Appears in the SIGGRAPH Asia 2011 Proceedings.

Single View Reflectance Capture using Multiplexed Scattering and Time-of-flight Imaging

Nikhil Naik [1] Shuang Zhao [2] Andreas Velten [1] Ramesh Raskar [1] Kavita Bala [2]

[1]MIT Media Lab [2]Cornell University

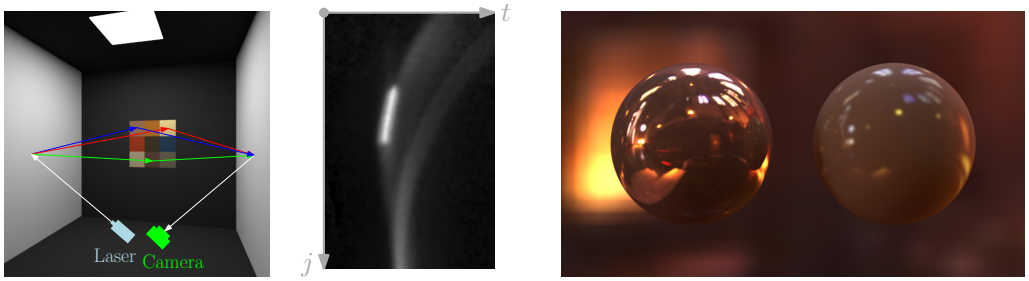

(a) (b) (c)

Figure 1: *By analyzing indirectly scattered light, our approach can recover reflectance from a single viewpoint of a camera and a laser projector. (a) To recover reflectance of patches on the back wall, we successively illuminate points on the left wall to indirectly create a range of incoming directions. The camera observes points on the right wall which in turn are illuminated by a range of outgoing directions on the patches. (b) Using space-time images ("streak images") from a time-of flight camera, we recover the reflectances of multiple patches simultaneously even though multiple pairs of patch and outgoing-direction contribute to the same point on the right wall. The figure shows a real streak image captured with two patches on the back wall. (c) We render two spheres with the simultaneously recovered parametric reflectance models of copper (left) and jasper (right) in simulation.*

Abstract

This paper introduces the concept of time-of-flight reflectance estimation, and demonstrates a new technique that allows a camera to rapidly acquire reflectance properties of objects from a single viewpoint, over relatively long distances and without encircling equipment. We measure material properties by indirectly illuminating an object by a laser source, and observing its reflected light indirectly using a time-of-flight camera. The configuration collectively acquires dense angular, but low spatial sampling, within a limited solid angle range - all from a single viewpoint. Our ultra-fast imaging approach captures space-time "streak images" that can separate out different bounces of light based on path length. Entanglements arise in the streak images mixing signals from multiple paths if they have the same total path length. We show how reflectances can be recovered by solving for a linear system of equations and assuming parametric material models; fitting to lower dimensional reflectance models enables us to disentangle measurements.

We demonstrate proof-of-concept results of parametric reflectance models for homogeneous and discretized heterogeneous patches, both using simulation and experimental hardware. As compared to lengthy or highly calibrated BRDF acquisition techniques, we demonstrate a device that can rapidly, on the order of seconds, capture meaningful reflectance information. We expect hardware advances to improve the portability and speed of this device.

Keywords: computational photography, multipath light transport, reflectance acquisition, global illumination, time of flight

Links: DL PDF

1 Introduction

Acquiring material properties of real-world materials has a long and rich history in computer graphics; existing techniques directly image the sample being measured to acquire different properties including tabulated reflectance functions, spatially varying reflectances, and parametric models (see [Weyrich et al. 2009] for a survey of state-of-the-art techniques.). These reflectance functions, are necessary for relighting, material editing, and rendering, as well as for matching and material identification.

In this paper, we present a new acquisition approach to reflectance measurement. Our approach is unique in two ways: we exploit ultra-fast time-of-flight (ToF) imaging to achieve rapid acquisition of materials; and we use indirect observation to acquire many samples simultaneously, and in fact, even permit around-the-corner measurement of reflectance properties. The key insight of this research is to exploit ultra-fast imaging to measure individual light transport paths, based on the distance traveled at the speed of light. This capability uniquely lets us separately measure the direct (0-bounce), 1-bounce, 2-bounce, and more, light paths; in comparison, traditional approaches use controlled laboratory settings to minimize the impact of multi-bounce light transport, or must explicitly separate direct and indirect lighting from all bounces.

We make the following contributions:

a) We present a new technique for reflectance acquisition by separating light multiplexed along different transport paths. Our approach uses indirect viewing with 3-bounce scattering coupled with time-of-flight imaging to capture reflectances. Our proof-of-concept system demonstrates first steps towards rapid material acquisition.

b) We exploit space-time images captured by a time-of-flight camera, that image different light transport paths over time. The inherent challenge is to decode material measurements in the presence of mixing over angular and spatial dimensions (we call this "entanglement"), where light arrives along multiple paths at the same point at the same time.

c) When there is no entanglement, it is possible to directly pick out the specular peak from streak images, enabling easy measurement of a material's gloss parameters. In the presence of entanglement, we show how fitting parametric models to our data successfully disentangles measurements.

d) Using time-of-flight (ToF) principles and indirect measurement, ours is the first solution to rapidly and remotely recover reflectance, even when the surface is not directly visible (around-the-corner). To the best of our knowledge this is the first approach to use light transport measurements in this manner. Our approach allows for remote reflectance capture without having to be close, or in contact, with the target material. No special instrumentation of the target material or scene is required.

There are several limitations to the current work. The acquisition is limited to a subset of the 4D space of a reflectance function expressed via lower-order parametric fits. The signal to noise ratio is not very high because the energy received after three bounces can be low. There are limits to the spatial and temporal resolution of the device, thus limiting the size of patches and the maximum sharpness of the reflectance function. For simplicity, we have assumed that the source and receiver surfaces (left and right walls) have a known diffuse reflectance; though this is not a fundamental limitation. Currently we use only a single wavelength laser, due to cost reasons, but getting lasers with different wavelength would permit spectral measurements.

Acquiring material properties enables a range of applications including image relighting, material editing, and material classification. We expect that improvements in laser technologies and time-resolved imaging will improve the stability, efficiency, and portability of this approach. We believe this design could in the future enable rapid, *"in-the-wild"* measurement of real-world scenes, without instrumentation.

2 Related Work

Nicodemus et al. [1977] introduced the Bidirectional Reflectance Distribution Function (BRDF), which characterizes light reflection from surfaces. BRDF acquisition has received much attention with efforts to acquire isotropic and anisotropic BRDFs, and spatially varying BRDFs and BTFs (bidirectional texture functions)[Ward 1992; Marschner et al. 1999; Dana et al. 1999; Lensch et al. 2001; Lensch et al. 2003; Matusik et al. 2003; Ngan et al. 2005; Lawrence et al. 2006; Ghosh et al. 2009; Ghosh et al. 2010a; Dong et al. 2010; Wang et al. 2009]. These acquisition processes are often lengthy, requiring extensive laboratory settings with calibration, and hours or even days of acquisition. Weyrich et al. [2009] present a detailed survey of the state of the art, and introduce a taxonomy of six major designs for material acquisition.

Capturing BRDFs requires measuring a large number of lighting–viewing combinations for the material. To decrease acquisition costs, many techniques focus on matching to parametric models (e.g., [Sato et al. 1997; Yu et al. 1999; Lensch et al. 2001; McAllister 2002; Gardner et al. 2003; Lensch et al. 2003; Goldman et al. 2010]). Various approaches decrease acquisition costs and increase coverage by exploiting properties of BRDFs including reciprocity, separability, spatial smoothness, and compressibility [Zickler et al. 2005; Sen et al. 2005; Garg et al. 2006; Wang et al. 2008; Dong et al. 2010]. Lighting configurations and variations have been considered including polarization and structured illumination[Wenger et al. 2005; Sen et al. 2005; Ma et al. 2007; Ghosh et al. 2010b].

Most BRDF techniques directly, or indirectly through a mirror surface [Weyrich et al. 2009], view and image the material sample with some notable exceptions [Han and Perlin 2003; Hawkins et al. 2005; Kuthirummal and Nayar 2006]. Hawkins et al. [2005] developed the dual light stage, and imaged a diffuse environment to increase angular measurements.

There has been a long line of research in computer vision in modeling interreflections to accurately model shape and reflectance [Forsyth and Zisserman 1990; Nayar et al. 1991]. Recently there has been interest in recovering scene properties such as geometry and albedo from multiple bounces of light [Nayar et al. 2006; Seitz et al. 2005; Sen et al. 2005; Kirmani et al. 2009; Liu et al. 2010; Bai et al. 2010], and also on recovering shape and material simultaneously [Holroyd et al. 2008].

In contrast to prior techniques we do not directly image the surface sample, but rather, indirectly image a diffuse surface and use time-of-flight principles to detect and measure all the bounces of light arriving at the diffuse surface, after interacting with the samples we want to measure. This approach enables rapid (on the order of seconds) acquisition of reflectances of multiple patches (tens of materials) simultaneously, over a range of angular measurements.

3 Time-of-flight Reflectance Acquisition

We now describe our time-of-flight acquisition system. We first describe the geometry of our setup. We derive the terms for acquiring a single patch, introduce streak images, and the problem of entanglements. We show how to generalize our acquisition to multiple patches, and discuss the coverage of our acquisition device.

3.1 Geometry of acquisition

We describe the canonical geometric setup to acquire reflectances in Figure 2-(a). The source S and receiver R are both assumed to be known Lambertian materials, and P is the patch being measured. In our equations, $\mathbf{s}$, $\mathbf{r}$, $\mathbf{p}$ indicate points on S, R, P, respectively. In addition, the laser illuminates the source S, and a camera views the surface R. Thus, we do not image P directly, but rather measure it indirectly.

Around the corner viewing: In the second configuration shown in Figure 2-(b) the patch being measured P is not directly visible to the camera. In this case, the source and the receiver are the same real-world surface (and of course, have the same reflectance properties). The laser shines on a part of the surface that is not being imaged by the camera (to avoid dynamic range issues). Depending on the orientation of the patch P, the angular coverage of directions is quite similar to the configuration in Figure 2-(a).

Mathematically, this configuration works in exactly the same way as the previous configuration, with the appropriate setting of angles and distances. In the following text, we illustrate our ideas using Figure 2-(a), but our physical experimental setup more closely matches Figure 2-(b).

3.2 Imaging a single patch

For each $\mathbf{s}$, light is reflected from P to R. A camera captures the light reflected along the path: $\mathbf{s} \rightarrow \mathbf{p} \rightarrow \mathbf{r}$. As shown in Figure 2-(a), given a point $\mathbf{p}$ with normal $\mathbf{N}_p$, θ_{ps} is the angle made by $\mathbf{s}$ at

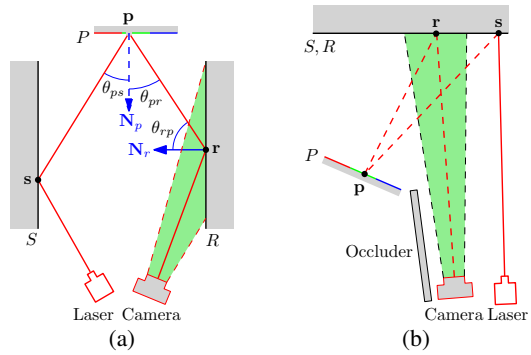

Figure 2: *We explore two acquisition setups, left: canonical setup, right: around-the-corner setup. To measure the patch P, the laser illuminates S, and the camera images R.*

$\mathbf{p}$ with respect to $\mathbf{N}_p$, namely $\cos\theta_{ps} = \mathbf{N}_p \cdot \frac{(\mathbf{s}-\mathbf{p})}{||(\mathbf{s}-\mathbf{p})||}$. Similarly, we define θ_{pr} for each $\mathbf{r}$.

Ignoring time, a regular camera viewing R would measure:

$$\frac{I_0\,\rho_S\,\rho_R}{\pi} f_r(\mathbf{s},\mathbf{p},\mathbf{r}) \frac{\cos\theta_{ps}\cos\theta_{pr}\cos\theta_{rp}}{\|\mathbf{s}-\mathbf{p}\|^2\,\|\mathbf{r}-\mathbf{p}\|^2} \tag{1}$$

where $f_r(\mathbf{s},\mathbf{p},\mathbf{r})$ is the BRDF at $\mathbf{p}$ along path $\mathbf{s}\to\mathbf{p}\to\mathbf{r}$, I_0 is the laser intensity, and ρ_S and ρ_R are the diffuse reflectance of S and R respectively. We introduce the geometry term $g(\mathbf{s},\mathbf{p},\mathbf{r})$, which entirely depends on the geometry of the patches, and the known reflectances:

$$g(\mathbf{s},\mathbf{p},\mathbf{r}) = \frac{I_0\,\rho_S\,\rho_R}{\pi} \frac{\cos\theta_{ps}\cos\theta_{pr}\cos\theta_{rp}}{\|\mathbf{s}-\mathbf{p}\|^2\,\|\mathbf{r}-\mathbf{p}\|^2} \tag{2}$$

Given a patch P with finite extent, a regular camera imaging the patch at location $\mathbf{r}$ would capture:

$$\int_{\mathbf{p}\in P} g(\mathbf{s},\mathbf{p},\mathbf{r}) f_r(\mathbf{s},\mathbf{p},\mathbf{r})\,\mathrm{d}\mathbf{p} \tag{3}$$

Using a time-of-flight camera: The availability of an ultra-fast lighting and imaging system is one of the key differences of our approach with respect to traditional measurement approaches. A fast laser switches rapidly, enabling fast scanning of S. The time-of-flight camera measures light at picosecond time intervals; thus, light traversing different path lengths arrive at different instants of time that are measurable by the imaging device. The distance traveled in 1 picosecond is 300 micrometers (speed of light is $3\times10^8 m/second$).

Consider a point $\mathbf{r}$: as light from the laser follows different paths, reflecting off different points $\mathbf{p}$, it will arrive at different times at the receiver. Apart from the spatial image, our camera is able to capture images over time (Currently, our camera captures images with only 1 spatial dimension and the temporal dimension.). We call these images "streak images". See Figure 3 for an illustrative streak image for a small patch P. Note the shape of the curve, depending on the path lengths of light from S, reflecting off P, to R. Note that with the time-of-flight camera it is possible to observe the specular peak directly from the streak image.

In comparison with a regular camera image we note that if we add all the image data across time for a single point $\mathbf{r}$ and a single source location $\mathbf{s}$ we obtain the regular camera image value for that $(\mathbf{s},\mathbf{r})$ combination.

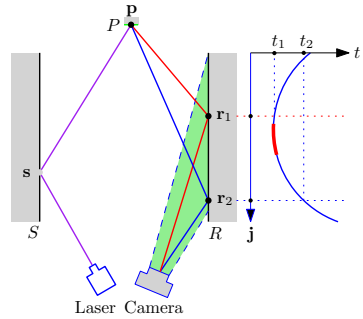

Figure 3: *This figure illustrates a simple setup (left) and the corresponding streak image (right). Red and blue colors on the curve indicate regions with higher and lower brightness, respectively. The bright region on the streak image contains information about the shape of the specular lobe of the BRDF of P.*

3.2.1 Streak image

We define the streak image to be a 2D image, Q. With an ultra-fast camera, we image R capturing the photon arrivals at R, and use this information to estimate the reflectances of P. Given the laser illumination s, the point $\mathbf{r}$ at a time t measures the following:

$$Q_s(\mathbf{r},t) = \int_{k\in K'} g(\mathbf{s},\mathbf{p},\mathbf{r}) f_r(\mathbf{s},\mathbf{p},\mathbf{r})\,\mathrm{d}\mathbf{p} \tag{4}$$

where, $K' \subseteq P$ consists of all points $\mathbf{p}$ such that the path length d for the laser light to arrive at $\mathbf{r}$ is the same. That is, at a given instant t light arrives at $\mathbf{r}$ from all points $\mathbf{p}$ which have an equal path length along $\mathbf{s}\to\mathbf{p}\to\mathbf{r}$:

$$d = \|\mathbf{s}-\mathbf{p}\| + \|\mathbf{r}-\mathbf{p}\| = c\cdot t$$

where c is the speed of light.

In the case where P is very small, the streak image has the shape shown in Figure 3. When P is a patch with finite extent, we get an integral of all the points on P (with the same total path length) at each time instant, thus giving a thicker curve in the streak image.

3.2.2 Path separation

The ultra-fast camera separates out light with different bounces. Thus at the receiver R, direct light from $\mathbf{s}\to\mathbf{r}$ arrives first, then light from $\mathbf{s}\to\mathbf{p}\to\mathbf{r}$, and so on, for greater number of bounces of light. Ultra-fast imaging permits us to easily separate out these different light transport paths, thus, greatly simplifying material acquisition by letting us separate out terms that include only the BRDF we are interested in.

3.2.3 Entanglement

One technical challenge with streak images, is that light paths are not always separated. In the case where two different light paths arrive at the same point $\mathbf{r}$ at the same time t (because they have the same path length), there is a linear mixing observed at that point in the streak image. Figure 4 shows the locus of points which get mixed at r: an ellipse in 2D, an ellipsoid in 3D. All the measurements from this locus of points are added together at the imaging point for any given time instant.

The around-the-corner setup, as seen in Figure 2-(b), brings the foci of the ellipse closer together, which increases the curvature of the ellipse around the sample location. This ensures fewer entanglements as compared to the canonical setup in Figure 2-(a).

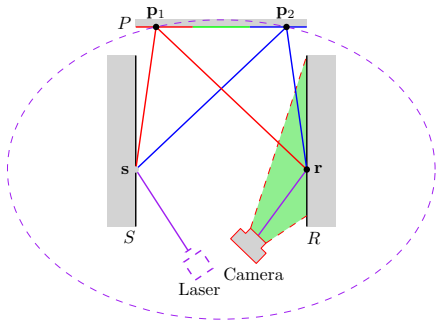

Figure 4: *Multiplexed scattering causes entanglements: Two light paths,* $\mathbf{s} \rightarrow \mathbf{p}_1 \rightarrow \mathbf{r}$ *and* $\mathbf{s} \rightarrow \mathbf{p}_2 \rightarrow \mathbf{r}$*, with the same path length, will get mixed in the streak image. The dashed ellipse (ellipsoid in 3D) shows all points with this path length.*

3.3 Multiple patch measurement

As a patch gets larger we can acquire more measurements, which improves the signal to noise ratio of our measurements. However, the larger the patch, the more entanglement occurs. Thus, there is a tradeoff in patch size between better signal and sharper signal (that preserves high frequency information). To avoid excessive blurring, we split the patch into multiple patches (of size $1cm \times 1cm$), and independently recover reflectance parameters for each subpatch. While entanglements are still possible, if the dimensions of the subpatch are not comparable with those of the source and receiver walls, there will be only a few such entanglements. In Section 4 we show how we can disentangle these observations by fitting to a parametric material model.

3.4 Coverage of our measurements

We indirectly observe BRDF measurements by imaging a Lambertian material R. This enables taking many BRDF measurements simultaneously, for a given point P, thus, accelerating acquisition. However, this approach is only practicable with the ultra-fast camera which is able to disambiguate between measurements across time, thus enabling accurate reconstruction. Note that given the geometry, there are restrictions on the range of incoming and outgoing directions that are measurable. As the laser sweeps over S, we measure a 2D set of values on R.

In the canonical setup shown in Figure 2-(a), the horizontal separation between the S and R surfaces, as well as the vertical length of the two surfaces decide the angular range of observations. In theory, if both S and R surfaces are infinitely long, and separated by an infinite distance, half of the hemisphere of all possible incoming directions as well as the other half of the hemisphere in the outgoing direction (mirror direction) will be sampled. However, due to the inverse-square fall-off of energy with path length, the practical setup cannot be arbitrarily large. Its dimensionality is limited by constraints on the intensity of the light source and the sensitivity of the camera imaging the receiver surface.

The around-the-corner setup in Figure 2-(b), can be thought of as the setup in Figure 2-(a) with the S and R surfaces folded together. So the same discussion of angular coverage applies to this case. The angular coverage of practical setups is discussed in Section 6.

4 Reconstructing Reflectance Values

In this section we describe how we recover reflectance values from the acquired data.

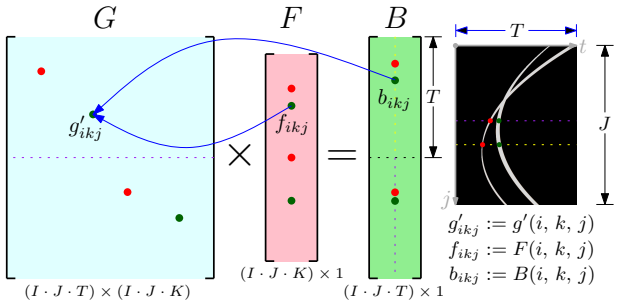

Figure 5: *The linear system to solve for reflectances is constructed using the streak images and scene geometry information. Consider two patches with a streak image for one source position* $\mathbf{s}_i$*. The streak image is sampled at two receiver positions* $\mathbf{r}_1$ *and* $\mathbf{r}_2$ *to create the observation vector* B*. Based on the time-slot an observation falls in, we calculate the corresponding "physical factor"* $g'(i, k, j)$*. The time-slot decides the row location of* g'_{ikj}*. The column location of* g'_{ikj} *is decided by the column location of the corresponding BRDF in* F*. The same process is repeated in the case of multiple source positions, by stacking up the streak-columns vertically, which in turn makes the* G *matrix taller.* B *and* G *are extremely sparse. In this figure, all the elements in* B *and* G*, except the colored dots, are zero.*

4.1 Discretizing the problem domain

We can estimate the geometry of the scene using our time-of-flight device. A basic method to achieve this is described in [Kirmani et al. 2009]. More recent algorithms are capable of reconstructing continuous surfaces and complex shapes without good knowledge of the object BRDF. Therefore, the various physical factors, θ_{ps}, θ_{pr}, θ_{rp}, $\|\mathbf{s} - \mathbf{p}\|$, and $\|\mathbf{r} - \mathbf{p}\|$, and d, can be determined. We discretize the problem as below: given I laser positions indexed by i, J receiver points indexed by j, K patch positions indexed by k, and T time slots (corresponding to the width of the streak images) indexed by m, we have $I \cdot J \cdot K$ unknown values of BRDF and $I \cdot J \cdot T$ measurements.

A single patch: Consider a small patch placed in the Figure 2-(a) configuration. We discretize the problem space for the patch, assuming that for any i, j and k, vectors $(\mathbf{s} - \mathbf{p})$ and $(\mathbf{r} - \mathbf{p})$, and the BRDF $f_r(\mathbf{s}, \mathbf{p}, \mathbf{r})$ are roughly constant. Discretizing Equation 4 we get:

$$Q_i(j, m) = A_k \, g(i, k, j) \, F(i, k, j) \tag{5}$$

where m is the discretized time taken by light traveling from the laser source to the camera along path $\mathbf{s} \rightarrow \mathbf{p} \rightarrow \mathbf{r}$, F is the discretized BRDF, A_k and $\mathbf{p}$ are the surface area and the center of the patch, respectively, and $g(i, k, j)$ is evaluated with all parameters at the center of the respective patches $\mathbf{s}_i$, $\mathbf{p}_k$ and $\mathbf{r}_j$. We introduce a modified geometry term $g'(i, k, j) = A_k \, g(i, k, j)$, where we fold the surface area A_k into the g term.

When there is no entanglement, $F(i, k, j)$ can be obtained by reading the intensity value $Q_i(j, m)$ on the streak image and dividing it by these known factors. With entanglement, the measurement in the streak image is a mixture of multiple paths, with different BRDF values along each path. Thus, it is not possible to directly invert these measurements. Therefore, we formulate a linear system of equations as below to solve for the unknown reflectance values.

Multiple patches: The case of K patches (with different BRDFs) is a simple generalization of the equations above. Again, when

there is no entanglement, each measurement corresponds to a single BRDF, but when there is entanglement there is a mixture of multiple (potentially different) BRDFs. Thus, the formulation for both single patch, and multipatch, is the same in the presence of entanglements.

4.2 Matrix formulation

We set up a linear system to solve for the BRDF for each patch. The angular coverage of the recovered BRDF depends on the geometry of the setup, primarily the dimensions of the source and receiver wall, as discussed in Section 3.4. The discretized system of equations is:

$$B(i,j,m) = G(i,k,j) \cdot F(i,k,j) + \nu \tag{6}$$

where, $B(i,j,m) := Q_i(j,m)$, and B, F, and G are the vectorized representations of observations, unknown BRDF values, and the physical factors respectively, and ν represents noise from the camera capture. The observation at each receiver position j represents one row $Q_i(j, 1:T)$. The observation vector B is constructed by stacking up the J columns.

The G matrix is populated with appropriate entries based on the scene geometry and the constraints due to entanglement (see Figure 5). Given a source position, and two receiver positions, we have all the time data for those two positions (shown as purple and yellow lines in the figure). The corresponding B vector consists of $2 \cdot T$ entries, the F vector consists of $2 \cdot K$ entries, and G is setup as described in the caption of the figure. Generalizing to I source positions, and J receiver positions, we arrive at the vector B of dimension $(I \cdot J \cdot T) \times 1$, matrix G of dimension $(I \cdot J \cdot T) \times (I \cdot J \cdot K)$, and vector F of dimension $(I \cdot J \cdot K) \times 1$. The matrix G as well as the observation vector B are extremely sparse. We will discuss the actual number of equations available next.

Rank: We assume that we have T time-slots of observations of the receiver R using the ultra-fast camera. Note that we assume that we limit the time range of observation to only include light that bounces from S to P to R (three bounces), eliminating lower and higher bounces of light. This is possible in our proposed setup using our time-of-flight imaging device.

Ideally, with no entanglements, if the path lengths of light arriving from each of the K patches is different, we will observe K separate responses at the receiver point R_j from $m = 1$ to T. The number of observations will be exactly equal to the number of unknowns, i.e. $I \cdot J \cdot K$, as there will be one unique observation corresponding to one triplet (i,j,k), and we can trivially invert the equation to acquire the BRDF value.

However any real-world geometry will contain a number of identical paths as shown in Figure 4. The light from different patches with identical pathlengths will add up in the corresponding bin $Q_i(j,m)$. Hence the number of observations corresponding to one laser and receiver position can be less than or equal to K. This makes the linear system underdetermined. Next, we describe how using reduced dimensional parametric models decreases the number of required measurements, thus, enabling recovery of reflectance parameters.

4.3 Parametric reflectance models

In order to solve the sparse underdetermined system defined earlier, we assume a low dimensional parametric model of the BRDF and recover the parameters of this BRDF. We use the half-angle parameterization proposed by Rusinkiewicz [1998], and use the dBRDF proposed in [Ashikhmin 2007], and used in [Ghosh et al. 2010b] to measure distributions of the BRDF. Ashikhmin et al. [2007] show that using such a fitting process for limited cone data can be effective. We compute the half angle vector $\mathbf{h}$ for each measurement and parameterize the BRDF as $f_r = k_d/\pi + k_s\, p(\mathbf{h})$ where the unknowns k_d, k_s are the diffuse and specular reflectance respectively, and $p(\mathbf{h})$ is a distribution parameterized by the half angle vector. Various distributions $p(\mathbf{h})$ have been published in graphics literature [Ngan et al. 2005; Matusik et al. 2003; Ashikmin et al. 2000]. Since our measurements have relatively limited cones of angles around the zero half-angle, we assume isotropic BRDFs and fit the following Ashikhmin-Shirley model described in Ngan et al. [2005]:

$$f_r = \frac{k_d}{\pi} + k_s \frac{n+1}{8\pi} \frac{(N \cdot H)^n}{(V \cdot H)\max((N \cdot L),\,(N \cdot V))}$$

We ignore the Fresnel term in our fit, which is reasonable given our configuration and range of angles covered. Thus, our BRDF estimation problem reduces to estimating 3 unknowns per patch i.e. k_d, k_s and n. Thus, the total number of unknowns for K patches reduce from $I \cdot J \cdot K$ to $3K$.

4.4 Solving for reflectances

When entanglement does not occur, for many half angle values it is possible to use streak images to directly measure out the BRDF without having to resort to fitting to a parametric representation. However, in the presence of entanglements, assuming a low dimensional parametric model, we have large number of observations and only a few unknowns per patch.

To solve the linear system of equations, we sort the columns of the matrix G by half angle, $(N \cdot H)$, values for each patch in ascending order. This ensures that the BRDF segment corresponding to each patch is a continuous segment in F. This helps to make the optimization process easier as the BRDF of each patch is now a single segment in vector F.

The observation vector B and the matrix G are very sparse as the actual number of non-zero observations in B are less than or equal to $I \cdot J \cdot K$. We use this fact to delete all zeros from B and the corresponding all-zero rows from G. This process reduces the size of the linear system considerably. The size of G reduces from $(I \cdot J \cdot T) \times (I \cdot J \cdot K)$ to $V \times (I \cdot J \cdot K)$, where V is dependent on the geometry of the setup. In our experiments, V is much less than $I \cdot J \cdot T$, and is of the order of $T \sim 500$.

To solve for the BRDF parameters, we apply unconstrained nonlinear optimization using the $fminunc$ function from the Matlab Optimization Toolbox. The optimization procedure uses the BFGS Quasi-Newton method with a mixed quadratic and cubic line search procedure. A detailed description of the algorithm can be found in the Matlab Optimization Toolbox User's guide [Mathworks 2011]. The optimization is performed over the $3K$ parameters to minimize the error metric using the L^2norm of $B - GF'$, where F' is the BRDF vector calculated using the estimated parameters.

We start with intuitive initial guesses of the parameters (n in the value of hundreds, and k_d and k_s at approximately 10 %) to the optimization process. We find that our optimization converges in a few iterations to the global minimum. Figure 8 shows the parametric fits obtained for multiple patches.

5 Experiments

Modifying existing time-of-flight (ToF) cameras to report a full time profile with sufficient SNR for our purpose is quite challenging. Our emphasis in this paper is on the computational and al-

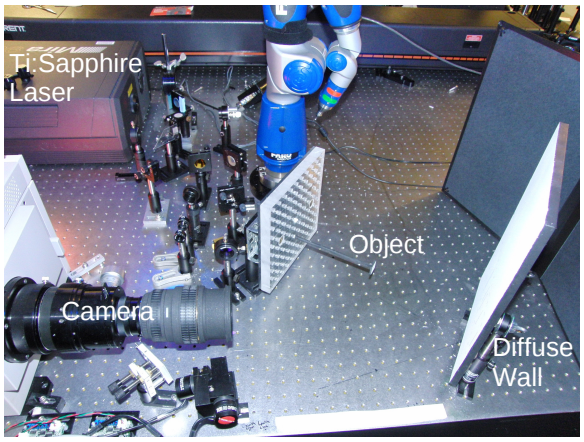

Figure 6: *The time-of-flight camera setup involves a laser projector via a Ti:Sapphire laser with steering mirrors and a pico-second accurate camera. This setup corresponds to Figure 2-(b). Reflectance of the "object" is recovered by aiming the camera and the laser projector at the diffuse wall.*

gorithmic aspects. Following experimental setup was used for our proof-of-concept implementation.

Our light source is a mode-locked 795 nm Ti:Sapphire laser. The camera is a Hamamatsu C5680 streak camera that captures a series of 512 one-dimensional images, i.e., lines of 672 pixels in the scene, with a time resolution of about 2 picoseconds. Figure 6 shows our setup in the configuration corresponding to Figure 2-(b).

6 Results

We now present the results of our time-of-flight reflectance measurement system. We first evaluate our results using a simulation of the device, to focus on demonstrating its capabilities with a good hardware implementation, and to validate against ground truth values. Finally we demonstrate acquisition using our hardware setup.

6.1 Simulation results and validation

In our simulation setup we simulated the working of a time-of-flight camera in the indirect viewing configurations of Figures 2-(a) and (b). The results are the same in both configurations, so we only present our results for the (b) configuration.

6.1.1 Simulation setup

We simulate our physical setup from Figure 6. S, P, and R form the left, back, and right planes of a box. S and R are of size $25 \times 25cm^2$, and are separated by a distance 15 cm. We consider $I = 15$ laser positions on S, $J = 25$ receiver positions on R. We specify the number of patches for each experiment as they arise. We measure $T = 512$ time slots, where each time slot corresponds to 1.6678 picoseconds. Both the source and receiver are assumed to be Lambertian surfaces with known reflectance, $k_d = 0.97$.

We simulate different measurement scenarios by introducing single and multiple patches of different materials assuming a parametric Ashikhmin BRDF model (see Section 4.3). Using the geometry information and the supplied BRDFs, we generate the streak images. To simulate finite patch areas, we sample each patch at ten random points on the surface and average them at different receiver points.

Our simulation provides us with the measurement vector B. We further add random noise to these observations to simulate the noise introduced by the camera and other external factors during the actual capture process. For this purpose, we employ a commonly applied noise model [Hasinoff et al. 2010; Schechner et al. 2007] consisting of a signal independent additive term, which includes dark current and amplifier noise, plus a signal dependent photon shot noise term.

The noise added to an observation $B(i, j, m)$ is given by $\nu_b = \nu_{\text{floor}} + \nu_{\text{photon}}$. Here ν_{floor} is a constant noise-floor given by $\nu_{\text{floor}} = 0.01 \cdot \max(B) \cdot \xi$ where B is the observation vector, and ξ is a random number in $\aleph(0; 1)$. ν_{photon} is the photon noise give by $\nu_{\text{photon}} = \eta \cdot B(i, j, m) \cdot \xi$, where η is a noise-parameter defined as a percentage.

6.1.2 Results

We simulate two scenarios using this setup: single patch and multiple patch. We choose nine materials from the database by Ngan et al [2005] to represent a wide variety of reflectances. We simulate the three different channels separately using the corresponding parameters. The first scenario contains a single unit-area ($1cm^2$) patch placed on $\mathbf{P}$. The second scenario simulates a two-dimensional grid of nine patches arranged in a 3×3 configuration.

Single patch simulations: We recover parametric BRDFs, as described in Section 4.3, for three different materials with different percentage of additive noise: $\eta = 0\%$, 1% and 10%. Table 1 shows the ground truth results, and the recovered parameter values for two materials (copper, and red plastic specular). For $\eta = 1\%$ the recovered parameters have very low error, and even for $\eta = 10\%$, the recovered parameter values are reasonable. The diffuse reflectances are most affected by noise, since they are relatively low. The supplementary material provides more comparisons.

Multi patch simulations: We use the same nine materials, now arranged in a 3×3 2-D grid of patches of unit area to create a 'multipatch' at $\mathbf{P}$ (see Figure 1). Again for this case, we recover the parameters for three different percentages of additive noise: $\eta = 0\%$, 1% and 10%. Figure 8 shows rendered spheres using the ground truth, and the recovered BRDFs with $\eta = 1\%$, and $\eta = 10\%$. Figure 8-(left) shows the streak image corresponding to the 9 materials. We can see that there is mixing in the streak image that we are able to separate out robustly using our reconstruction algorithm. See the supplementary material for more comparisons.

Angular coverage of simulation setup: The angular coverage in the simulated and physical setups is between 0° to 20° in terms of half-angles. Both the incoming and outgoing angles are in the range of 25° to 65° with respect to the patch normal. The setup dimensions are selected such that grazing angles are avoided.

6.2 Results using our experimental device

We now evaluate our prototype experimental device for the "around the corner" viewing mode using a high speed time-of-flight camera, and a picosecond accurate laser. The setup as shown in Figure 2-(b) is described in detail in Section 5.

Single patch data: For a single material, we image a small material patch (of size $1.5 \times 1.5\,cm^2$) using our acquisition setup. Figure 9-(1a) and (1b) show the streak images for two very different measured materials: copper and red plastic. 1(c) shows the rendered spheres using an environment map with the recovered parameters for the two materials, with the error plots (1d) and (1e). The results are taken at the wavelength of our laser at a wavelength band from 770 nm to 820 nm and centered at about 795 nm. In

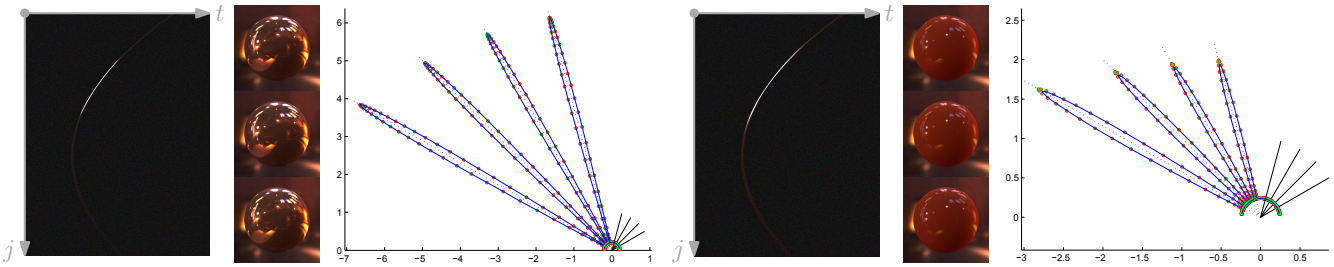

(1a) (1b) (1c) (2a) (2b) (2c)

Figure 7: *In single-patch settings, our approach can achieve near-perfect reconstruction with relatively low noise-levels. Two BRDFs were picked from Ngan et al.[2005]: (1) copper, (2) red plastic. Column (a) shows streak images; (b) shows rendered spheres under environment lighting using the ground truth (top) and the recovered BRDFs (middle with 1% of noise and bottom with 10%), respectively; (c) shows plots of the BRDFs in log-scale: the dotted red curve indicates the ground truth BRDF, the green curve represents the BRDF recovered with 1% of noise, and the blue curve shows that with 10%. See Table 1 for BRDF parameter values.*

Material	Ground Truth			Recovered (1% noise)			Error	Recovered (10% noise)			Error
	k_d	k_s	n	k_d	k_s	n	(log10)	k_d	k_s	n	(log10)
Copper	0.076	1.040		0.076	1.040		0.267	0.077	1.053		2.263
	0.041	0.609	40800	0.041	0.609	40800	-0.632	0.044	0.660	40800	2.871
	0.029	0.266		0.029	0.266		-0.702	0.027	0.252		2.298
Red plastic specular	0.242	0.058		0.242	0.058		-1.398	0.228	0.055		1.179
	0.031	0.045	8320	0.031	0.045	8320	-1.825	0.030	0.045	8320	0.170
	0.009	0.030		0.009	0.030		-1.880	0.009	0.032		0.845

Table 1: *Single patch simulated reconstruction: ground truth used to generate streak images, recovered values with* 1% *noise, and with* 10% *noise. The three rows per material are the RGB values. Our approach recovers the BRDF parameters accurately with 1% of noise. With 10% of noise added to the input, recovered BRDFs are still reasonable, although there were larger offsets on the diffuse components, which normally contain small values and are more sensitive to noise.*

principle RGB color channels could be realized, for example by using an optical parametric oscillator [Dudley et al. 1994] as light source or converting light from multiple laser sources into the visible. The angular coverage is usually between 0° to 15° in terms of half-angles.

The parameters fitted for red plastic are $k_d = 0.2175$, $k_s = 0.0508$ and $n = 7.448 \cdot 10^3$, and for copper are $k_d = 0.1010$, $k_s = 0.9472$, $n = 2.83 \cdot 10^4$. These values are in rough agreement with numbers cited by Ngan et al. [2005].

Multi patch data: We place two patches of copper and red plastic side-by-side and measure them using our experimental device. Figure 9-(2a) shows the streak image for both simultaneously. We recover the parametric BRDFs for them simultaneously, shown as rendered spheres in Figure 9-(2b), with error plots in 2(c) and 2(d). The parameters reconstructed for red plastic are $k_d = 0.3105$, $k_s = 0.0433$ and $n = 6.321 \cdot 10^3$, and for copper are $k_d = 0.1320$, $k_s = 0.8365$, $n = 3.120 \cdot 10^4$. Again, these values roughly agree with Ngan et al. [2005].

Validation using published gonioreflectometer data: We further evaluate our proposed method for the "around the corner" viewing mode using published gonioreflectometer data from the Cornell Reflectance Database [2001]. We employ a simulation setup with the geometry from Figure 2-(b). A single patch with the tabulated BRDF from the House Paint data is used as the input to generate the streak images. We then recover the parametric fit for the material using our algorithm. The recovered parameter values are: $k_d = (0.268, 0.431, 0.602)$, $k_s = (0.038, 0.041, 0.080)$, and $n = 11.6$. Renderings of spheres using the measured data and the recovered parameters match each other visually (see Figure 10).

7 Discussion

Our single-view point approach exploits the temporal dimension, but it introduces traditional problems of time-of-flight cameras in space-time resolution, signal to noise ratio and dynamic range.

Limited angular sampling: In a closed room, one can theoretically sample all the pairs of incoming and outgoing directions from an object point. In practice, we can sample only a subset of all half angles as the coverage is dependent on the field of view of the camera. Despite these limited angles, a 3-parameter BRDF model can be estimated. We also rely on the friendly reflectances of the sources and receivers, though this is not a fundamental limitation.

Space-time resolution: Our at-a-distance capture mechanism means that we cannot resolve small features with varying BRDF or surface normals. So our approach is suitable for coarsely segmented reflectance patches. Surfaces with rich surface details can be acquired accurately to the limit of the spatio-temporal resolution of the ToF device.

Time resolution limits our ability to perform linear inversion. Currently our camera only captures one spatial dimension, however, we can scan the laser in two dimensions over the scene to increase the sampling range of angles, and further sweep the camera.

Color: Since the laser operates at a single wavelength our images are monochrome and taken in the near infrared. Colored images could be taken with a white light supercontinuum source, a set of 3 lasers at different colors, or a tunable optical parametric oscillator.

Signal to noise ratio and capture time: The theoretical signal to noise ratio (SNR) of the combined streak camera system is about 1000:1. A common way to improve the SNR by several orders of magnitude is to bundle the same laser power into fewer pulses at a

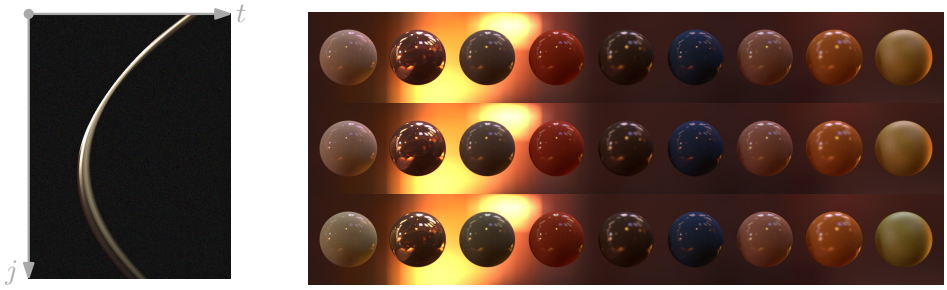

Figure 8: *With nine patches with different BRDFs picked from Ngan et al.[2005], our approach is able to obtain high-quality reconstructions: (left) the streak image; (right) rendered spheres under environment lighting using the ground truth (top) and the recovered BRDFs (middle with 1% of noise and bottom with 10%), respectively. See the supplementary material for BRDF parameter values.*

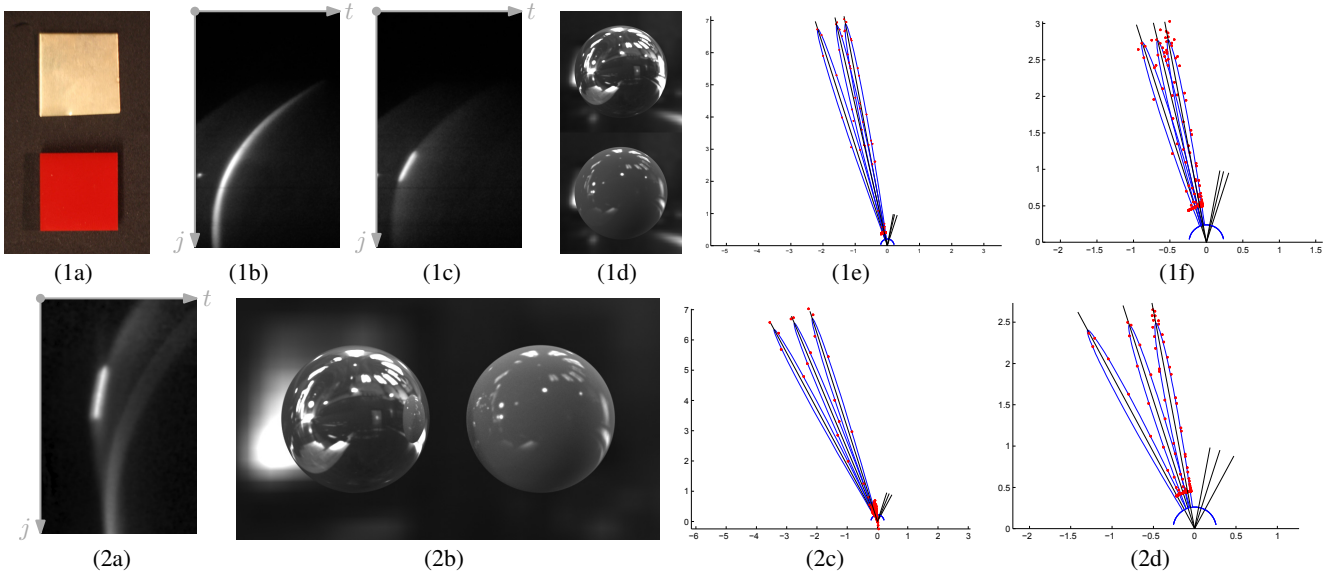

Figure 9: *For data using our experimental device our method obtains good results. Top row: single patch; Bottom row: two patches. (1a)Actual patches of copper and plastic used for acquisition; (1b) and (1c) Streak images taken by our ToF camera: a – copper, b – plastic; (1d) Spheres using the recovered BRDFs (copper – top, plastic – bottom); (1e) and (1f) Plots for the BRDFs: red dots indicate measured data points and blue curves are the recovered BRDFs. Bottom row: two patch results. (2a) Streak image for both materials together. (2b) Spheres rendered using recovered BRDFs (copper on left, plastic on right). (2c) and (2d): Error plots for the BRDFs.*

lower repetition rate but with the same pulse length.

Our acquisition time of 10-15 seconds per streak camera image is a result of our current hardware, which averages over a large number of frames to improve SNR. A commercially available laser with pulse energy of about 1 mJ could reduce the acquisition time to nanoseconds, while offering better SNR as well.

Dynamic range: Capturing specular peaks and weak diffuse reflections in a single photo is limited due to the camera dynamic range. We partially overcome this by using two different exposure photos.

Acquisition in the presence of ambient light: The time-of-flight cameras are well suited for "in-the-wild" acquisition in the presence of ambient light. Most ambient light is never detected by the sensor because of the short capture window. Even highly sensitive photon counting systems use these techniques to operate in daylight and over hundreds of thousands of kilometers; for example, laser links from Earth to the Moon and Mars, and commercial airborne LIDAR systems [Degnan 2002; Warburton et al. 2007].

Portability: While ultra-fast lasers have not yet reached a state of maturity to make them portable, this approach shows promise in creating portable and compact devices in the future. Our system can be extended for usage in unstructured environments with arbitrary geometry and lighting. Moreover, our algorithm in itself is not limited to a particular geometric configuration.

Laser speckle: Imaging devices using coherent light often suffer from laser specle noise. The laser coherence is, however, not maintained in multiple diffuse bounces. When the laser light returns to the camera after two or three bounces, it is no longer coherent and laser speckle is not observed.

Our work uniquely combines cutting edge research in ultra-fast optics with emerging topics in computer graphics. Our computational approach has been validated, but our physical prototype is a modification of electro-optic hardware which is expensive, currently non-portable, and may take years to become practical. But there are no specific fundamental challenges to improve these systems.

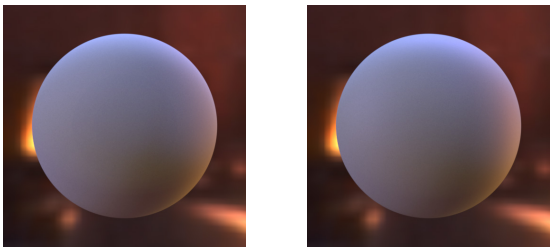

Figure 10: *Validation using published gonioreflectometer data from the Cornell Reflectance Database[2001]: (left) measured data, (right) the recovered BRDF.*

8 Conclusion

In this work, we have demonstrated a high speed photography device to acquire segmented scene reflectances using indirect reflections. We have identified the underlying constraints in using time-of-flight capture, including entanglement of light paths with the same path length. We demonstrate disentanglement for several patches, on the order of tens, by fitting to low dimensional parametric models.

While fast and smaller solid state lasers are coming, merging them with fast imaging devices is a clear logical step. We believe that this approach acquisition has potential to enable fast, portable, and remote BRDF capture devices. Without the need to instrument a scene, our work may spur applications like real-world material classification, real-time material editing, and relighting. Our approach can also be used when capturing complete BRDFs is not the ultimate goal but sampling a part of it for material detection and classification can suffice. In addition, around the corner recovery of material properties can enable radical applications, e.g., recovering malignant growth in endoscopy beyond the reach of a camera. Our approach also fits in the general spririt of computational photography to allow one to capture meaningful properties from a single camera viewpoint and then allow powerful post-capture operations, in this case to relight or edit materials.

Acknowledgements

We acknowledge Prof. Moungi Bawendi for making available his equipment and lab space. We would like to thank Rohit Pandharkar and Jason Boggess for their help at different stages of the project. The renderings of spheres are generated using the Mitsuba physically based renderer [Jakob 2010]. Funding was provided by the National Science Foundation under awards CCF-0644175, CCF-0811680 and IIS-1011919, and an Intel PhD Fellowship. Ramesh Raskar was supported by an Alfred P. Sloan Research Fellowship.

References

ASHIKHMIN, M. 2007. Distribution-based BRDFs. *University of Utah – Technical Report*.

ASHIKMIN, M., PREMOŽE, S., AND SHIRLEY, P. 2000. A microfacet-based BRDF generator. In *ACM SIGGRAPH 2000 Papers*, 65–74.

BAI, J., CHANDRAKER, M., NG, T.-T., AND RAMAMOORTHI, R. 2010. A dual theory of inverse and forward light transport. In *ECCV 2010*, 1–8.

CORNELL, 2001. Reflectance data. www.graphics.cornell.edu/online/measurements/reflectance.

DANA, K., VAN GINNEKEN, B., NAYAR, S., AND KOENDERINK, J. 1999. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics 18*, 1, 1–34.

DEGNAN, J. J. 2002. Asynchronous laser transponders for precise interplanetary ranging and time transfer. *Journal of Geodynamics 34*, 3-4, 551 – 594.

DONG, Y., WANG, J., TONG, X., SNYDER, J., LAN, Y., BEN-EZRA, M., AND GUO, B. 2010. Manifold bootstrapping for SVBRDF capture. In *ACM SIGGRAPH 2010 papers*, 98:1–98:10.

DUDLEY, J. M., REID, D. T., EBRAHIMZADEH, M., AND SIBBETT, W. 1994. Characterisics of a noncritically phasematched Ti:sapphire pumped femtosecond optical parametric oscillator. *Optics Communications 104*, 4,5,6 (January), 419–430.

FORSYTH, D., AND ZISSERMAN, A. 1990. Shape from shading in the light of mutual illumination. *Image Vision Comput. 8* (February), 42–49.

GARDNER, A., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2003. Linear light source reflectometry. In *ACM SIGGRAPH 2003 Papers*, 749–758.

GARG, G., TALVALA, E.-V., LEVOY, M., AND LENSCH, H. P. A. 2006. Symmetric photography: Exploiting data-sparseness in reflectance fields. In *Rendering Techniques 2006*, 251–262.

GHOSH, A., CHEN, T., PEERS, P., WILSON, C. A., AND DEBEVEC, P. 2009. Estimating specular roughness and anisotropy from second order spherical gradient illumination. *Computer Graphics Forum 28*, 4, 1161–1170.

GHOSH, A., CHEN, T., PEERS, P., WILSON, C. A., AND DEBEVEC, P. 2010. Circularly polarized spherical illumination reflectometry. In *ACM SIGGRAPH Asia 2010 papers*, 162:1–162:12.

GHOSH, A., HEIDRICH, W., ACHUTHA, S., AND O'TOOLE, M. 2010. A basis illumination approach to BRDF measurement. *Int. J. Comput. Vision 90* (November), 183–197.

GOLDMAN, D., CURLESS, B., HERTZMANN, A., AND SEITZ, S. 2010. Shape and spatially-varying BRDFs from photometric stereo. *IEEE PAMI 32*, 6 (june), 1060 –1071.

HAN, J. Y., AND PERLIN, K. 2003. Measuring bidirectional texture reflectance with a kaleidoscope. In *ACM SIGGRAPH 2003 Papers*, 741–748.

HASINOFF, S., DURAND, F., AND FREEMAN, W. 2010. Noise-optimal capture for high dynamic range photography. In *CVPR 2010*, 553 –560.

HAWKINS, T., EINARSSON, P., AND DEBEVEC, P. E. 2005. A dual light stage. In *Rendering Techniques*, K. Bala and P. Dutre, Eds., 91–98.

HOLROYD, M., LAWRENCE, J., HUMPHREYS, G., AND ZICKLER, T. 2008. A photometric approach for estimating normals and tangents. In *ACM SIGGRAPH Asia 2008 papers*, 133:1–133:9.

JAKOB, W., 2010. Mitsuba physically-based renderer. www.mitsuba-renderer.org.

KIRMANI, A., HUTCHISON, T., DAVIS, J., AND RASKAR, R. 2009. Looking around the corner using transient imaging. In *ICCV 2009*, 159–166.

KUTHIRUMMAL, S., AND NAYAR, S. K. 2006. Multiview radial catadioptric imaging for scene capture. In *ACM SIGGRAPH 2006 Papers*, 916–923.

LAWRENCE, J., BEN-ARTZI, A., DECORO, C., MATUSIK, W., PFISTER, H., RAMAMOORTHI, R., AND RUSINKIEWICZ, S. 2006. Inverse shade trees for non-parametric material representation and editing. In *ACM SIGGRAPH 2006 Papers*, 735–745.

LENSCH, H. P. A., KAUTZ, J., GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2001. Image-based reconstruction of spatially varying materials. In *Eurographics Workshop on Rendering*, 63–70.

LENSCH, H. P. A., KAUTZ, J., GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2003. Image-based reconstruction of spatial appearance and geometric detail. *ACM Transactions on Graphics 22*, 2, 234–257.

LIU, S., NG, T., AND MATSUSHITA, Y. 2010. Shape from second-bounce of light transport. In *ECCV 2010*, 280–293.

MA, W.-C., HAWKINS, T., PEERS, P., CHABERT, C.-F., WEISS, M., AND DEBEVEC, P. 2007. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Eurographics Symposium on Rendering*, 183–194.

MARSCHNER, S., WESTIN, S., LAFORTUNE, E., TORRANCE, K., AND GREENBERG, D. 1999. Image-Based BRDF measurement including human skin. In *Eurographics Workshop on Rendering*, 139–152.

MATHWORKS. 2011. Matlab Optimization Toolbox User's Guide *http://www.mathworks.com/help/pdf_doc/optim/optim_tb.pdf* Last accessed 4-September-2011. 6–1–6–18.

MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. In *ACM SIGGRAPH 2003 Papers*, 759–769.

MCALLISTER, D. 2002. *A Generalized Surface Appearance Representation for Computer Graphics*. PhD thesis, UNC Chapel Hill.

NAYAR, S. K., IKEUCHI, K., AND KANADE, T. 1991. Shape from interreflections. *Int. J. Comput. Vision 6* (August), 173–195.

NAYAR, S. K., KRISHNAN, G., GROSSBERG, M. D., AND RASKAR, R. 2006. Fast separation of direct and global components of a scene using high frequency illumination. In *ACM SIGGRAPH 2006 Papers*, 935–944.

NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental Analysis of BRDF Models. In *Eurographics Symposium on Rendering*, 117–226.

NICODEMUS, F. E., RICHMOND, J. C., AND HSIA, J. J. 1977. Geometrical considerations and reflectance. *National Bureau of Standards*, NBS Monograph 160 (October).

RUSINKIEWICZ, S. 1998. A new change of variables for efficient BRDF representation. In *Eurographics Workshop on Rendering*, 11–22.

SATO, Y., WHEELER, M., AND IKEUCHI, K. 1997. Object shape and reflectance modeling from observation. In *ACM SIGGRAPH 1997 papers*, 379–387.

SCHECHNER, Y., NAYAR, S., AND BELHUMEUR, P. 2007. Multiplexing for optimal lighting. *IEEE PAMI 29*, 8 (aug.), 1339 –1354.

SEITZ, S. M., MATSUSHITA, Y., AND KUTULAKOS, K. N. 2005. A theory of inverse light transport. In *ICCV 2005*, 1440–1447.

SEN, P., CHEN, B., GARG, G., MARSCHNER, S. R., HOROWITZ, M., LEVOY, M., AND LENSCH, H. P. A. 2005. Dual photography. In *ACM SIGGRAPH 2005 Papers*, 745–755.

WANG, J., ZHAO, S., TONG, X., SNYDER, J., AND GUO, B. 2008. Modeling anisotropic surface reflectance with example-based microfacet synthesis. In *ACM SIGGRAPH 2008 papers*, 41:1–41:9.

WANG, J., DONG, Y., TONG, X., LIN, Z., AND GUO, B. 2009. Kernel nystrom method for light transport. In *ACM SIGGRAPH 2009 papers*, 29:1–29:10.

WARBURTON, R. E., MCCARTHY, A., WALLACE, A. M., HERNANDEZ-MARIN, S., HADFIELD, R. H., NAM, S. W., AND BULLER, G. S. 2007. Subcentimeter depth resolution using a single-photon counting time-of-flight laser ranging system at 1550 nm wavelength. *Opt. Lett. 32*, 15 (Aug), 2266–2268.

WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *ACM SIGGRAPH 1992 papers*, 265–272.

WENGER, A., GARDNER, A., TCHOU, C., UNGER, J., HAWKINS, T., AND DEBEVEC, P. 2005. Performance relighting and reflectance transformation with time-multiplexed illumination. In *ACM SIGGRAPH 2005 Papers*, 756–764.

WEYRICH, T., LAWRENCE, J., LENSCH, H. P. A., RUSINKIEWICZ, S., AND ZICKLER, T. 2009. Principles of appearance acquisition and representation. *Foundations and Trends in Computer Graphics and Vision 4*, 2, 75–191.

YU, Y., DEBEVEC, P., MALIK, J., AND HAWKINS, T. 1999. Inverse global illumination: recovering reflectance models of real scenes from photographs. In *ACM SIGGRAPH 1999 papers*, 215–224.

ZICKLER, T., ENRIQUE, S., RAMAMOORTHI, R., AND BELHUMEUR, P. 2005. Reflectance sharing: Image-based rendering from a sparse set of images. In *Eurographics Symposium on Rendering*, 253–264.