

MIT Open Access Articles

Dynamic Display of BRDFs

The MIT Faculty has made this article openly available. *Please share* how this access benefits you. Your story matters.

Citation: Hullin, Matthias B., Hendrik P. A. Lensch, Ramesh Raskar, Hans-Peter Seidel, and Ivo Ihrke. "Dynamic Display of BRDFs." *Computer Graphics Forum* 30, no. 2 (April 28, 2011): 475-483.

As Published: <http://dx.doi.org/10.1111/j.1467-8659.2011.01891.x>

Publisher: Wiley Blackwell

Persistent URL: <http://hdl.handle.net/1721.1/80346>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike 3.0



Dynamic Display of BRDFs

Matthias B. Hullin^{†1}, Hendrik P. A. Lensch², Ramesh Raskar³, Hans-Peter Seidel^{1,4} and Ivo Ihrke^{4,1}

¹MPI Informatik ²Universität Ulm ³MIT Media Lab ⁴Universität des Saarlandes

Abstract

This paper deals with the challenge of physically displaying reflectance, i.e., the appearance of a surface and its variation with the observer position and the illuminating environment. This is commonly described by the bidirectional reflectance distribution function (BRDF). We provide a catalogue of criteria for the display of BRDFs, and sketch a few orthogonal approaches to solving the problem in an optically passive way. Our specific implementation is based on a liquid surface, on which we excite waves in order to achieve a varying degree of anisotropic roughness. The resulting probability density function of the surface normal is shown to follow a Gaussian distribution similar to most established BRDF models.

Categories and Subject Descriptors (according to ACM CCS): I.3.1: Three-dimensional displays, I.3.7: Virtual reality, I.3.7: Colour, shading, shadowing, and texture, I.3.3: Display algorithms

1. Introduction

We are living in a time where the boundaries between real and virtual worlds are gradually blurred out. As technology keeps on evolving, somewhere in the distant future the output of computer displays may become visually indistinguishable from the real world.

In this paper, we focus on one major clue that helps distinguish real from virtual objects: reflectance. If we move around an object, and its appearance, the highlights, etc. change consistently with what we are used to and in accordance with the surrounding world, this object is more likely to be perceived as “real”.

In order to achieve full realism for computer generated content, it therefore stands to reason that the display of the future will behave more like a showcase window through which the the real and virtual worlds can interact with each other, rather than being a strict output device. So far, all types of computer displays have shown pixels of different colours that should ideally be as invariant to the viewing and lighting conditions as possible. Our goal is to physically mimic the characteristic way different surfaces reflect light, often described in terms of the Bidirectional Reflectance Distribution Function (BRDF), and display *materials* instead of colours (Figure 1).

In the recent years, many methods have been developed to fabricate materials with custom reflectance and sub-surface




Figure 1: Left: Idea of a “reflectance display” that reacts to its environment like real-world materials do. Right: Reflection of a checkerboard pattern in a surface that can exhibit different degrees of anisotropic roughness.

scattering properties. Yet, to our knowledge these properties have not been displayed dynamically. In Section 2, we discuss the prior work in related fields. The contribution of this paper is a very first step towards the dynamic display of materials by means of a physical device that can be programmed to exhibit a range of reflectance distributions. Before we elaborate on our own approach, we lay the founda-

[†] hullin@mpi-inf.mpg.de

tions in Section 3 with a general definition of the problem, along with a few sketches to its solution.

Of major importance to the appearance of real-world materials is their microstructure, e.g. in the form of surface roughness. By shaping surfaces, specific reflectance distribution functions can be achieved, as in the case of polished or brushed steel surfaces, or sandblasted glass. Our approach to dynamically displaying reflectance is also based on roughness modulation: we start with a liquid surface and excite surface waves on it. The space of possible appearances is defined by the reflectance of the base material, and the achievable surface structures that are in turn governed by the physics of wave propagation on liquids. In Section 4, we provide the underlying theory, and make a few basic predictions that will later be checked in experiment.

In Section 5, we present two prototypes that modulate the angular variation of reflected light in an optically passive way. We show in Section 6 that our devices can produce a range of anisotropic BRDFs that match our theoretical expectations. Using implementations of the same principle on different scales, we show that miniaturisation is not just possible but also desirable. Naturally, as the first devices of their kind, our prototypes are limited in terms of expressivity and practical use. We discuss these limitations in Section 7, and provide directions for future improvement.

A full derivation of the probability density function of reflection directions from a sine-shaped height field is provided as an appendix.

2. Related Work

From the very early age of computer graphics research, it has been recognised that reflectance models are a crucial ingredient for realistic rendering. Torrance and Sparrow [TS67] had been the first to provide a shading model based on microfacet geometry, Phong [Pho75] and Blinn [Bli77] showed its first applications in computer graphics. As the technical possibilities grew and the demand for physical accuracy increased, researchers began fitting model parameters against sparse reflectance measurements [War92] and provided databases of measured BRDFs [CUR96, MPBM03]. Ngan et al. [NDM05] related many of the previously introduced analytical BRDF models to dense reflectance measurements taken from real materials.

We use the established Ward BRDF [War92] as a reference model for the reflectance distributions exhibited by our device.

Rendering with real-world environment lighting can yield a great degree of realism with moderate technical and artistic effort, as first demonstrated by Debevec [Deb98]. Raskar et al. [RWLB01] used computer-controlled lighting to make real objects look in a desired way. Their “shader lamps” require total darkness for best performance. On the observer

Display approach	Optics	Light	Geometry /Material	Viewpoint
Virtual reality	active	virtual	virtual dynamic	simulated
[JMY*07, IKS*10]	active	virtual	virtual dynamic	real
Augmented reality [RWLB01]	active	virtual	real static	real
[Deb98]	active	real	virtual dynamic	simulated
[CNR08, KN08, HLHR09]	active	real	virtual dynamic	real
[FRSL08] and fabrication	passive	real	real static	real
BRDF display (ours)	passive	real	real dynamic	real

Table 1: Light and observer-dependent rendering and display techniques at a glance.




Figure 2: Left: optically active setup with a camera, a processing stage (here: negation) and a projector. Right: optically passive setup consisting of an imaging lens and a diffrusor sheet to achieve a blurred image.

side, the virtual viewpoint can be controlled by the user through different means, e.g. through game input devices or head tracking as in many early Virtual Reality installations. In the recent years, autostereoscopic displays have been constructed (optionally combined with head tracking) to achieve freely viewable 3D imagery [JMY*07, IKS*10] and even more based on lenticular or parallax barrier principle [HL10, KHL10]. Various active devices have been demonstrated that combine light field sensing and/or display with intermediate processing [CNR08, KN08, HLHR09].

The last three years have spawned a large amount of work dedicated to fabricating materials and objects with custom properties. The efforts include milling of height fields to reproduce reflectance distributions [WPMR09], printing of spatially varying BRDFs on paper [MAG*09], fabrication of subsurface scattering materials [HFM*10, DWP*10], surface reliefs that show lighting-dependent images [AM10], and even custom deformability [BBO*10]. The reflectance field assembly by Fuchs et al. [FRSL08] stacked purely passive “pixels”, each encapsulating a full 4D transmittance field, albeit at very limited resolution.

As the key contribution of this work, we see the definition of the BRDF display as a dynamic alternative to fabrication. We outline the main characteristics that any BRDF display should possess, and demonstrate a design that meets all principal requirements. To our knowledge, our device is the first to be both optically passive and programmable (see Table 1).

3. Displaying Reflectance

A BRDF display is a surface that can be programmed to exhibit varying reflectance distributions. We propose the fol-

lowing (non-exhaustive) list of criteria to assess the capabilities of a given device:

Principal Criteria

- P1 **View and light dependence.** The reflectance of a BRDF display must vary with the viewing and lighting direction in a physically plausible manner.
- P2 **Immediate response** to illumination changes and observer position.
- P3 **High dynamic range,** or the capability to deal with a very wide range of incident intensities.
- P4 **Programmability.** The display has to be controlled by a computer in a nonpermanent way.
- P5 **Light efficiency and contrast.** Any setup will fail to convince if it has insufficient light throughput, especially when competing against undesired effects such as first-surface reflections.

BRDF Parameters/Properties

- B1 **Bell-shaped highlights.** For many glossy real-world materials (as well as most analytical BRDF models, for that matter), the highlights follow Gaussian or power-of-cosine distributions, or sums thereof.
- B2 **Anisotropy,** where the reflectance varies with the tangent orientation of the sample.
- B3 **Modulation of absolute reflectivity.**
- B4 **Colour modulation.**
- B5 **Multi-lobe BRDFs,** for instance the rather popular model of diffuse+glossy+specular.
- B6 **“Off-normal” highlights,** e.g., when the average microfacet orientation does not coincide with the macroscopic normal, e.g. scale/sawtooth structures, hair.
- B7 **Retroreflection,** as has been observed for diffuse surfaces [PWKJ07].

Higher-Level Texture Layout

- T1 **Spatial extent.** In order for an observer to appreciate the reflectance function, a display needs to have a certain minimum size to display the shape of highlights.
- T2 **Spatial variation,** where the display is formed by an array of individually controllable “texels”.
- T3 **Normal variation,** where the perception of reflectance is further supported by displaying non-flat surfaces.

3.1. Approaching the Problem in a Passive Way

While our task is rather well-defined—modulate the angular distribution of light reflected on a surface—it is hard to come up with “the” one ideal path to its solution. In fact, a host of different approaches can be imagined, each one with its own set of advantages and drawbacks.

The very first design decision is whether the display should be optically active or passive (Figure 2). Active systems are very flexible. They naturally meet Criterion P4 and, depending on the implementation, are potentially compatible with all of B1–T3. However, the performance of today’s

light field sensing and output devices is bound to collide with P2 and P3. Purely optical setups, on the other hand, offer the fastest possible response (all “processing” being done at the very speed of light) and virtually unlimited dynamic range. We argue that, since real-world reflectance is optically passive, it must be possible to mimic it by passive means. The following families of passive setups can be imagined:

Holography. In principle, holographic techniques could be applied to display different information, i.e. reflectances, for different *viewing* directions. These techniques are referred to as angle-multiplexed holograms, e.g. [Mok93]. Volume holographic storage devices [Orl00] are reacting to selective directional *illumination*. Holographic *wavelength* multiplexing [RLY92] could even store wavelength-dependent BRDFs [HHA*10]. However, to our best knowledge, the combination of these different holographic techniques to form a reflectance field display has not been demonstrated.

Integral Photography inspired approach. Using an optical multiplexer, the 4D space spanned by the incident and outgoing hemispheres is mapped to a plane, where the corresponding radiance values can be modulated by a 2D mask. [FRSL08] implemented a similar idea in transmission; a reflective counterpart could for instance be imagined using lenslet and mirrorlet arrays. While this approach allows for almost arbitrary modulation, its resolution is inherently limited. As [FRSL08] demonstrated, an impression of continuity (and hence, the *look* of reflectance) is hard to achieve using a lenslet array that shoots out rays from different locations, unless a massive array of identical pixels is being viewed from a large distance, or an elaborate system of diffractors is employed. More importantly, though, multiplexing approaches trade resolution for efficiency, and inherently lose a lot of light when sharp highlights need to be resolved. Considering the above, any implementation of the Integral Photography idea will have to struggle to meet Criteria P5, T1 and T2. Also, a dynamic but optically passive device has yet to be demonstrated.

Redistribution. Since real BRDFs span only a small subspace of all possible 4D distributions, general 4D modulation may not be necessary. A different, more natural approach to our problem is to redistribute the available light dynamically. Inspired by how real materials function on a microscopic level, a desired reflectance distribution can be obtained through a controlled scattering process [WPMR09]. There are numerous (static and dynamic) ways of achieving this, and the principle by itself does not collide with any of our criteria. What appears particularly interesting in this context is the idea of stacking functional layers to extend the space of achievable BRDFs. For instance, by mounting a switchable diffusor (LC-TEC FOS-25x30-PSCT, a cholesteric LC panel) on top of a mirror, we were able to switch between two different reflectance profiles: mirroring (with slight haze) and diffuse. We demonstrate another multi-layer application in Section 6.3.

Our own design, which will be covered in the rest of this paper, belongs to the *redistribution* family of BRDF displays. We use a liquid surface as the reflecting base geometry, and reshape it over space and time by inducing surface waves. The design allows for dynamic modulation of the angular spread of the reflected light over a limited range, but not its intensity. Our prototypes meet Criteria P1–P5, B1–B2 and T1, and further miniaturisation may allow for B5 and T2 as well.

4. Characterisation of Reflectance by Surface Waves

The operation principle of our device is to excite surface waves in a medium that supports relatively free travel of these waves. For our experiments, we use an interface between air and water.

We are *not* aiming at producing standing waves that would generate oscillating, yet stationary, microgeometry. Instead we rely on *time-averaging* of travelling waves. If the generated height field varies fast enough, this results in an impression of a static microfacet distribution at every surface point. Mathematically, this is akin to averaging over a static height field of infinite extent.

From the earliest days, microfacet-based models have been expressing reflectance as the product of a *probability density function* (PDF) of reflection directions, and additional material, geometry and normalisation terms (see e.g. Eq. 11 in [TS67]). As a physical surface and Fresnel reflector, our device naturally takes care of all of these, so we can focus on controlling the PDF.

In the following, we discuss the PDF generated by a single sine wave of small amplitude. For small angles ϕ , e.g.

$$\phi < 5^\circ \approx 0.0873 \text{ rad}, \quad (1)$$

the trigonometric functions can be approximated to an error of less than 0.4% by the first term of their Taylor series:

$$\sin(\phi) \approx \tan(\phi) \approx \phi_{(\text{rad})} \quad \text{and} \quad \cos(\phi) \approx 1. \quad (2)$$

Also, we can safely assume that interreflections are absent.

4.1. Single Sine Wave in One Dimension




Figure 3: On a half-wave height field $h(x) = a \sin(kx)$ (red), an incident ray (blue) with an angle β to the vertical (green) is reflected.

Assume a sine wave in x direction as depicted in Figure 3:

$$h(x) = a \sin(kx), \quad \text{where } k = \frac{2\pi}{\lambda}, \quad (3)$$

a being the amplitude, k the wave number and λ the wavelength of the excited wave. The angle $\alpha(x)$ of the surface normal at position x , measured in a mathematically positive sense with respect to the vertical axis, is related to the slope of the function as follows:

$$\tan \alpha(x) = h'(x) = \frac{dh}{dx} = ak \cos(kx). \quad (4)$$

Eq. 1 is met if the roughness $ak =: \rho$ is small. Then, a light ray incident at $(x, h(x))$ under a fixed angle β to the vertical will be mirrored into the reflection angle δ :

$$\delta(x) = 2\alpha(x) - \beta \stackrel{(2)(4)}{\approx} 2\rho \cos(kx) - \beta \quad (5)$$

The PDF $f_{\Delta\beta}(\delta)$, i.e., the likelihood of a ray in an *ensemble* of rays incident under the angle β to be reflected into the angle δ , can now be imagined as the limit case of a value histogram of $\delta(x)$: the probability for δ to lie within an infinitesimal interval $d\delta$ is the combined measure of all intervals dx which are mapped to $d\delta$.

For symmetry reasons, it is sufficient to look at the first half-wave of our height field ($x \in [0, \pi/k]$), where $\delta(x)$ is monotonically decreasing and therefore bijective, so that $f_{\Delta\beta}(\delta)$ can be obtained as the derivative of the inverse of $\delta(x)$, normalised by the measure of the interval:

$$\begin{aligned} f_{\Delta\beta}(\delta) &= \frac{k}{\pi} \left| \frac{dx(\delta)}{d\delta} \right| \approx \frac{1}{\pi} \left| \frac{d}{d\delta} \cos^{-1} \left(\frac{\delta + \beta}{2\rho} \right) \right| \\ &= \frac{1}{\pi \sqrt{4\rho^2 - (\delta + \beta)^2}} \quad (6) \end{aligned}$$

A plot of this function can be seen in Figure 4 (blue curve). Without the assumption of Eqs. 2, the PDF turns out slightly bulkier. Please refer to the appendix for a full derivation.

4.2. Multiple Sine Waves in One or Two Dimensions




Figure 4: PDFs for sum-of-sinusoid functions in the linear limit (small amplitude, normal incidence, $\rho = 0.5$). As more sinusoidal terms are added up, the distribution converges to a Gaussian profile.

When we shoot an ensemble of rays at our height field, the exact shape of the surface is not of importance. In fact, we can treat the orientation of microfacets as a random variable that follows a probability distribution of Eq. 6. As n sinusoidal terms are superimposed in our device, this corresponds to an addition of random variables $\hat{\Delta}_\beta^1 + \dots + \hat{\Delta}_\beta^n$. If we can ensure that the $\hat{\Delta}_\beta^i$ are *independent and identically distributed* (iid), the central limit theorem states that the PDF of their sum is the n -fold convolution of the individual PDFs, and that it approaches a Gaussian distribution for a large number of terms [GS01]. In practice, we observe a satisfactory bell shape already for 5 superimposed sinusoidal waves, see Fig. 4.

Note that the distribution $f_{\hat{\Delta}_\beta}$ only depends on ρ , i.e. the *product* of amplitude and wave number. Since the wavenumber is directly related to the excitation frequency through the dispersion relation of the medium, we can generate identical distributions (same ρ) using different combinations of wave number and amplitude.

Independence, in our setting, translates to a vertical motion of every point on the dynamic height field surface that is as non-repetitive as possible. We choose excitation frequencies that relate like large prime numbers to approximate this.

The variance of a superposition of n identical distributions is related to the roughness as $\sigma_n^2 = n\sigma_0^2 = n \cdot 2\rho^2$, where $\sigma_0^2 = 2\rho^2$ is the variance of the single-sine distribution (Eq. 6). Within the linear approximation, we can thus generate Gaussian reflectance profiles of a desired variance by scaling all amplitudes of the sinusoidal terms uniformly.

In our device we are using orthogonally travelling planar waves in the x - and y -directions. At sufficiently small ρ values, the waves decouple and the two-dimensional PDF of the reflection directions simply becomes the product of two distributions $f_{\hat{\Delta}_\beta^x}$ and $f_{\hat{\Delta}_\beta^y}$, both of the form in Eq. 6.

4.3. Connection to Analytical BRDF Models

Since the variances on the x - and y -axes can be chosen independently, our display can reproduce BRDFs with elliptical highlights as popularised by Ward [War92]. Note that our variance σ_n^2 describes the reflection angle whereas Ward's $\alpha_{x,y}^2$ is related to the half-angle. Hence, the distributions are comparable when $\sigma_{x,y}^2 \approx 4\alpha_{x,y}^2$.

While the surface waves modulate the specular or glossy part of these models, the diffuse term can be realised by using a ground plane made of Labsphere Spectralon, an almost Lambertian reflector. The colour of the diffuse reflection can be influenced by a transmissive filter; we demonstrate this by dyeing the water prior to modulating the water surface.

In conclusion, our device is capable of displaying microfacet BRDF models with anisotropic Gaussian microfacet distribution. The parameters of the model can be directly related to the parameters that control our device.




Figure 5: Resonance profile of one of the wave generators in Setup 1, as observed through deflection of a laser beam. Trajectories were recorded in 1 Hz steps.

5. Construction of Devices

We built two incarnations of the same principle on different scales. Both setups consist of a flat water surface on which a pair of actuators excites crossing planar waves. We use voice coils that are fed with an amplified audio signal from the computer. Our signal source is the free software Puredata [Puc] running a patch that synthesises a stereo signal from sine wave terms of different frequency and amplitude.

Setup 1 (large, slow) consists of the ripple tank system WA-9897 by Pasco, Inc., a device designed for demonstration experiments in physics classes. We use a pair of ripple generators (each with a bar-shaped lever and modified to accept audio input) on a flat water tank with wedge-shaped soft foam beaches at the borders to suppress reflections, and a surface of approximately $23 \times 23 \text{ cm}^2$.

Setup 2 (small, fast) is a downscaled version built from a pair of 2.5-inch hard disk drives (Figure 6). The platters and the controller boards as well as part of the aluminium frames were removed, leaving only the arm assemblies in place. To each arm we attached a small bar-shaped piece of plastic to dip into the water, and mounted both frames crosswise. A small water receptacle (approx. $2 \times 2 \text{ cm}^2$ in size) is placed underneath the actuators.

6. Results

6.1. Characterisation

By deflecting a laser beam, we characterised the achieved surface normal variation for the “slow” Setup 1. The response of the actuator and its coupling to the water surface varies with the frequency of the signal. In Figure 5, we see that the efficiency peak is located around 20 Hz when the system is in contact with water. The resonance frequency of the uncoupled actuators is approximately 40 Hz. In Figure 7, we investigate the 1-dimensional case as described in Section 4.1. The driving signal is a sum of sinusoids with different amplitudes and frequencies. The amplitudes of each term were individually adjusted to yield a constant roughness parameter ρ on the water surface by adjusting the peak deflection angle of a laser beam. The resulting brightness profiles agree surprisingly well with the theoretical predictions. Due to its smaller size, Setup 2 responds considerably faster than Setup 1. We managed to excite water waves at frequencies as high as 800 Hz, although viscous damping limits




Figure 6: Various views of Setup 2 built from a pair of discarded 2.5-inch hard disk drives. From left to right: Components, top view, close-up on actuator 2, water surface and checkerboard target.

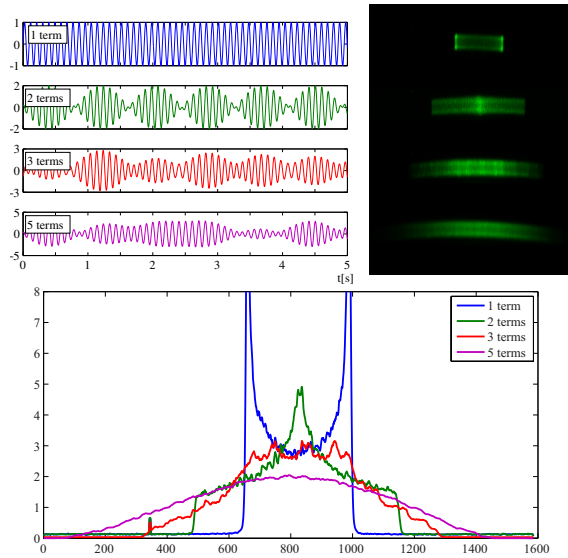


Figure 7: We verified the insights from Section 4.1 by superimposing sinusoids on the water surface of Setup 1 and observing the deflection of a laser beam. Top, left: idealised temporal profile of water wave; right: photos of laser beam. Bottom: intensity profiles. Note the similarity to Figure 4.

the reach of such high-frequent waves to a few millimetres or centimetres. We found the setup to be most efficient for frequencies around 200Hz, with deflection angles of up to 30° for single sine waves, or $\rho \approx 0.13$. If distortions of the trajectory can be tolerated, deflection angles of 50° ($\rho \approx 0.20$) are achievable. Above that, the actuators become unstable but droplet formation is not observed even for much higher amplitudes. As we relate our deflection angle to Ward’s BRDF model through the variance (Section 4.2), we obtain a range of $0 \leq \alpha_{x,y} < 0.14$, which overlaps with the values measured by Ward ($0.04 \leq \alpha_{x,y} \leq 0.26$).

6.2. Reflectance

In Figure 8, we compare the reflectance of our Setup 1 reflecting a checkerboard pattern with a raytraced simula-




Figure 8: Photos of Setup 1 (left, checkerboard pitch 10 mm) and raytraced simulations (right) of various degrees of anisotropic blur.




Figure 9: Real-time blur as displayed by Setup 2. The checkerboard scale is 1 mm.

tion of a comparable setting. The photos are long-exposure shots, which, given the slow response of the device, were required in order to achieve satisfactory temporal averaging. The anisotropic blur of the reflection is very similar in nature to the simulated result.

The smaller Setup 2 allows for comparable results at a much faster speed, enabling the observer (or a video camera) to directly perceive the angular blur. Figure 9 shows four representative frames from the accompanying video where the amplitudes in X and Y direction are manually adjusted in real time.

Setup 2 shows that miniaturisation brings a lot of benefits, since the achievable frequencies are approximately reciprocal to the geometric scale. Scaling down the setup by another order of magnitude may for instance enable temporal multiplexing of different lobes.

6.3. Diffuse + Specular

Also using Setup 2, we placed a diffuse reflector underneath the water surface and dyed the liquid. As can be seen in Figure 10, this adds greatly to the range of displayable BRDFs, however in our case the shallow water layer leads to an increase of viscous damping.




Figure 10: Macrophotos of Setup 2 (checker size 1 mm). By adding a diffuse white substrate and injecting coloured inks (here done manually), we obtain a combination of a dynamic anisotropic glossy and a static diffuse lobe.



Figure 11: Reflection of a human eye in Setup 2b filled with liquid metal. Left: resting. Right: in motion. The distortions are caused by surface tension.

6.4. Liquid Metal

In order to obtain a higher reflectivity especially for near-normal directions, we replaced the water in Setup 2 with an eutectic alloy of gallium, indium and tin. The nontoxic substance has a melting point of -18°C and is therefore liquid at room temperature. The increased viscosity, mass density, strong surface tension and incessant formation of an oxide layer make it difficult to control the surface shape. In particular, planar waves are rather hard to obtain even at very small amplitudes. However, we can still control a slight variation in the reflectance (Figure 11).

7. Discussion

The limitations of our device can be broadly classified into two categories: practical issues of our prototypical implementation, and fundamental limitations of the general design. In the following we discuss the major practical limitations of our prototype and suggest alternative ways of implementation.

1. Our display exhibits a limited range of surface roughness. With higher amplitudes, the nonlinearities in our physical system become hard to predict. By investigating the dominating effects and modelling them in the predictive model, it could be possible to achieve wider scattering profiles.
2. We are currently limited to BRDFs consisting of a single white Gaussian lobe in conjunction with a fixed coloured diffuse component. Further miniaturisation, e.g. using piezos as actuators, might enable temporal multiplexing of different lobes. The ink-based colouring layer currently used to model the diffuse component can be replaced by a passive transfective display technology.
3. The Fresnel reflection factor is currently the one of water. By adding refractive index altering agents to the medium in our device, a predictable change of reflectivity could be achieved that is closer to many real-world solid materials like plastics, resins or coatings.
4. Currently, we are only using two orthogonally crossing planar waves to generate BRDFs separable in the two dimensions. A more general implementation of the prin-

ciple could make use of multiple point-like actuators, generating spherical waves, implementing Huygens principle. This way, more general reflectance distributions could be generated, for instance a rotation of the tangent frame, which is currently missing in our design.

5. The use of liquids constrains us to horizontal mounting. Exchanging the water for solid jelly-like substances or elastic films might allow for arbitrary mounting angles. This may require the theoretical model to account for effects such as internal stresses and strains.

The only fundamental limitation we are aware of is the lack of a possibility to vary the normal of the simulated macro-surface. For this, we would require sawtooth-shaped waves that can hardly propagate on a strongly dispersive medium such as a solid or liquid surface. Therefore, Criteria B6 and T3 are most likely not achievable with the proposed system.

In conclusion, this paper has introduced the concept of a reflectance display. We have discussed theoretical requirements, advantages and drawbacks of different potential implementations. We have characterised a promising design both theoretically and practically by building two prototypes at different scales. An analytic link between Ward-like anisotropic BRDFs and the class of BRDFs displayable by our device has been established. In experiment, we have verified that dynamic time-averaged microfacet distributions can give the impression of dynamically changing, programmable reflectance in real time. Our prototype meets the principal criteria for a BRDF display and offers room for many extensions. We are confident that this presents a first step towards future hyper-realistic displays. However, much work remains to be done.

Acknowledgements

The authors would like to thank Michael Wand and Martin Bokeloh for their valuable comments. This work has been partially funded by the German Research Foundation (DFG), Cluster of Excellence “Multi-Modal Computing and Interaction” and Emmy Noether fellowship LE 1341/1-1.

References

- [AM10] ALEXA M., MATUSIK W.: Reliefs as images. *ACM Trans. Graph. (Proc. ACM SIGGRAPH)* (2010).
- [BBO*10] BICKEL B., BÄCHER M., OTADUY M. A., LEE H. R., PFISTER H., GROSS M., MATUSIK W.: Design and fabrication of materials with desired deformation behavior. *ACM Trans. Graph. (Proc. ACM SIGGRAPH)* (2010).
- [Bli77] BLINN J. F.: Models of light reflection for computer synthesized pictures. *Proc. ACM SIGGRAPH* (1977).
- [CNR08] COSSAIRT O., NAYAR S. K., RAMAMOORTHY R.: Light Field Transfer: Global Illumination Between Real and Synthetic Objects. *ACM Trans. Graph. (Proc. ACM SIGGRAPH)* (2008).
- [CUR96] CURET: Columbia Utrecht Texture Database. <http://www1.cs.columbia.edu/CAVE/software/curet/>, 1996.
- [Deb98] DEBEVEC P.: Rendering synthetic objects into real scenes. *Proc. ACM SIGGRAPH* (1998).
- [DWP*10] DONG Y., WANG J., PELLACINI F., TONG X., GUO B.: Fabricating spatially-varying subsurface scattering. *ACM Trans. Graph. (Proc. ACM SIGGRAPH)* (2010).
- [FRSL08] FUCHS M., RASKAR R., SEIDEL H.-P., LENSCH H. P. A.: Towards passive 6D reflectance field displays. *ACM Trans. Graph. (Proc. ACM SIGGRAPH)* (2008).
- [GS01] GRIMMETT G., STIRZAKER D.: "Probability and Random Processes". Oxford University Press, 2001.
- [HFM*10] HAŠAN M., FUCHS M., MATUSIK W., PFISTER H., RUSINKIEWICZ S.: Physical reproduction of materials with specified subsurface scattering. *ACM Trans. Graph. (Proc. ACM SIGGRAPH)* (2010).
- [HHA*10] HULLIN M. B., HANIKA J., AJDIN B., SEIDEL H.-P., KAUTZ J., LENSCH H. P. A.: Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions. *ACM Trans. Graph. (Proc. ACM SIGGRAPH)* (2010).
- [HL10] HIRSCH M., LANMAN D.: Build your own 3d display. In *ACM SIGGRAPH 2010 Courses* (2010).
- [HLHR09] HIRSCH M., LANMAN D., HOLTZMAN H., RASKAR R.: BiDi screen: A thin, depth-sensing LCD for 3D interaction using lights fields. *ACM Trans. Graph. (Proc. ACM SIGGRAPH Asia)* (2009).
- [IKS*10] ITO K., KIKUCHI H., SAKURAI H., KOBAYASHI I., YASUNAGA H., MORI H., TOKUYAMA K., ISHIKAWA H., HAYASAKA K., YANAGISAWA H.: Sony RayModeler: 360° autostereoscopic display. In *SIGGRAPH Emerging Technologies* (2010).
- [JMY*07] JONES A., MCDOWALL I., YAMADA H., BOLAS M., DEBEVEC P.: An interactive 360° light field display. In *SIGGRAPH Emerging Technologies* (2007).
- [KHL10] KIM Y., HONG K., LEE B.: Recent researches based on integral imaging display method. *3D Research 1* (2010), 17–27. 10.1007/3DRes.01(2010)2.
- [KN08] KOIKE T., NAEMURA T.: BRDF display: interactive view dependent texture display using integral photography. In *IPT/EDT '08: Proceedings of the 2008 Workshop on Immersive Projection Technologies/Emerging Display Technologies* (New York, NY, USA, 2008), ACM, pp. 1–4.
- [MAG*09] MATUSIK W., AJDIN B., GU J., LAWRENCE J., LENSCH H. P. A., PELLACINI F., RUSINKIEWICZ S.: Printing spatially-varying reflectance. *ACM Trans. Graph. (Proc. ACM SIGGRAPH Asia)* (2009).
- [Mok93] MOK F. H.: Angle-Multiplexed Storage of 5000 Holograms in Lithium Niobate. *Optics Letters* 18, 11 (1993), 915–917.
- [MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. *ACM Trans. Graph. (Proc. SIGGRAPH)* 22, 3 (2003), 759–769.
- [NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of BRDF models. In *Proc. of the Eurographics Symposium on Rendering* (2005), pp. 117–226.
- [Ori00] ORLOV P.: Volume Holographic Data Storage. *Communications of the ACM* 43, 11 (2000), 46–54.
- [Pho75] PHONG B. T.: Illumination for computer generated pictures. *Commun. ACM* 18, 6 (1975), 311–317.
- [Puc] PUCKETTE M. S.: PureData (Pd). <http://puredata.info>.
- [PWKJ07] PAPETTI T. J., WALKER W. E., KEFFER C. E., JOHNSON B. E.: Coherent backscatter: measurement of the

retroreflective BRDF peak exhibited by several surfaces relevant to lidar applications. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* (Sept. 2007), vol. 6682.

[RLY92] RAKULJIC G. A., LEYVA V., YARIV A.: Optical Data Storage by using Orthogonal Wavelength-Multiplexed Volume Holograms. *Optics Letters* 17, 20 (1992), 1471–1473.

[RWLB01] RASKAR R., WELCH G., LOW K.-L., BANDYOPADHYAY D.: Shader lamps: Animating real objects with image-based illumination. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques* (London, UK, 2001), Springer-Verlag, pp. 89–102.

[TS67] TORRANCE K. E., SPARROW E. M.: Theory for off-specular reflection from roughened surfaces. *J. Opt. Soc. Am.* 57, 9 (1967), 1105–1112.

[WAR92] WARD G. J.: Measuring and modeling anisotropic reflection. *Computer Graphics (Proc. SIGGRAPH)* 26, 2 (1992), 265–272.

[WPMR09] WEYRICH T., PEERS P., MATUSIK W., RUSINKIEWICZ S.: Fabricating microgeometry for custom surface reflectance. *ACM Trans. Graph. (Proc. ACM SIGGRAPH)* (2009).

A. PDF of a Sine-Shaped Height Field

Given a height field $h(x)$, the reflection direction δ of a light ray incident at $((x, h(x)))$ under a fixed angle β to the vertical is given as

$$\delta(x) \stackrel{(4)(5)}{=} 2 \tan^{-1} h'(x) - \beta. \quad (7)$$

Now, consider the wave to generate a microfacet distribution. In our experiments we achieve this situation by time-averaging over the waves traveling by. As seen from Eq. 7, the reflection directions directly depend on the derivative of the height field, $h'(x)$. Thus, knowledge of the distribution of the height field derivatives yields the probability density function of reflection directions caused by the height field.

This can be achieved by considering a random process which describes light rays as they hit the height field at random positions. The problem then becomes to compute the probability density of a function of a random variable. The initial distribution (of hit points) is assumed to be uniform over the domain of a half-wave (Eq. 3):

$$f_X(X) = \begin{cases} k/\pi & X \in [0, \pi/k] \\ 0 & \text{else} \end{cases}, \quad (8)$$

to which Eq. 7 is applied to obtain the directional distribution. The choice of a unit distribution is valid if the amplitude of the wave is small such that the microfacets can be considered as lying in the plane of the macro-surface, only exhibiting the orientations of the sine wave. From probability theory [GS01] the probability density of the random variable Y generated by applying a function g to another random variable X is

$$f_Y(y) = f_X(g^{-1}) \cdot \left| \frac{dg^{-1}}{dy}(y) \right|. \quad (9)$$

In our case this yields

$$f_\Delta(\delta) = \frac{k}{\pi} \left| \frac{d\{\delta(x)\}^{-1}(\delta)}{d\delta} \right|, \quad (10)$$

where $\{\delta(x)\}^{-1}(\delta)$ is the inverse of the direction function, Eq. 7, and f_Δ is the probability density of the reflection directions,

$$\{\delta(x)\}^{-1}(\delta) = \frac{\cos^{-1}\left(\frac{\tan((\delta+\beta)/2)}{ak}\right)}{k}, \quad (11)$$

and

$$\frac{d\{\delta(x)\}^{-1}(\delta)}{d\delta} = -\frac{1}{2} \frac{1 + \tan((\delta+\beta)/2)^2}{ak^2 \sqrt{1 - \frac{\tan((\delta+\beta)/2)^2}{a^2 k^2}}}. \quad (12)$$

Thus, the probability density function for the reflection direction due to a sine wave is given by:

$$f_\Delta(\delta) = \frac{1}{2\pi} \frac{1 + \tan((\delta+\beta)/2)^2}{ak \sqrt{1 - \frac{\tan((\delta+\beta)/2)^2}{(ak)^2}}}, \quad (13)$$

where $|\cdot|$ has been dropped because $\frac{d\{\delta(x)\}^{-1}(\delta)}{d\delta}$ is strictly negative. Amplitude and wave number enter the equation as a common factor $\rho = ak = 2\pi a/\lambda$. As a consequence, the BRDF is invariant to changes in a and k that leave ρ constant. Changing the incident angle β only results in a shift of the distribution and does not affect its shape.

A.1. Mean and Variance

In order to quantify the amount of angular blur caused by a given PDF, we compute its variance. Since there are many complicating factors in the non-linear regime, we constrain ourselves to small values of ρ . This allows us to use the angle-linearised version of the PDF, Eq. 6, for which we can derive the variance in closed form.

The mean is found at $-\beta$, as expected. For the variance σ_0^2 , we exploit the shift-invariance of the PDF and set $\beta = 0$:

$$\begin{aligned} \sigma_0^2 &= \int_{-2\rho}^{2\rho} \delta^2 \cdot f_{\Delta_0} d\delta \\ &= \int_{-2\rho}^{2\rho} \frac{\delta^2}{\pi \sqrt{4\rho^2 - \delta^2}} d\delta \\ &= \left[\frac{1}{2\pi} \cdot \left(4\rho^2 \tan^{-1} \frac{\delta}{\sqrt{4\rho^2 - \delta^2}} - \delta \sqrt{4\rho^2 - \delta^2} \right) \right]_{-2\rho}^{2\rho} \\ &= \left[\frac{1}{2\pi} \cdot 4\rho^2 \tan^{-1} \frac{\delta}{\sqrt{4\rho^2 - \delta^2}} \right]_{-2\rho}^{2\rho} \end{aligned} \quad (14)$$

where the integration bounds are due to the domain of the inverse tangent. To compute the integral we have to take limits from the left at the right boundary and from the right at the left boundary of the interval $[-2\rho \dots 2\rho]$.

$$\lim_{\delta \rightarrow \pm 2\rho} \frac{1}{2\pi} \cdot 4\rho^2 \tan^{-1} \left(\frac{\delta}{\sqrt{4\rho^2 - \delta^2}} \right) = \pm \rho^2 \quad (15)$$

The variance of the angle-linearised directional distribution is thus given by $\sigma_0^2 = 2\rho^2$.