontal to vertical. Our high resolution makes dithering possible without significant visual resolution loss, and for the pattern in Figure 1 we dithered two anisotropic orientations to allow intersection of the horizontal and vertical reflectance regions. The pattern in the third column of Figure 11 demonstrates isotropic reflectance with different lobe widths. The different regions change their intensity as a function of the illumination angle. The background is a perfect mirror (a lobe of zero width) which appears bright when illuminated from the mirror direction (second row), and dark from any other direction (third row). The last column shows a pattern fabricated with an anti-mirror reflectance over a background with a narrow isotropic lobe. Thus, at small incident angles the kids are darker and the background is bright, while the kids region becomes brighter as the illumination angle increases.

The images in Figures 1 and 11 were captured under white illumination. Thus all reflectance functions involve a specular spike. The supplementary video (and the first example on the third column of Figure 11) show a bright spot which appears when the illumination direction is exactly the mirror direction.

5 Limitations

Photolithography fabrication suffers from a number of limitations summarized below.

1. The fabricated wafers produce a grayish reflectance and do not allow for color albedo.
2. Photolithography is restricted to planar surfaces and cannot be adapted to other object geometries.
3. The reflectance functions which can be designed using our approach have very high spatial resolution but they are not as general as the ones produced by previous geometric optics approaches [Weyrich et al. 2009].
4. Using optical lithography the smallest feature size is $2\mu\text{m}$. As we show in our analysis in Sec. 3.1, this implies that the maximum reflectance lobe width of our prototype is about 30° . Note that this limitation can be alleviated by using more sophisticated technologies such as Electron Beam Lithography.
5. Our analysis ignored shadowing and masking. However the etching depth in our prototype is on the order of a quarter of the wavelength, which is very small relative to the feature area (130nm etching depth compared to at least $2\mu\text{m}$ features). Thus occlusion effects are negligible, and the Fourier transform model is accurate except at extreme grazing angles. Shadowing and masking can be potentially accounted for using the models of [Sancer 1969; He et al. 1991]
6. A single layer prototype is cable of canceling the specular spike only at a single wavelength. A common fabrication error which leads to a non zero specular spike is ‘over-etching’.




Figure 10: Images of the reflectance of a planar dots pattern on our prototype wafer. We show closeups (extracted from different images) of parts of the wafer with different reflectance functions under white illumination and monochromatic 550nm light. The reflected dots are blurred by the surface reflectance function and the observed image resembles the shape of designed reflectance function. Columns 2-4 show isotropic reflectances of decreasing widths, the bottom dot is a mirror reflection. Columns 5-7 show anisotropic and anti-mirror reflectances.



Figure 11: Fabricated patterns viewed under white illumination at varying directions. First row: Design mask and the reflectance types involved (synthetic). Second and third rows: images of pattern under two different lighting directions (real). See supplementary video.

meaning that the etched features are wider than planned, and as a result a zero mean isn't achieved at any wavelength. In practice some of the wafers we fabricated were accurate enough (Fig. 9, 1st row) while others suffered from such over etching errors (Fig. 9, 2nd row).

6 Conclusions

Surface reflectance at micro scale is dominated by wave effects and can not be explained using a geometric ray model. In this work we exploited interference effects while accounting for the capabili-

ties of existing micro fabrication technology. We were able to produce custom designed reflectance functions with high spatial resolution of 220dpi, which provides an order of magnitude improvement compared to the few centimeter wide features of previous geometric optics approaches.

Acknowledgments: We thank E. Adelson for insightful discussions, Y. Wang, O. Bitton, D. Oron and R. Tenne for experimental help, M. Irani and M. Zontak for helpful comments. Funding was provided by BSF, ERC, Intel ICRI-CI and NSF CGV 1116303.

References

- ALLDRIN, N., AND KRIEGMAN., D. 2006. A planar light probe. In *CVPR*, 2324–2330.
- ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. 2000. A microfacet-based BRDF generator. In *ACM SIGGRAPH*, 65–74.
- BECKMANN, P., AND SPIZZICHINO, A. 1963. *The scattering of electromagnetic waves from rough surfaces*. International series of monographs on electromagnetic waves. Pergamon Press.
- BENTON, S. A., AND BOVE, V. M. 2008. *Holographic Imaging*. Wiley-Interscience.
- CUYPERS, T., HABER, T., BEKAERT, P., OH, S. B., AND RASKAR, R. 2012. Reflectance model for diffraction. *ACM Trans. Graph.* 31, 5, 122.
- DONG, Y., WANG, J., PELLACINI, F., TONG, X., AND GUO, B. 2010. Fabricating spatially-varying subsurface scattering. *ACM Trans. Graph.* 29, 4 (July), 62:1–62:10.
- FINCKH, M., DAMMERTZ, H., AND LENSCH, H. P. A. 2010. Geometry construction from caustic images. In *ECCV*, Springer-Verlag, 464–477.
- GOODMAN, J. W. 1968. *Introduction to Fourier Optics*. McGraw-Hill Book Company.
- HAŠAN, M., FUCHS, M., MATUSIK, W., PFISTER, H., AND RUSINKIEWICZ, S. 2010. Physical reproduction of materials with specified subsurface scattering. *ACM SIGGRAPH* 29, 3, 61:1–61:10.
- HE, X. D., TORRANCE, K. E., SILLION, F. X., AND GREENBERG, D. P. 1991. A comprehensive physical model for light reflection. *SIGGRAPH* 25, 4, 175–186.
- IWATA, F., AND TSUJIUCHI, J. 1974. Characteristics of a photorealist hologram and its replica. *Appl. Opt.* 13, 6 (Jun), 1327–1336.
- JOHNSON, M. K., COLE, F., RAJ, A., AND ADELSON, E. H. 2011. Microgeometry capture using an elastomeric sensor. *ACM SIGGRAPH* 30, 4, 46:1–46:8.
- KISER, T., EIGENSATZ, M., NGUYEN, M. M., BOMPAS, P., AND PAULY, M. 2012. Architectural caustics controlling light with geometry. In *Advances in Architectural Geometry*.
- KOENDERINK, J., DOORN, A. V., DANA, K., AND NAYAR, S. 1999. Bidirectional Reflectance Distribution Function of Thoroughly Pitted Surfaces. *ICCV* 31, 2/3, 129–144.
- KRESS, B. C., AND MEYRUEIS, P. 2009. *Dynamic Digital Optics*. John Wiley & Sons, Ltd, 217–252.
- LEVIN, A., GLASNER, D., XIONG, Y., DURAND, F., FREEMAN, W., MATUSIK, W., AND ZICKLER, T. 2013. High spatial resolution BRDFs with metallic powders using wave optics analysis. *MIT CSAIL TR 2013-008*.
- LUCENTE, M., AND GALYEAN, T. A. 1995. Rendering interactive holographic images. In *SIGGRAPH*, 387–394.
- MALZBENDER, T., SAMADANI, R., SCHER, S., CRUME, A., DUNN, D., AND DAVIS, J. 2012. Printing reflectance functions. *ACM Trans. Graph.* 31, 3, 20:1–20:11.
- MATSUSHIMA, K. 2005. Computer-generated holograms for three-dimensional surface objects with shade and texture. *Appl. Opt.* 44, 22 (Aug), 4607–4614.
- MATUSIK, W., AJDIN, B., GU, J., LAWRENCE, J., LENSCH, H. P., PELLACINI, F., AND RUSINKIEWICZ, S. 2009. Printing spatially-varying reflectance. *ACM SIGGRAPH Asia* 28, 5 (Dec.), 128:1–128:9.
- NAYAR, S., KIKEUCHI, AND KANADE, T. 1991. Surface Reflection: Physical and Geometrical Perspectives. *PAMI* 13, 7 (Jul), 611–634.
- OREN, M., AND NAYAR, S. 1994. Generalization of Lambert's Reflectance Model. In *ACM SIGGRAPH*, 239–246.
- PAPAS, M., JAROSZ, W., JAKOB, W., RUSINKIEWICZ, S., MATUSIK, W., AND WEYRICH, T. 2011. Goal-based caustics. *Eurographics* 30, 2 (Apr.), 503–511.
- PATOW, G., AND PUEYO, X. 2005. A survey of inverse surface design from light transport behavior specification. *Comput. Graph. Forum* 24, 4, 773–789.
- PATOW, G., PUEYO, X., AND VINACUA, A. 2007. User-guided inverse reflector design. *Comput. Graph.* 31, 3 (June), 501–515.
- PONT, S. C., AND KOENDERINK, J. J. 2005. Reflectance from locally glossy thoroughly pitted surfaces. *Computer Vision and Image Understanding* 98, 2, 211–222.
- REN, P., WANG, J., SNYDER, J., TONG, X., AND GUO, B. 2011. Pocket reflectometry. In *ACM SIGGRAPH*, 45:1–45:10.
- RUSINKIEWICZ, S. 1998. A new change of variables for efficient BRDF representation. In *Rendering Techniques (Proc. Eurographics Workshop on Rendering)*.
- SANCER, M. 1969. Shadow-corrected electromagnetic scattering from a randomly rough surface. *IEEE Transactions on Antennas and Propagation* 17, 5 (sep), 577 – 585.
- SINZINGER, S., AND JAHNS, J. 2006. *Microoptics*. John Wiley and Sons.
- STAM, J. 1999. Diffraction shaders. In *ACM SIGGRAPH*, 101–110.
- TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory for off-specular reflection from roughened surfaces. *J. Opt. Soc. Am.* 57, 9 (Sep), 1105–1112.
- TUMBLIN, J., AGRAWAL, A., AND RASKAR, R. 2005. Why i want a gradient camera. In *CVPR*, vol. 1, IEEE, 103–110.
- ULICHNEY, R. 1987. *Digital halftoning*. MIT press.
- WALKER, S. J., AND JAHNS, J. 1990. Array generation with multilevel phase gratings. *J. Opt. Soc. Am. A* 7, 8 (Aug), 1509–1513.
- WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *ACM SIGGRAPH*, 255–264.
- WEYRICH, T., DENG, J., BARNES, C., RUSINKIEWICZ, S., AND FINKELSTEIN, A. 2007. Digital bas-relief from 3D scenes. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 26, 3 (Aug.).
- WEYRICH, T., PEERS, P., MATUSIK, W., AND RUSINKIEWICZ, S. 2009. Fabricating microgeometry for custom surface reflectance. *ACM SIGGRAPH* 28, 3 (Aug.), 32:1–32:6.
- WOLFF, L., NAYAR, S., AND OREN, M. 1998. Improved Diffuse Reflection Models for Computer Vision. *IJCV* 30, 1 (Oct), 55–71.
- YAROSLAVSKY, L. 2004. *Digital Holography and Digital Image Processing*. Kluwer Academic Publishers.
- ZIEGLER, R., BUCHELI, S., AHRENBERG, L., MAGNOR, M. A., AND GROSS, M. H. 2007. A bidirectional light field - hologram transform. *Comput. Graph. Forum* 26, 3, 435–446.

A Reflectance and Scattering From Rough Surfaces

We provide a review of the standard derivation of the scattering of light from rough surfaces [Stam 1999; Beckmann and Spizzichino 1963]. The main useful property we explain below is that the reflectance equals to the Fourier transform of a signal whose phase is proportional to the depth of the surface.

Figure 12 illustrates the basic setup. We consider a roughly planar surface $\mathbf{S} = (x, z(x))$. Without loss of generality assume its tangent is coincident with the $z = 0$ plane. \mathbf{S} is illuminated by a coherent point light source at direction $\mathbf{l} = (\mathbf{l}_x, \mathbf{l}_z)$ and viewed from direction $\mathbf{v} = (\mathbf{v}_x, \mathbf{v}_z)$. In our notation \mathbf{l} and \mathbf{v} are unit vectors, while the actual position of the light and the viewer are $\mathbf{p}_l = r_l \mathbf{l}$ and $\mathbf{p}_v = r_v \mathbf{v}$ with $r_l, r_v \rightarrow \infty$. The coherent source emits a spherical wave, thus the complex field reaching a 3D point \mathbf{p} is

$$U_0(\mathbf{p}) = \frac{1}{\|\mathbf{p} - \mathbf{p}_l\|} e^{ik\|\mathbf{p} - \mathbf{p}_l\|} \quad (19)$$

where $k = 2\pi/\lambda$ is the wave number and λ the wavelength of light. The surface reflects the incoming illumination (Figure 12(a)). Based on Huygens' principle, each surface point $\mathbf{p}_s = (x, z(x))$ emits a new spherical wave multiplied by the incoming field $U_0(\mathbf{p}_s)$. The reflected field at a viewing point \mathbf{p}_v is the superposition of all these spherical waves (Kirchhoff integral):

$$\begin{aligned} U_r(\mathbf{p}_v) &= \gamma \int \frac{1}{\|\mathbf{p}_v - \mathbf{p}_s\|} U_0(\mathbf{p}_s) e^{ik\|\mathbf{p}_v - \mathbf{p}_s\|} d\mathbf{s} \\ &= \gamma \int \frac{1}{\|\mathbf{p}_l - \mathbf{p}_s\| \|\mathbf{p}_v - \mathbf{p}_s\|} e^{ik(\|\mathbf{p}_l - \mathbf{p}_s\| + \|\mathbf{p}_v - \mathbf{p}_s\|)} d\mathbf{s} \end{aligned} \quad (20)$$

where γ is a scale factor related to the foreshortening between \mathbf{v} , \mathbf{l} –when light arrives at an oblique angle, the same energy is spread over a wider surface area.

When the illumination and viewing points are sufficiently far, a Taylor expansion shows that the distances can be approximated as

$$\|\mathbf{p}_s - \mathbf{p}_l\| \approx r_l - \mathbf{l}^T \mathbf{p}_s, \quad \|\mathbf{p}_s - \mathbf{p}_v\| \approx r_v - \mathbf{v}^T \mathbf{p}_s. \quad (21)$$

The projection $\mathbf{l}^T \mathbf{p}_s$ is visualized in Figure 12(b).

If we substitute Eq. (21) in Eq. (20) and rely on the fact that $\frac{1}{\|\mathbf{p}_l - \mathbf{p}_s\|}$, $\frac{1}{\|\mathbf{p}_v - \mathbf{p}_s\|}$ are constant to a first order approximation, we can express the reflected field at \mathbf{p}_v up to a multiplicative constants as:

$$\begin{aligned} U_r(\mathbf{p}_v) &= \gamma \int e^{-ik((\mathbf{v} + \mathbf{l})^T \mathbf{p}_s)} d\mathbf{s} \\ &= \gamma \int e^{-ik((\mathbf{v}_x + \mathbf{l}_x)x + (\mathbf{v}_z + \mathbf{l}_z)z(x))} dx. \end{aligned} \quad (22)$$

Some derivation, omitted for brevity (see e.g. [Goodman 1968; Stam 1999; Beckmann and Spizzichino 1963]) shows that the proper foreshortening scaling is

$$\gamma = \frac{(1 + \langle \mathbf{l}, \mathbf{v} \rangle)}{\lambda(\mathbf{v}_z + \mathbf{l}_z)}. \quad (23)$$

In practice, due to the finite aperture of the lens or the finite size of the eye, the integral in Eq. (22) should be evaluated only over a finite spatial support Δ_d , thus Eq. (22) reads:

$$U_r(\mathbf{p}_v) = \gamma \int_{x \in \Delta_d} e^{-ik((\mathbf{v}_x + \mathbf{l}_x)x + (\mathbf{v}_z + \mathbf{l}_z)z(x))} dx. \quad (24)$$




Figure 12: Illumination setup: A surface illuminated from direction \mathbf{l} and viewed from \mathbf{v} . (a) A point light source at \mathbf{p}_l emits a spherical wave. When that wave interacts with the surface, each surface point emits a new spherical wave centered at the interaction point. (b) When the source is sufficiently far the spherical wave approaches a plane wave. The phase difference between the field at a surface point \mathbf{p}_s and the field at the origin is the projection of \mathbf{p}_s on \mathbf{l} , given by $\mathbf{l}^T \mathbf{p}_s$.

To understand Eq. (24) let $\mathbf{h} = (\mathbf{h}_x, \mathbf{h}_z) = (\mathbf{l} + \mathbf{v})/2$ denote the half vector (note that unlike the standard convention \mathbf{h} is not a unit vector). We also denote by $a(x)$ the surface modulation function, that is, a signal whose phase is proportional to the surface height:

$$a(x) = e^{-ik2\mathbf{h}_z z(x)}. \quad (25)$$

Thus, the reflected field is proportional to $\mathfrak{A}(\omega_x)$, the Fourier transform (Fraunhofer diffraction [Goodman 1968]) of the surface modulation $a(x)$ restricted to the domain Δ_d :

$$U_r(\mathbf{p}_v) = \gamma \int_{x \in \Delta_d} e^{-ik(2\mathbf{h}_x x)} a(x) dx \quad (26)$$

$$= \gamma \mathfrak{A}\left(\frac{2\mathbf{h}_x}{\lambda}\right). \quad (27)$$

We observe the intensity of the field and thus the reflectance under coherent illumination is

$$R_c(\mathbf{v}, \mathbf{l}) = \left| \gamma \mathfrak{A}\left(\frac{2\mathbf{h}_x}{\lambda}\right) \right|^2. \quad (28)$$

B Surface construction algorithm

As mentioned in Sec. 3.1 we attempt to find surface tiling R_o whose power spectrum provides a good approximation to the theoretical expectation R_E . We preform the optimization using a genetic algorithm search.

The algorithm maintains a population of $\{R_o^j\}_{j=1}^m$ candidates, which is initialized using random samples from Sampling Process 1. For each candidate we define a score:

$$e^j = \|R_o^j - R_E\|. \quad (29)$$

We iterate the following process

1. Mutate each sample R_o^j by resampling $\epsilon\%$ of its parameters, either resampling step heights z_j or widths a_j , to obtain a new set of m candidates $\{R_o^{j\text{new}}\}_{j=1}^m$.

2. Evaluate the quality of each new sample:

$$e^{j\text{new}} = \|R_o^{j\text{new}} - R_E\| \quad (30)$$

3. Select the best m candidates out of the $2m$ members in the union of old and new populations.

The genetic optimization produces multiple R_o^j candidates which provide reasonable approximations to R_E . Each R_o^j covers an area $\Delta_d \times \Delta_d$. To reduce periodic artifacts we tile neighboring $\Delta_d \times \Delta_d$ regions with different R_o structures.

We assumed a coherence area $\Delta_c = 16\mu m$ and attempted to tile dots (pixels) of size $\Delta_d = 112\mu m$, that is a dot of size $\Delta_d \times \Delta_d$ includes 7×7 coherent areas.

C Expected Reflectance Derivation

Below we provide complete proofs for the expected reflectance claims mentioned in the main paper body. To connect to the proof, we reformulate some of them slightly differently.

Claim 1 Consider a surface modulation function \mathbf{a} , whose phase is a piecewise flat signal sampled according to Sampling Process 1. The reflectance is given by

$$R_E(\mathbf{h}) = E[\mathcal{J}(2\mathbf{h}_x/\lambda)] \quad (31)$$

with

$$E[\mathcal{J}(\omega_x)] \approx \quad (32)$$

$$(1 - |\tau|^2)E_{p_a} \left[a^2 \cdot \text{sinc}^2 \left(\frac{\omega_x}{a^{-1}} \right) \right] + |\tau|^2 \varrho \Pi \left(\frac{\omega_x}{\Delta_c^{-1}} \right) \quad (33)$$

where $\tau = E_{p_z} [e^{-ik2\mathbf{h}_z z_j}]$ and $\varrho = \Delta_c E_{p_a} [a]$.

Proof of claim 1: To compute $E[\mathcal{J}(\omega_x)]$ we first compute the expectation of the unblurred power spectrum $E[|\mathfrak{A}(\omega_x)|^2]$, and then, by the linearity of the expectation:

$$E[\mathcal{J}(\omega_x)] = E[|\mathfrak{A}(\omega_x)|^2] \otimes \Pi_{0, \Delta_a/\lambda}. \quad (34)$$

To simplify analysis assume we sample a fixed number of n steps $n = \Delta_d/E_{p_a} [a]$. The exact dot area cover by a random sample of n steps is a random variable $\tilde{\Delta}_d$ whose expectation is Δ_d :

$$\tilde{\Delta}_d = \sum_{j=1}^n a_j, \quad E[\tilde{\Delta}_d] = \Delta_d = nE_{p_a} [a]. \quad (35)$$

We compute the expectation in two steps, first we think of step centers and widths x_j, a_j as observed values and compute the expected spectrum over all possible depth assignments z_j sampled independently from p_z . In the second step we compute expectations with respect to x_j, a_j and $\tilde{\Delta}_d$.

We define a set of random variables associated with the surface modulation function at each of the steps:

$$\psi_j = e^{-ik2\mathbf{h}_z z_j} \Pi_{x_j, a_j}. \quad (36)$$

Their corresponding Fourier transforms are denoted Ψ_j . Thus,

$$\mathbf{a} = \sum_{j=1}^n \psi_j \quad \mathfrak{A} = \sum_{j=1}^n \Psi_j. \quad (37)$$

We can express the expected power spectrum as a sum of two terms, variance and squared mean. Using the fact that for a fixed sample of x_j, a_j, Ψ_j are independent random variables we get:

$$E_{p_z} [|\mathfrak{A}|^2] = \sum_{j=1}^n E_{p_z} [|\Psi_j - E_{p_z} [\Psi_j]|^2] + \left| \sum_{j=1}^n E_{p_z} [\Psi_j] \right|^2. \quad (38)$$

The variance term (first term in Eq. (38)) gives rise to the sinc^2 lobe part of the reflectance function, the first term of Eq. (33), while the squared mean term (second term in Eq. (38)) leads to the impulse part of the reflectance function, the second term of Eq. (33).

We start by computing the variance term. Since ψ_j is a rect centered at x_j its Fourier transform is a phase shifted sinc

$$\Psi_j = e^{-2\pi i(x_j \omega_x + \frac{2\mathbf{h}_z z_j}{\lambda})} a_j \text{sinc}(\omega_x a_j). \quad (39)$$

Taking expectations with respect to p_z we note

$$E_{p_z} [\Psi_j] = \tau e^{-2\pi i(x_j \omega_x)} a_j \text{sinc}(\omega_x a_j). \quad (40)$$

Thus

$$\begin{aligned} E_{p_z} [|\Psi_j - E_{p_z} [\Psi_j]|^2] &= E_{p_z} \left[\left| e^{-ik2\mathbf{h}_z z_j} - \tau \right|^2 \right] a_j^2 \text{sinc}^2(\omega_x a_j) \\ &= \left(E_{p_z} \left[\left| e^{-ik2\mathbf{h}_z z_j} \right|^2 \right] - 2\tau^* E_{p_z} \left[e^{-ik2\mathbf{h}_z z_j} \right] + |\tau|^2 \right) a_j^2 \text{sinc}^2(\omega_x a_j) \\ &= (1 - |\tau|^2) a_j^2 \text{sinc}^2(\omega_x a_j). \end{aligned} \quad (41)$$

The expression in Eq. (41) no longer depends on the centers x_j . After taking the expectation with respect to a_j we see that the variance term is indeed equal to the first term in Equation Eq. (33):

$$\sum_{j=1}^n E [|\Psi_j - E[\Psi_j]|^2] = n(1 - |\tau|^2) E_{p_a} \left[a^2 \cdot \text{sinc}^2 \left(\frac{\omega_x}{a^{-1}} \right) \right]. \quad (42)$$

We now compute the second part of Eq. (38), the squared mean. Due to the linearity of the Fourier transform $\sum_{j=1}^n E[\Psi_j]$ is the Fourier transform of $\sum_{j=1}^n E[\psi_j]$. For each sample of x_j, a_j

$$E_{p_z} [\psi_j] = E_{p_z} [e^{-ik2\mathbf{h}_z z_j} \Pi_{x_j, a_j}] = \tau \Pi_{x_j, a_j} \quad (43)$$

since the steps Π_{x_j, a_j} are disjoint and their union covers the full surface area, $\sum_j E_{p_z} [\psi_j]$ is a flat signal of amplitude τ and width $\tilde{\Delta}_d$. The mean is also constant with respect to the selection of x_j, a_j . Thus, the Fourier transform of the mean over a spatial support Δ_d is a narrow sinc around the zero frequency:

$$\sum_{j=1}^n E[\Psi_j] = \tau \tilde{\Delta}_d \text{sinc} \left(\frac{\omega_x}{\tilde{\Delta}_d^{-1}} \right). \quad (44)$$

Substituting Eqs. (44) and (42) in Eq. (38) yields

$$\begin{aligned} E [|\mathfrak{A}(\omega_x)|^2] &\approx \quad (45) \\ n(1 - |\tau|^2) E_{p_a} \left[a^2 \cdot \text{sinc}^2 \left(\frac{\omega_x}{a^{-1}} \right) \right] &+ |\tau|^2 \tilde{\Delta}_d^2 \text{sinc}^2 \left(\frac{\omega_x}{\tilde{\Delta}_d^{-1}} \right). \end{aligned} \quad (46)$$

To conclude the proof, we need to compute $E[\mathcal{J}(\omega_x)]$ from $E[|\mathfrak{A}(\omega_x)|^2]$. We follow Eq. (34) and blur $E[|\mathfrak{A}(\omega_x)|^2]$ with a rect of width $\Delta_c^{-1} = \Delta_a/\lambda$. The first term of Eq. (46) is relatively smooth and we assume it does not change that much by blurring, that is:

$$\begin{aligned} n(1 - |\tau|^2) E_{p_a} \left[a^2 \cdot \text{sinc}^2 \left(\frac{\omega_x}{a^{-1}} \right) \right] &\otimes \Pi_{0, \Delta_c^{-1}} \approx \\ n(1 - |\tau|^2) E_{p_a} \left[a^2 \cdot \text{sinc}^2 \left(\frac{\omega_x}{a^{-1}} \right) \right]. \end{aligned} \quad (47)$$

The second term of Eq. (46) is a very narrow squared sinc of width $\tilde{\Delta}_d$. Assuming that $\tilde{\Delta}_d \gg \Delta_c$, then $\Delta_c^{-1} \gg \tilde{\Delta}_d^{-1}$. Hence, convolving the narrow sinc with the wider rect of width Δ_c^{-1} results roughly in a rect of width Δ_c^{-1}

$$|\tau|^2 \tilde{\Delta}_d^2 \text{sinc}^2 \left(\frac{\omega_x}{\tilde{\Delta}_d^{-1}} \right) \otimes \Pi_{0, \Delta_c^{-1}} \approx |\tau|^2 \tilde{\Delta}_d^2 \frac{\tilde{\Delta}_d^{-1}}{\Delta_c^{-1}} \Pi_{0, \Delta_c^{-1}}. \quad (48)$$

We further use $E[\tilde{\Delta}_d] = \Delta_d = n E_{p_a}[a]$ (Eq. (35))

$$E \left[|\tau|^2 \tilde{\Delta}_d^2 \text{sinc}^2 \left(\frac{\omega_x}{\tilde{\Delta}_d^{-1}} \right) \otimes \Pi_{0, \Delta_c^{-1}} \right] \approx n |\tau|^2 \Delta_c E_{p_a}[a] \Pi_{0, \Delta_c^{-1}}. \quad (49)$$

Combining Eqs. (47) and (49) yields the desired result (Eq. (33)) up to a global scaling factor:

$$E[\mathfrak{J}(\omega_x)] \approx n(1 - |\tau|^2) E_{p_a} \left[a^2 \cdot \text{sinc}^2 \left(\frac{\omega_x}{a^{-1}} \right) \right] + n |\tau|^2 \rho \Pi \left(\frac{\omega_x}{\Delta_c^{-1}} \right). \quad (50)$$

□

Claim 2 Let $\mathfrak{J} = \{\zeta_1, \dots, \zeta_m\}$, be a set of m depth values such that $\tau = \frac{1}{m} \sum_{j=1}^m e^{-ik2\mathbf{h}_z \zeta_j} = 0$. The expected power spectrum of a set of m steps of width a_0 , whose depths are a random permutation of $\{\zeta_1, \dots, \zeta_m\}$ is

$$E[\mathfrak{J}(\omega_x)] \approx \left(1 - \text{sinc} \left(\frac{\tilde{\omega}_x}{(ma_0)^{-1}} \right) \right)^2 \text{sinc} \left(\frac{\omega_x}{a_0^{-1}} \right)^2 \quad (52)$$

with $\tilde{\omega}_x \equiv \omega_x \pmod{a_0^{-1}}$

Proof of claim 2: As in the proof of claim 1, $E[\mathfrak{J}]$ can be computed by blurring $E[|\mathfrak{A}|^2]$. However in this case $E[|\mathfrak{A}|^2]$ does not have much high frequency content. Thus it is enough to compute $E[|\mathfrak{A}|^2]$ and approximate $E[\mathfrak{J}] \approx E[|\mathfrak{A}|^2]$.

The Fourier transform is given by

$$\mathfrak{A}(\omega_x) = \sum_j e^{-2\pi i(x_j \omega_x + \frac{2\mathbf{h}_z z_j}{\lambda})} a_0 \text{sinc}(\omega_x a_0). \quad (53)$$

Where x_j denotes the center of the j 'th step. For ease of notation we will assume $\mathbf{h}_z \approx 1$ as in the case of small incident angles.

Let us denote

$$\mathfrak{A}_0(\omega_x) = \sum_j e^{-2\pi i(x_j \omega_x + \frac{2z_j}{\lambda})} \quad (54)$$

then $\mathfrak{A}(\omega_x) = \mathfrak{A}_0(\omega_x) a_0 \text{sinc}(\omega_x a_0)$. Since the sinc has a constant width a_0^{-1} and is not a random variable, to prove the claim we need to show that

$$E[|\mathfrak{A}_0(\omega_x)|^2] \propto \left(1 - \text{sinc} \left(\frac{\tilde{\omega}_x}{(ma_0)^{-1}} \right) \right)^2. \quad (55)$$

First, we define

$$s_{\omega_x} = \sum_j e^{-2\pi i \omega_x x_j} = m \cdot \text{sinc} \left(\frac{\tilde{\omega}_x}{(ma_0)^{-1}} \right). \quad (56)$$

Next, we want to compute the expected phase at step j_2 conditioned on the fact that step $j_1 \neq j_2$ was set to $z_{j_1} = \zeta_j$. This implies that position j_2 can have any of the other $m - 1$ values in a uniform probability. Since $\sum_{j=1}^m e^{-ik2\zeta_j} = 0$,

$$E[e^{-ik2z_{j_2}} | z_{j_1} = \zeta_j] = -\frac{1}{m-1} e^{-ik2z_{j_1}}. \quad (57)$$

Using Eqs. (56) and (57), we are ready to compute Eq. (55)

$$\begin{aligned} & E[|\mathfrak{A}_0(\omega_x)|^2] \\ &= E \left[\left(\sum_{j_1} e^{-2\pi i(x_{j_1} \omega_x + \frac{2z_{j_1}}{\lambda})} \right) \left(\sum_{j_2} e^{2\pi i(x_{j_2} \omega_x + \frac{2z_{j_2}}{\lambda})} \right) \right] \\ &= \frac{1}{m} \sum_j E \left[\sum_{j_1} e^{-2\pi i(x_{j_1} \omega_x + \frac{2z_{j_1}}{\lambda})} \left(\sum_{j_2} e^{2\pi i(x_{j_2} \omega_x + \frac{2z_{j_2}}{\lambda})} \right) \Big| z_{j_1} = \zeta_j \right] \\ &= \frac{1}{m} \sum_j \sum_{j_1} e^{-2\pi i(x_{j_1} \omega_x + \frac{2\zeta_j}{\lambda})} \\ &\quad \cdot \left(e^{2\pi i(x_{j_1} \omega_x + \frac{2\zeta_j}{\lambda})} + \sum_{j_2 \neq j_1} e^{2\pi i x_{j_2} \omega_x} E \left[e^{ik(2z_{j_2})} \Big| z_{j_1} = \zeta_j \right] \right) \\ &= \frac{1}{m} \sum_j \sum_{j_1} e^{-2\pi i(x_{j_1} \omega_x + \frac{2\zeta_j}{\lambda})} \\ &\quad \cdot \left(e^{2\pi i(x_{j_1} \omega_x + \frac{2\zeta_j}{\lambda})} - \sum_{j_2 \neq j_1} e^{2\pi i z_{j_2} \omega_x} \frac{1}{m-1} e^{ik2\zeta_j} \right) \\ &= \frac{1}{m} \sum_j \sum_{j_1} e^{-2\pi i(x_{j_1} \omega_x + \frac{2\zeta_j}{\lambda})} e^{ik2\zeta_j} \\ &\quad \cdot \left(e^{2\pi i x_{j_1} \omega_x} - \frac{1}{m-1} \sum_{j_2 \neq j_1} e^{2\pi i x_{j_2} \omega_x} \right) \\ &= \frac{1}{m} \sum_j \sum_{j_1} e^{-2\pi i x_{j_1} \omega_x} \left(e^{2\pi i x_{j_1} \omega_x} - \frac{s_{\omega_x} - e^{2\pi i x_{j_1} \omega_x}}{m-1} \right) \\ &= \frac{1}{m} \sum_j \sum_{j_1} e^{-2\pi i x_{j_1} \omega_x} \left(e^{2\pi i x_{j_1} \omega_x} \left(1 + \frac{1}{m-1} \right) - \frac{s_{\omega_x}}{m-1} \right) \\ &= \frac{1}{m} \sum_j m \left(1 + \frac{1}{m-1} \right) - \sum_{j_1} e^{-2\pi i x_{j_1} \omega_x} \left(\frac{s_{\omega_x}}{m-1} \right) \\ &= \frac{1}{m} \sum_j \frac{m^2}{m-1} - \left(\frac{|s_{\omega_x}|^2}{m-1} \right) \\ &= \frac{m^2}{m-1} - \left(\frac{m^2 \text{sinc}^2 \left(\frac{\tilde{\omega}_x}{(ma_0)^{-1}} \right)}{m-1} \right) \\ &= \frac{m^2}{m-1} \left(1 - \text{sinc}^2 \left(\frac{\tilde{\omega}_x}{(ma_0)^{-1}} \right) \right). \end{aligned} \quad (58)$$

Eq. (58) provides Eq. (55), completing the proof □

