Multibounce Light Transport Analysis using Ultrafast Imaging for Material Acquisition

by

Nikhil Naik

B.Tech. Electronics & Telecommunication, College of Engineering Pune, 2010
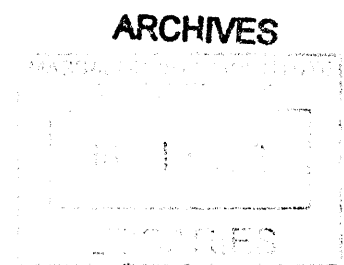
Submitted to the Program in Media Arts & Sciences,
School of Architecture and Planning,
in partial fulfillment of the requirements for the degree of

Master of Science in Media Arts & Sciences

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2012

Author . . . .
Program in Media Arts & Sciences,
School of Architecture and Planning,
May 12, 2012

Certified by. . . . .
Ramesh Raskar
Associate Professor
Program in Media Arts & Sciences
Thesis Supervisor

Accepted by . . . .
Mitchel Resnick
LEGO Papert Professor of Learning Research
Academic Head, Program in Media Arts and Sciences

Multibounce Light Transport Analysis using Ultrafast Imaging for Material Acquisition

by

Nikhil Naik

Submitted to the Program in Media Arts & Sciences,
School of Architecture and Planning,
on May 12, 2012, in partial fulfillment of the
requirements for the degree of
Master of Science in Media Arts & Sciences

Abstract

This thesis introduces a novel framework for analysis of multibounce light transport using time-of-flight imaging for the applications of ultrafast reflectance acquisition and imaging through scattering media. Using ultrafast imaging and ultrafast illumination, we analyze light indirectly scattered off materials to provide new insights into the important problem of material acquisition. We use an elegant matrix based representation of light transport, which enables scene reconstruction using standard optimization techniques. We demonstrate the accuracy and efficiency of our methods using various simulations as well as an experimental setup.

In particular, we develop the concept of 'in the wild' reflectance estimation using ultrafast imaging. We demonstrate 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-fight camera. As compared to lengthy or highly calibrated reflectance acquisition techniques, we demonstrate a device that can rapidly and simultaneously capture meaningful reflectance information of multiple materials.

Furthermore, we use this framework to develop a method for imaging through scattering media using ultrafast imaging. We capture the diffuse scattering in the scene with a time-of-flight camera and analyze the multibounce light transport to recover albedo and depth information of planar objects hidden behind a diffuser. The methods developed in this thesis using ultrafast imaging can spur research with novel real-time applications in computer graphics, medical imaging and industrial photography.
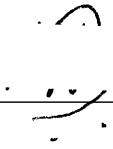
Thesis Supervisor: Ramesh Raskar
Title: Associate Professor
Program in Media Arts & Sciences

Multibounce Light Transport Analysis using Ultrafast Imaging for Material Acquisition

by

Nikhil Naik

The following person served as a reader for this thesis:
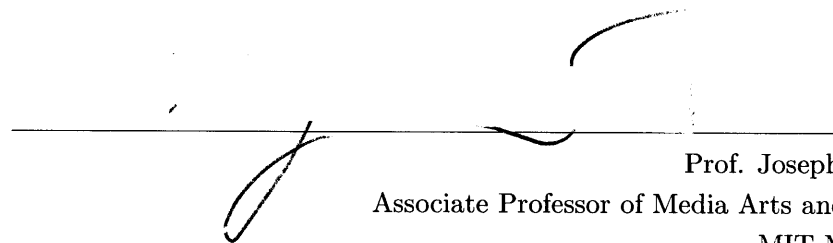
Reader: _____

Prof. Kavita Bala
Associate Professor of Computer Science
Cornell University

Multibounce Light Transport Analysis using Ultrafast Imaging for Material Acquisition

by

Nikhil Naik

The following person served as a reader for this thesis:

Reader: _____

Prof. Joseph Paradiso

Associate Professor of Media Arts and Sciences

MIT Media Lab

Acknowledgments

This work would not have been possible without the help of my mentors and colleagues. I would like to thank my advisor Prof. Ramesh Raskar for his guidance, support, and encouragement over the past two years. I also consider myself fortunate to get the opportunity to work closely with Prof. Kavita Bala. I would like to thank her for her invaluable guidance. Furthermore, I thank Prof. Joe Paradiso for helpful comments and suggestions.

I am also very thankful to Andreas Velten, Otkrist Gupta, Roarke Horstmeyer, Abhijit Bendale, Rohit Pandharkar, Di Wu, Christopher Barsi, Douglas Lanman, Taya Leary and all the members of Camera Culture group for their help and support. I would like to thank Priya for her helpful critique and support. Finally I would like to thank my family, in particular my parents and sister, for their love and care.

Contents

6

List of Figures

13

14

List of Tables

Chapter 1

Introduction

In the following sections we describe the motivations for our research, and layout the major contributions of our work.

1.1 Motivation

New imaging platforms, sensors as well as illumination techniques often enable novel ways of capturing, processing and interpreting images. Examples include the use of X-ray imaging for medical radiography, archeology, and the use of infrared imaging for night vision and nondestructive testing. In this thesis, we introduce an ultrafast imaging and an ultrafast illumination platform for the applications of reflectance capture and imaging through scattering media.

Specifically we make use of a picosecond resolution 'streak camera' and a femtosecond pulsed laser, tools which have been traditionally employed for studying femtochemistry experiments. However, unlike femtochemistry experiments looking at microscopic samples, we repurpose this imaging tool to develop a framework for potentially capturing light in room-sized natural scenes and analyzing the 'space-time' images obtained using the camera.

Ultrafast cameras typically have a time resolution in the order of a few picoseconds, enabling them to capture up to a trillion (10^{12}) frames per second. Light travels a distance of $0.3\,mm$ in one picosecond ($10^{-12}\,s$). Therefore, ultrafast cameras are essentially able to capture light as it travels. It provides us with the unique ability of separating and analyzing different light paths, including light reaching the camera directly from the source as well as light bouncing off different objects in a scene.

In this thesis, we are interested in looking at light indirectly captured by the camera in the form of light scattered off the surface of different objects in a scene. Using an ultrafast camera, we are able to separate out and analyze the multiple bounces of light. This analysis of 'multibounce light transport' provides new insights into the important problems of reflectance acquisition and imaging through scattering media.

Reflectance acquisition, or measuring and modeling how different materials reflect light, is a very important problem in Computer Graphics. It has many applications including photo-realistic rendering, image relighting as well as material identification and matching. We show that ultrafast imaging and multibounce light transport analysis can not only develop a faster, more efficient reflectance acquisition technique but also allow novel applications by allowing 'in the wild' acquisition.

Additionally, imaging through scattering media is another important problem where analysis of diffuse scattering can provide novel insights. This problem has applications in industrial imaging, medical imaging, especially of tissues, as well as atmospheric imaging. For these applications, traditional methods use coherent ballistic photons or penetrating wavelengths. We demonstrate a novel method which makes use of scattered photons to reconstruct a scene imaged through a diffuser.

1.2 Novelty and Contributions

This thesis makes important contributions to analyzing multibounce light transport using ultrafast imaging as well as many specific contributions to the problems of reflectance acquisition and imaging through scattering media.

a) Analysis of Multibounce Light Transport: We capture and analyze indirect bounces of light off objects to obtain information about their appearance and reflectance properties. In traditional techniques, the indirect bounces of light are typically ignored or blocked by the imaging setup. We instead treat scattered light as a useful tool for the applications of appearance capture by using ultrafast imaging.

b) Linear System Representation of Light Transport: We model the multibounce light transport as a linear system by separating the geometric contributions from material

17

appearance. Using this light transport matrix allows us to develop elegant and fast methods of solving for material properties using a wide array of powerful optimization and linear algebra tools. We demonstrate the accuracy and efficiency of this approach which can inspire usage of similar light transport analysis models for different applications.

c) Ultrafast Reflectance Acquisition: We present a new acquisition approach to reflectance measurement using 'demultiplexed scattering' and ultrafast imaging. Our approach is unique in two ways. We exploit ultrafast imaging to enable rapid acquisition of material reflectance properties; and we use indirect observation to acquire many materials simultaneously, and in fact, even measure reflectance properties of non-line-of-sight objects. Our method uses advanced opto-electronic devices and analysis of multibounce light transport to pave the way for uninstrumented reflectance capture outside laboratory settings. These 'in the wild' acquisition methods can have many novel applications in industrial photography and medical imaging. Chapters 3 and 4 describe this method in detail.

d) Imaging Through Scattering Media: We present a new approach for imaging through scattering media by imaging the light reflected off materials behind a scattering medium using time-of-flight (ToF) cameras. Unlike many prior techniques which ignore secondary and higher bounces of light, we demonstrate that the analysis of scattered photons is useful for recovering albedo and depth information of scenes imaged through a diffuser. We analyze real world scenes and develop a mathematical framework which can inspire novel applications in tissue imaging as well as atmospheric imaging. Chapters 5 and 6 describe this method in detail.

Chapter 2

Background and Related Work

The related work is roughly grouped into prior work in ultrafast imaging, followed by reflectance acquisition and imaging through scattering media.

2.1 Ultrafast Imaging and Illumination Techniques

In most practical applications, ultrafast imaging at picosecond resolution is performed using streak cameras which deposit photons across a spatial dimension, rather than integrating them in a single pixel. Picosecond electron optical streak cameras became practical in the early 1960s. Campillo and Shapiro [2] provide a review of the early designs of these cameras. We also make use of pulsed femtosecond illumination. This illumination technique is typically employed in biological tissue imaging techniques such as Optical Coherence Tomography (OCT) [3] and two-photon microscopy [4].

Recently, there has been a lot of interest to use picosecond cameras and femtosecond pulsed illumination in free space, which enables imaging at speeds fast enough to observe radiometric changes over time. Such imaging methods have provided the opportunity for analyzing scenes in many innovative ways that are not possible with a standard still cameras.

The use of this imaging setup in free space was first proposed by Raskar and Davis[5]. Kirmani et al. demonstrated that three bounce analysis of a ToF camera can recover hidden 1-0-1 planar barcodes [6]. Multipath analysis was used to track the motion of an occluded target by Pandharkar et al. [7]. The algorithms to reconstruct geometry of occluded objects "around a corner" were developed by Velten et al. [8].

This camera setup was also used for capturing 2D videos at a trillion frames per second

by Velten et al. [9]. Wu et al. [10] used these ToF videos for analyzing direct and global light transport and demonstrated its computer vision applications such as edge detection.

2.2 Ultrafast Reflectance Acquisition

Nicodemus et al. [11] 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) [12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23]. These acquisition processes are often lengthy, requiring extensive laboratory settings with calibration and hours or even days of acquisition. Weyrich et al. [24] 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., [16, 25, 26, 15, 27, 28, 29]). Various approaches decrease acquisition costs and increase coverage by exploiting properties of BRDFs including reciprocity, separability, spatial smoothness, and compressibility [30, 31, 32, 33, 22]. Lighting configurations and variations have been considered including polarization and structured illumination [34, 31, 35, 36].

Most BRDF techniques directly, or indirectly through a mirror surface [24], view and image the material sample with some notable exceptions [37, 38, 39]. For instance, Hawkins et al. [38] 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 [40, 41]. Recently there has been interest in recovering scene properties such as geometry and albedo from multiple bounces of light [42, 43, 31, 44, 45], and also on recovering shape and material simultaneously [46].

In contrast to prior techniques we do not directly image the surface sample. Rather, we indirectly image a diffuse surface and use ToF 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) and simultaneous
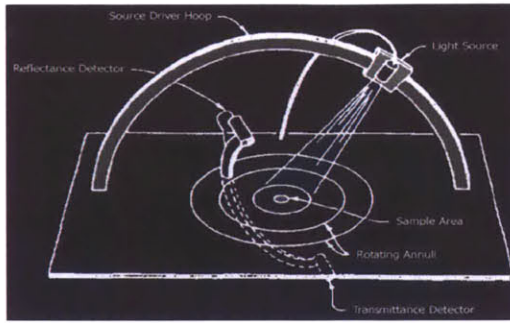
Gonioreflectometer-Schematic
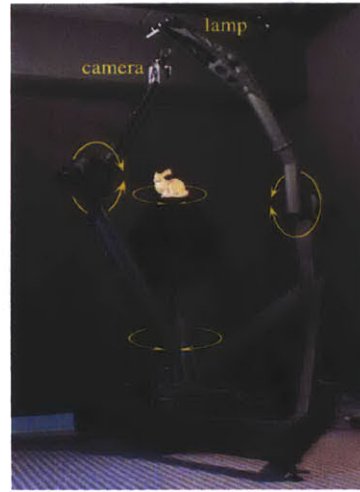
Gonioreflectometer - Physical Setup
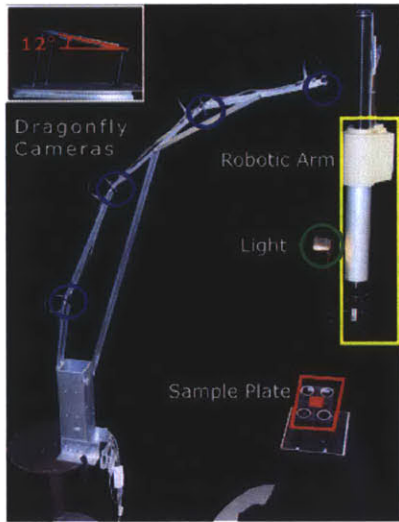
Sun et al. 2006
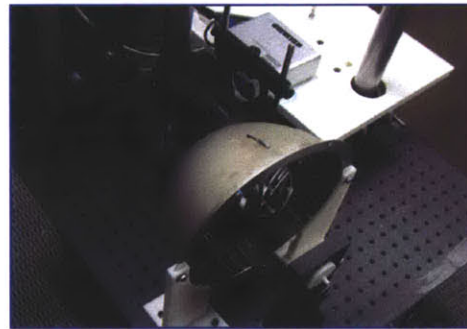
Ghosh et al. 2010
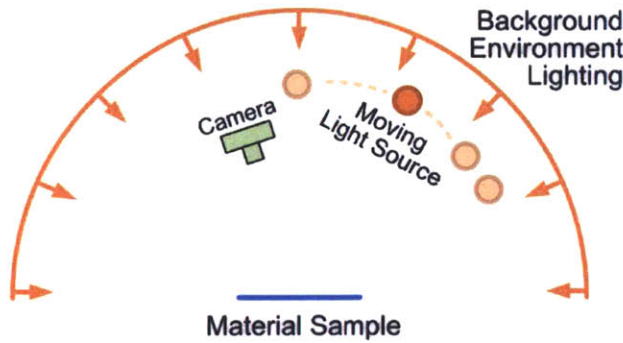
Dong et al. 2010

Figure 2-1: BRDFs are typically acquired using Gonioreflectometers which require extensive instrumentation and long calibration & acquisition time. Hence recent research has focused on novel acquisition devices which can address these problems. The figure shows gonioreflectometers and some of the important related acquisition techniques.

acquisition of reflectances of multiple patches (tens of materials) over a range of angular measurements. Chapter 3 explains our contributions in detail.

2.3 Imaging through Scattering Media

Imaging, at the microscopic level of living cells to the astronomical level of the atmosphere, is often hindered by natural scattering media, like tissue or fog. Methods to circumvent such scatering media are currently a challenge in imaging and optics. In addition to better imaging, progress in this area presents interesting applications in machine vision, including looking through smoke and even imaging what lies behind a curtain.

Light traveling through a discrete scattering medium gets split into two components: a coherent and an incoherent component. The coherent (ballistic) component propagates through the medium undeviated in the forward direction. The incoherent (diffuse) component undergoes multiple light scattering or a random walk in the random medium.

Reconstruction of images from diffuse scattered light through a scattering medium, known as the 'inverse scattering problem', is a difficult experimental and theoretical problem [47]. Therefore traditional techniques have focused on reducing or eliminating the diffuse component and making use of the ballistic component.

Wang et al. [48] were first to look at biological tissues, like human breast tissue, using time-gating. Other methods make use of fluorescence absorption [49] to reduce the multiple scattered light by its preferential absorption in the medium itself.

Ballistic photon imaging of biological in the near-infrared frequency has given rise to the whole field of Optical Coherence Tomography (OCT). However in heavily scattering media, ballistic photons have a very low probability of occurrence, and photons with a relatively few scattering events are also rare. Therefore OCT suffers from the problems of low signal-to-noise ratio (SNR) and expensive electro-optical devices. Schmitt et al. [50] provides an extensive survey of OCT.

Unlike OCT, Diffuse Optical Tomography (DOT) [51], makes use of diffusely scattered light using various light propagation models. In DOT, a scattering medium to be imaged is illuminated by Near-IR light using an array of sources and the scattered light is observed using an array of detectors either in reflection mode or transmission mode. The images obtained are used to infer the localized optical properties such as scattering and absorption

coefficients of the illuminated material. DOT faces significant challenges in obtaining high resolution images of complicated structures, due to the complex light transport inside these media as well as the ill-posed nature of the inverse problem.

In contrast to these prior approaches, our work does not ignore the diffuse component of light. Our work also does not try to understand global light transport using probabilistic models. Instead, we use ultrafast imaging to look at both ballistic and diffuse light transport and we model it as a linear system using a light transport matrix. We separate out the geometric component from the scene characteristics and pose an optimization problem to solve for albedo and depth of objects imaged through scattering media. Chapter 5 explains our contributions in detail.

Chapter 3

Ultrafast Reflectance Acquisition

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 [24] 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 identifying materials.

We present a new approach to measuring reflectance. Our approach is unique in two ways: it exploits ToF imaging to achieve rapid acquisition of materials; and it uses 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 ultrafast 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 multibounce light transport; or they must explicitly separate direct and indirect lighting from all bounces.

3.1   Contributions

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 ToF imaging to capture reflectances. Our proof-of-concept system

demonstrates first steps towards rapid material reflectance acquisition.

b) We exploit space-time images captured by a ToF camera that images 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. 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 ToF principles and indirect measurement, our solution is the first 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.

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 enable rapid, "in-the-wild" measurement of real-world scenes without instrumentation.

We now describe our ToF 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.

3.2 BRDF : Definition and Acquisition

In computer graphics, photo-realistic rendering, as the name suggest, strives to generate physically accurate pictures. For this purpose, an analysis of global and local illumination is performed. The global illumination analysis computes the light reaching any particular point in the scene from its environment. The local illumination analysis computes the relation between the incoming and outgoing point at a particular scene point. The bi-directional reflectance distribution function (BRDF), is a tool for describing the distribution

25
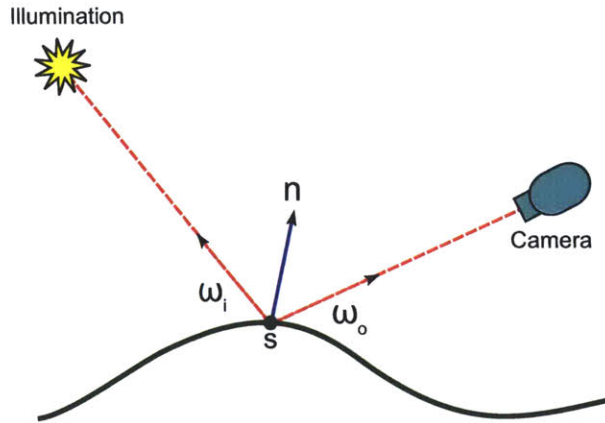
Figure 3-1: BRDF Definition: Diagram showing vectors used to define the BRDF of a scene point s. (Source: Wikipedia)

of reflected light at the surface of a scene point.

Given an incoming light ray, the BRDF calculates a ratio of light reflected in a particular outdoor direction from a point. Note that the BRDF model assumes that light does not scatter inside the surface to which the point belongs. Sub-surface scattering is covered by the bi-directional sub-surface scattering reflectance distribution function (BSSRDF).

Mathematically, BRDF at a point s, as seen in Figure 3-1 is defined as follows -

$$f_r(\omega_i, \omega_o) = \frac{dL_o(\omega_o)}{L_i(\omega_i) \cdot \cos \theta_i} \tag{3.1}$$

where ω_i, ω_o are vectors in incoming and outgoing direction direction, L is radiance, and θ_i is the angle between ω_i and the surface normal n, respectively.

Traditionally BRDFs are acquired using gonioreflectometers. A gonioreflectometer consists of a light source capable of illuminating the material to be measured from different directions and a sensor which captures light reflected from the material from different angles. Both light source and sensor can rotate in a hemisphere, thereby capturing the four dimensions of the BRDF.

As the number of possible incoming and outgoing angles tend to be quite large, BRDF acquisition using gonioreflectometers is a time-consuming process. Typical acquisition process requires hours of capture time followed by post-processing for representing BRDFs.

3.2.1 Need for Better Acquisition Techniques

As mentioned before, traditional BRDF acquisition using gonioreflectometers is a slow, cumbersome process. Naturally this has lead researchers to develop novel designs and algorithms for improving the acquisition process. A brief survey of related work in this area is presented in Section 2.2.

However, state of the art BRDF acquisition techniques still have a scope for improvement. These acquisition techniques require instrumenting the environment of the material being acquired, in one way or the other. In fact, many of them require large, complicated setups with extensive hardware, as it can be seen in Figure 2.2. Moreover, they require lengthy calibration routines before the actual acquisition process. Most techniques need proximity to the acquisition sample and ability to directly illuminate the sample.

So our motivation is to use the capabilities of ultrafast imaging to address these problems as well as enable novel applications for BRDF acquisition. In summary, our goal is to work towards a reflectance measurement technique which can -

1. Reduce calibration & acquisition time

2. Eliminate environment instrumentation

3. Enable acquisition without proximity/contact

We hope that this work will inspire real-time "in the wild" reflectance acquisition and enable many novel applications thereof.

3.3 Using Time-of-Flight

In this section we explain the fundamentals of ToF imaging and the key idea of using indirect observation and multibounce light transport for BRDF acquisition.

The traditional light transport models in computer graphics and vision assume that the speed of light is infinite. This assumption allows them to model the irradiant flux i.e. the rate of incidence of photons to be independent of time. This model is known as the 'steady-state' model of light transport. In contrast, the 'transient' light transport model takes the finite speed of light ($\sim 3 \times 10^8\, m/s$). Therefore light traveling along different paths

27
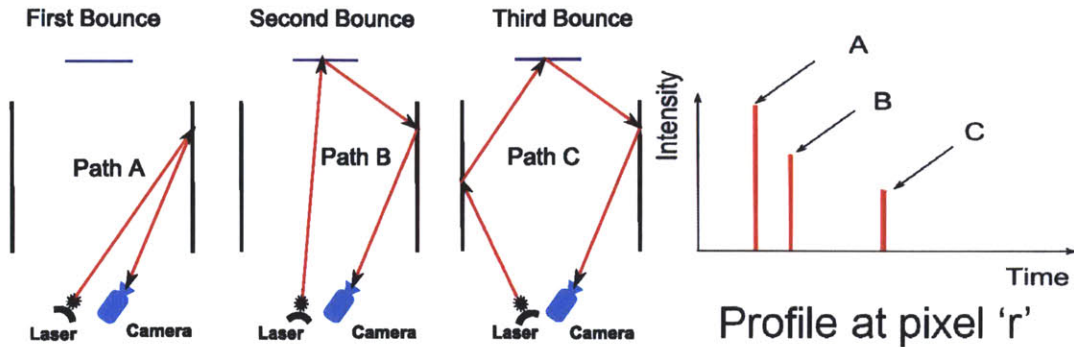
Figure 3-2: The ToF camera enables us to separate multiple bounces of light. We pick out the third bounce from path C which corresponds to indirect observation

in a scene takes different intervals of time to travel. This model gives rise to a time-varying irradiance function.

Figure 3-2 demonstrates this concept using a plot of the intensity versus time profile at a point in the scene for different light paths. A laser shines on different points in the scene and the intensity (or irradiance) versus time profile at the point r is measured. The intensity falls off from paths A to C as the total distance traveled by photons (or 'path length') increases, due to the inverse square law of power. Moreover the time profile moves towards the right as light takes more time to travel to point r as path length increases.

Of course, this would be the response of an imaging setup which consists of a camera with time-resolution higher than the speed of light, which would be able to resolve between any two light paths. In reality, most commercial high-speed cameras are too slow as compared to the speed of light. The scientific ultrafast imaging systems which can be modified for real world experiments with natural scenes are picosecond 'streak cameras'. Light travels a distance of $0.3\,mm$ in one picosecond ($10^{-12}\,s$). This time resolution is sufficient to resolve between different individual light paths in a natural scene. Furthermore, this ultrafast imaging hardware can be combined with pulsed illumination using a femtosecond impulse laser with pulse widths of the order of $10^{-15}\,s$.

3.3.1 Multibounce Light Transport

An important advantage of analysis of transient light transport that it allows us to clearly differentiate between multiple bounces of light. As seen in Figure 3-2, the first bounce

28
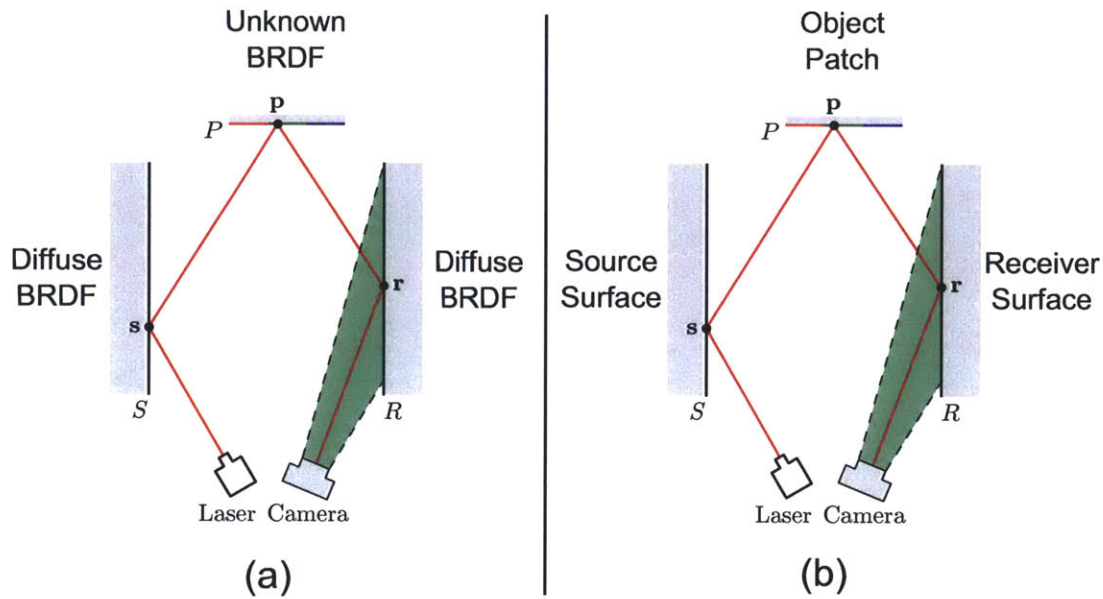
Figure 3-3: (a) The geometric configuration of the BRDF acquisition consisting of three planar surfaces. (b) The reflectance profiles of the three surfaces. The left and right surfaces have diffuse BRDFs, while the middle surface has an unknown BRDF.

(path A), second bounce (path B) and third bounce (path C) all have very different path lengths and any one of them can be handpicked for analysis separately. This enables direct observation of properties such as distance, and the difference between primary,secondary and other reflections. In particular, we concentrate on the tertiary bounce, seen in path C. If we have a pulsed laser source illuminating the surface on the left, some of the light scattered off the left surface will illuminate the middle surface. Furthermore, the light which reflects off the middle surface to the right surface will be imaged by an ultrafast camera. This setup essentially allows for both indirect illumination and indirect observation of the middle surface.

3.3.2   Why Indirect Observation?

Consider a geometric configuration depicted in Figure 3-3-(a), in which we denote the left surface as 'source surface' and the right surface as 'receiver surface'. The middle surface is a planar patch of a material with unknown BRDF while the source and receiver surfaces have known diffuse BRDFs. For a path traversed by a ray as seen in Figure 3-3-(b), the intensity observed by the camera is a product of the BRDFs of the three surfaces and the fall-off due to inverse square law (we'll derive the exact formulation later). As the BRDFs of source
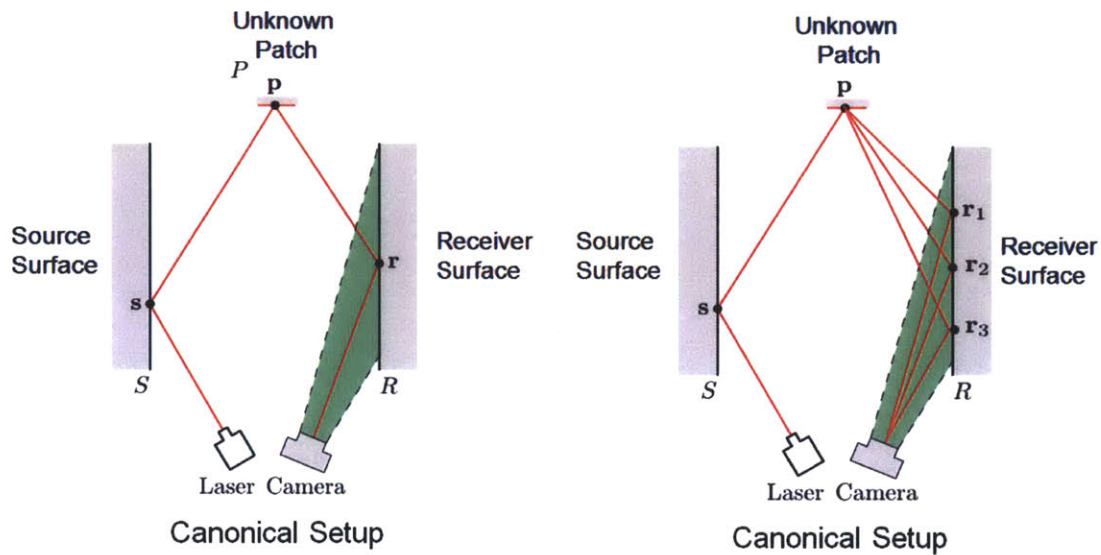
Figure 3-4: Indirect observation: The light path on the left figure corresponds to one BRDF observation. A single photo of the surface allows multiple observations as seen on the right

and receiver surfaces are known, this measurement corresponds to a BRDF observation for a unique incoming and outgoing angle for the unknown patch.

As the camera is imaging the complete right surface, we obtain multiple BRDF observations corresponding to one incoming angle and many outgoing angles from a single image, as seen in Figure 3-4-(a). Another set of observations, with a different incoming angle, can be obtained by simply choosing a different illumination point on the source surface. Therefore indirect observation gives us a large number of unique BRDF observations using a few images obtained with different laser positions.

3.3.3 Exploiting Time-of-Flight

It can be easily seen that the canonical setup from Figure 3-3-(a) can be used by a normal still camera to obtain BRDF measurements for a single unknown material patch in the middle. So this configuration does not really provide any important advantage over many of the prior techniques, described in Section 2.2. Now consider an unknown surface made up of multiple materials, as seen in Figure 3-5-(left). A normal still camera will integrate the two rays coming from patches p_1 and p_2 at the receiver point r. Therefore we cannot tease out the contributions from the individual patches, which correspond to separate BRDF measurements of the two patches.
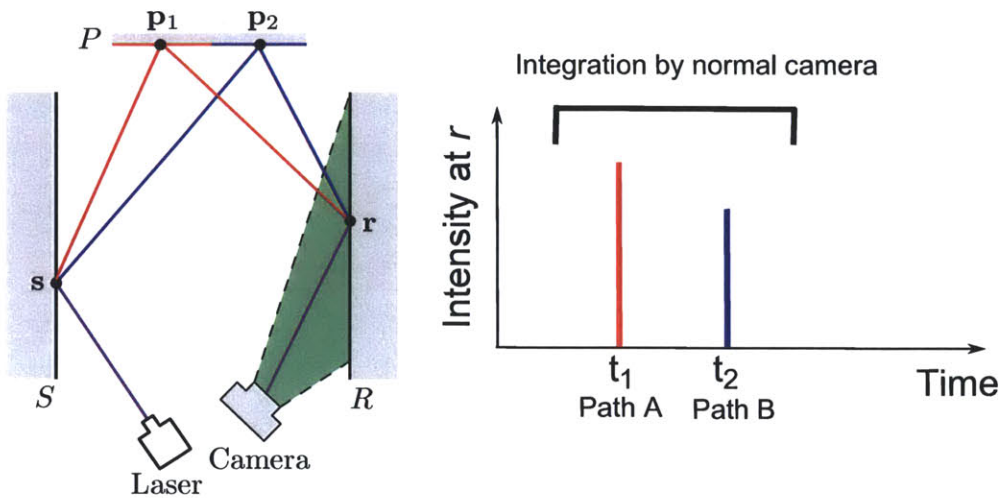
Figure 3-5: Time-of-flight imaging allows to disambiguate BRDF observations from multiple patches at a single point. The space-time response on the right demonstrates that the observations from the two patches are separated in time. A regular camera will integrate these two responses.

However these two rays have different path lengths. They arrive at point r at two different instances of time t_1 and t_2 as seen in Figure 3-5-(right). A ToF camera can be used to separate these observations and obtain BRDF measurements for the two patches 'simultaneously'. In this manner, ToF imaging combined with indirect observation allows simultaneous measurement of many BRDF values. Furthermore, the capture time is theoretically in picoseconds as the ToF images require only picoseconds of exposure. The following sections describe theoretical formulation of this problem and methods for ultrafast and simultaneous measurement of multiple BRDFs.

3.4   Geometry of Acquisition

Now we formally describe the canonical geometric setup to acquire reflectances, as seen in Figure 3-6-(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, s, r, 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 3-6-(b) the patch being measured P is not directly visible to the camera. In this case, the source
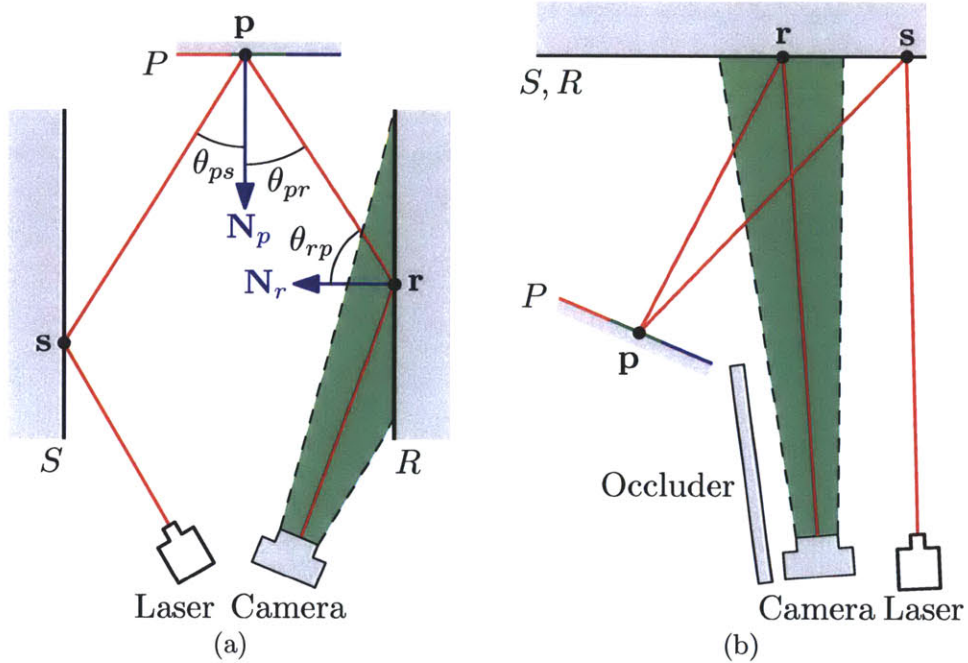
Figure 3-6: We explore two acquisition setups, left: canonical setup, right: around-the-corner setup. To measure the patch P, the laser illuminates S, and the camera images R.

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 3-6-(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 3-6-(a), but our physical experimental setup more closely matches Figure 3-6-(b).

3.5 Streak Image Formation

Consider a scene shown in Figure 3-6-(a). The femtosecond laser(L) emits a short light pulse and is pointed towards the Source(S) to form a laser spot s. The light reflected by the source which is incident on the patch P, gets reflected back and a part of it reaches the receiver wall(R). A streak camera focused on R captures the time varying image on the wall at a time resolution of 2 picoseconds. The streak camera has one spatial dimension. For
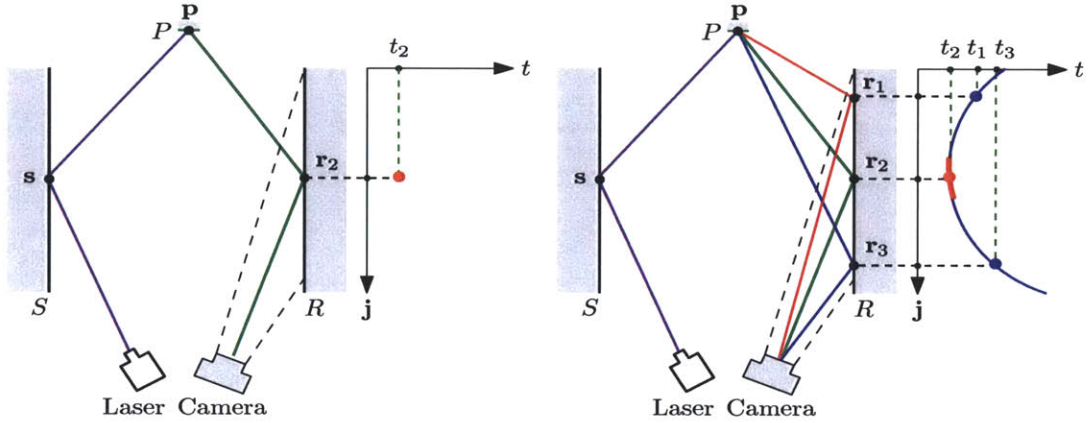
Figure 3-7: The planar geometry gives rise to characteristic hyperbolic streak images. The specular peak is directly observed at point r_2.

each laser position s, a $2D$ streak image (1 spatial and 1 temporal dimension) is recorded.

The camera captures the light reflected along the path: $\mathbf{s} \to \mathbf{p} \to \mathbf{r}$. As shown in Figure 3-6-(a), given a point $\mathbf{p}$ with normal $\mathbf{N}_p$, θ_{ps} is the angle made by $\mathbf{s}$ at $\mathbf{p}$ 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 $\mathbf{r}$.

For a single infinitesimally small patch P, the reflected spherical wavefront propagating from the patch reaches different points on the wall at different times creating a hyperbolic curve in the streak image. We now describe the mathematical formulation of the imaging process.

3.5.1 Bounce Reduction

As seen in Figure 3-6-(a), the optical path for light travel consists of four segments : (1) from the laser L to the source position s, (2) from s to the unknown patch p, (3) from p to a point on the receiver r, and (4) finally from r to the camera C where it is recorded. The first and the fourth segment are directed segments and do not involve diffuse scattering. This fact allows us to pre-calibrate for these segments and effectively reduce the tertiary scattering problem to a primary (single) scattering problem.

More concretely, suppose the camera records the streak image $I_C(p,t)$, where p is the pixel coordinate and t is time. In I_C, $t = 0$ corresponds to the instant the laser pulse is emitted from L. Then I_C is related to the intensity $I_R(r,t)$ of light incident on the receiver

33

plane by the transformation

$$I_R(r, t) = I_C(H(r), t - ||s - L|| - ||C - r||).\tag{3.2}$$

Here H is the projective transformation (homography) mapping coordinates on R to camera coordinates. The time shift by the distance from camera to screen, $||C - r||$, varies hyperbolically with the pixel coordinate r. Since the receiver R has a known planar geometry, H, $||L - B||$ and $||C - r||$ can be computed in advance. Note there is no $\cos(\theta)$ factor or $1/r^2$ fall off in the above formula as the camera integrates over more pixels for oblique and distant patches. For this to hold, it is also important that R is Lambertian, as we assume.

In summary, equation 3.2 reduces the problem to a single scattering problem, with an unfocused point source at s emitting a pulse at $t = 0$ and an unfocused virtual array of receivers on R recording the intensity of the reflected wavefront, $I_R(r, t)$.

3.5.2 Image Formation Model

The homography correction from Equation 3.2 allows us to consider a simplified scenarios of three surfaces, S, P and R. The surface S is illuminated by the laser at s. The surface R (receivers) can be assumed to host a virtual array of ultrafast photo-detectors. The virtual photo-detectors create an image $I_R(r, t)$ as intensity pattern of the incoming light as a function of time, t, and the position r. Hence the image, $I_R(r, t)$, is the intensity observed at $r \in R$ at time t. Experimentally, the virtual photo-detectors are realized by using a Lambertian planer surface R observed by a streak camera with picosecond time resolution (Figure 3-6.

Ignoring time, a regular camera viewing R would measure:

$$\frac{I_0 \, \rho_S \, \rho_R}{\pi} f_r(\mathbf{s}, \mathbf{p}, \mathbf{r}) \frac{\cos \theta_{ps} \cos \theta_{pr} \cos \theta_{rp}}{||\mathbf{s} - \mathbf{p}||^2 \, ||\mathbf{r} - \mathbf{p}||^2} \tag{3.3}$$

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:

$$g(\mathbf{s}, \mathbf{p}, \mathbf{r}) = \frac{I_0 \, \rho_S \, \rho_R}{\pi} \frac{\cos \theta_{ps} \cos \theta_{pr} \cos \theta_{rp}}{||\mathbf{s} - \mathbf{p}||^2 \, ||\mathbf{r} - \mathbf{p}||^2} \tag{3.4}$$

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

$$\int_{\mathbf{p}\in P} g(\mathbf{s},\mathbf{p},\mathbf{r})f_r(\mathbf{s},\mathbf{p},\mathbf{r})\,\mathrm{d}\mathbf{p} \tag{3.5}$$

We define the streak image to be a 2D image, I_R. With an ultrafast 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:

$$I_R(\mathbf{r},t) = \int_{\mathbf{p}\in P'} g(\mathbf{s},\mathbf{p},\mathbf{r})f_r(\mathbf{s},\mathbf{p},\mathbf{r})\,\mathrm{d}\mathbf{p} \tag{3.6}$$

where, $P' \subseteq P$ 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 3-7. 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.

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.

3.5.3 Shape Characteristics

As mentioned before in 3.5.1, the optical path is a sum of four separate segments. 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. This is depicted in Figure 3-7. In the figure, the total optical path length for receiver point r_2 is least amongst all the points imaged by the camera, including r_1 and r_3. Therefore it is located on the left of all the other points on the time axis in the streak image I_R. Moreover, the point r_2 is the mirror direction for the source s, which means that the specular peak is observed directly from the streak image at this point.
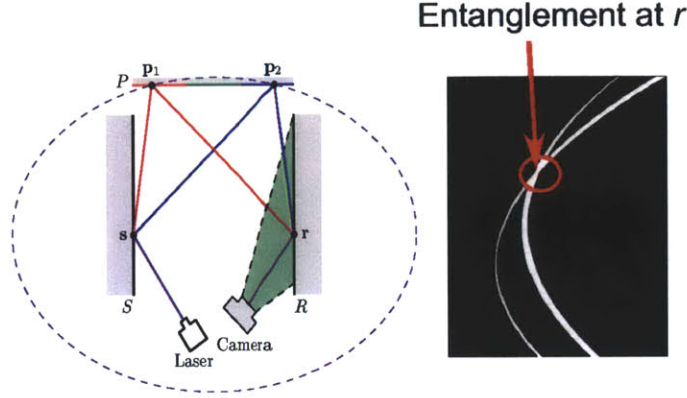
Figure 3-8: Multiplexed scattering causes entanglements: Two light paths, $\mathbf{s} \to \mathbf{p_1} \to \mathbf{r}$ and $\mathbf{s} \to \mathbf{p_2} \to \mathbf{r}$, with the same path length, will get mixed in the streak image. The dashed ellipse (ellipsoid in 3D) shows all points with this path length.

For a fixed laser position, s, and sensor location, r, at a time t, the allowed values of p all lie on an ellipsoid with focal points s and r, given by the equation $t = \|\mathbf{r} - \mathbf{p}\| + \|\mathbf{s} - \mathbf{p}\|$.

If we fix s and p this equation describes a two sheeted hyperboloid in (r, t)-space:

$$t - \|\mathbf{s} - \mathbf{p}\| = \|\mathbf{r} - \mathbf{p}\| = \sqrt{(x - u)^2 + (y - v)^2 + z(x, y)^2} \tag{3.7}$$

where (u, v) are the two coordinates of r in the plane of the receiver wall. In particular, each point on the hidden surface p will contribute a **hyperboloid** to the image $I_R(u, v, t)$. The hyperboloids will have different shapes, depending on the depth $z(x, y)$, and will be shifted along the t-axis. Smaller depth $z(x, y)$ increases eccentricity and leads to higher curvature at the vertex of the hyperboloid.

3.5.4 Path Separation

The ultrafast camera separates out light with different bounces. 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. Ultrafast imaging permits us to easily separate out these different light transport paths therefore greatly simplifying material acquisition by letting us separate out terms that include only the BRDF we are interested in.
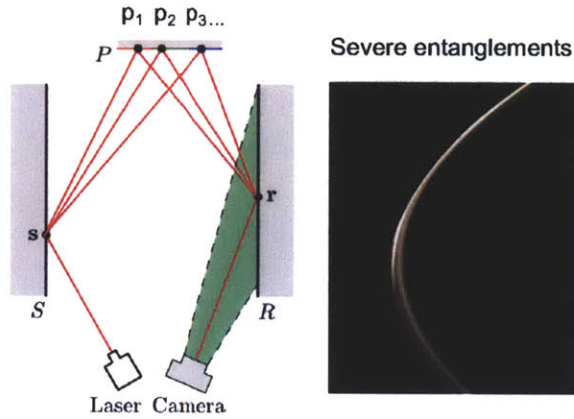
Figure 3-9: The number of entanglements increase dramatically with increase in number of patches. The streak image shows a thick streak which is a combination of streaks from multiple patches.

3.5.5 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 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 3-8 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 3-6-(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 3-6-(a).

3.6 Multiple Patch Measurement

As a patch gets larger we can acquire more measurements, which improves the SNR of our measurements. However, the larger the patch, the greater the entanglement. Thus, there is a trade-off 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 3.8 we show how we can disentangle these observations by fitting to a parametric material model.

3.7   Coverage of Our Measurements

By imaging a Lambertian material R, We indirectly observe BRDF measurements. This process enables simultaneous BRDF measurements, for a given point P, thus accelerating acquisition. This approach is only practicable with the ultrafast camera which can disambiguate between measurements across time, 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 3-6-(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 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 3-6-(b), can be thought of as the setup in Figure 3-6-(a) with the S and R surfaces folded together. So the same discussion of angular coverage applies to this case.

3.8   Reconstructing Reflectance Values

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

3.8.1   Discretizing the Problem Domain

We can estimate the geometry of the scene using our ToF device. A basic method to achieve this is described by Kirmani et al. [52]. 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 $\|\mathbf{r} - \mathbf{p}\|$,
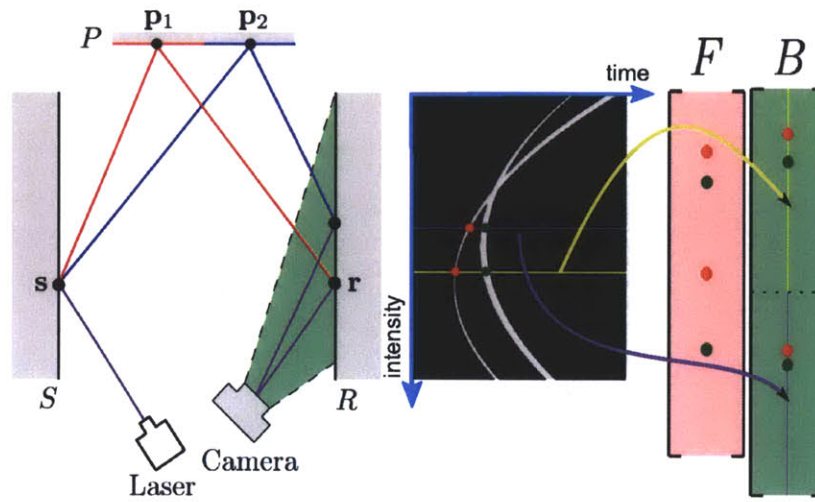
(a)
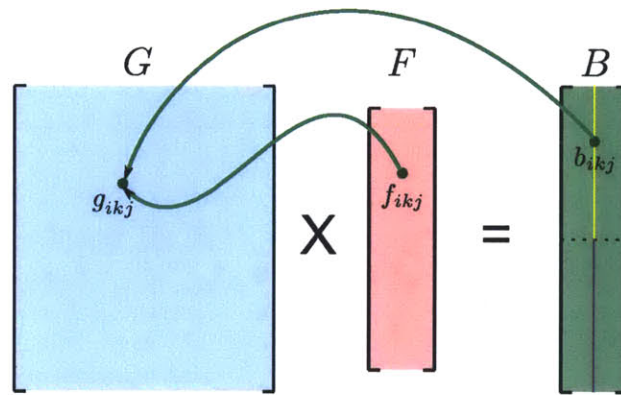
(b)
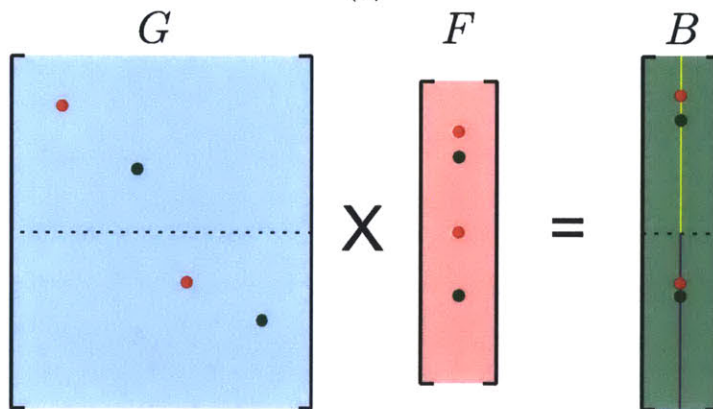
(c)

Figure 3-10: We form the linear system using acquired streak images & geometry information. The location of responses corresponding to unknown patches in the streak image are used to populate G, F and B.

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 3-6-(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 3.6 we get:

$$Q_i(j, m) = A_k \, g(i, k, j) \, F(i, k, j) \tag{3.8}$$

where m is the discretized time taken by light traveling from the laser source to the camera along path $\mathbf{s} \to \mathbf{p} \to \mathbf{r}$, 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. Accordingly, the formulation for both single patch, and multipatch, is the same in the presence of entanglements.

3.8.2 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 3.7. The discretized system of equations

is:

$$B(i, j, m) = G(i, k, j) \cdot F(i, k, j) + \nu \qquad (3.9)$$

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 3-11). 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 T time-slots of observations of the receiver R using the ultrafast camera. 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 ToF 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 3-8. The light from different patches with identical path lengths 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. Therefore the linear system becomes underdetermined. Next, we describe how using reduced dimensional parametric models decrease the number of required measurements, enabling recovery of reflectance parameters.
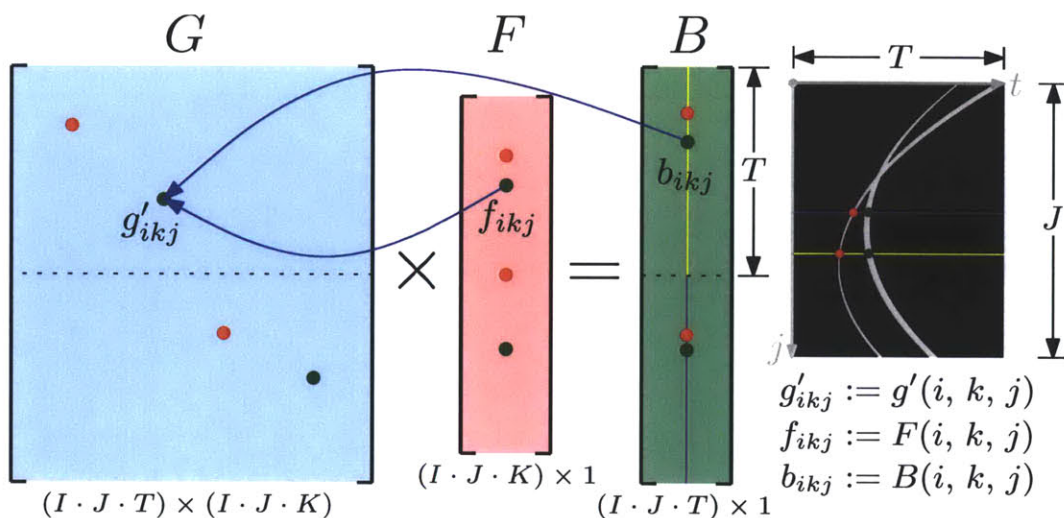
Figure 3-11: The linear system to solve for reflectances is constructed using the streak images and scene geometry information. Consider two patches with a streak image for one source position $\mathbf{s}_i$. The streak image is sampled at two receiver positions $\mathbf{r}_1$ and $\mathbf{r}_2$ to create 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.

3.8.3 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 [53], and use the dBRDF proposed by Ashikhmin et al. [54], and used in [36] to measure distributions of the BRDF. Ashikhmin et al. [54] 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 reflectances 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 [18, 17, 55]. 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. [18]:

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

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 $3K$.

3.8.4 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 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 approximately $T \sim 500$.

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 [56]. 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 4-4 shows the parametric fits obtained for multiple patches.

Chapter 4

Ultrafast Reflectance Acquisition: Results

We now present the results of our ToF reflectance measurement system. We first evaluate our results using a simulation of the device, specifically 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.

4.1 Simulation Results and Validation

In our simulation setup we simulated the working of a ToF camera in the indirect viewing configurations of Figures 3-6-(a) and (b). The results are the same in both configurations, so we only present our results for the (b) configuration.

4.1.1 Simulation Setup

We simulate our physical setup from Figure 4-6. S, P, and R form the left, back, and right planes of a box. S and R are of size 25×25 cm^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 3.8.3).
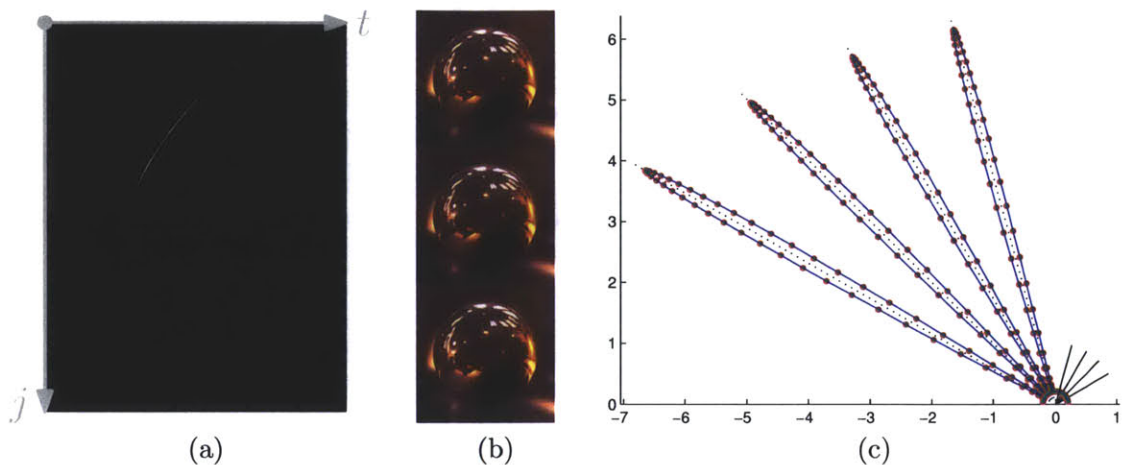
Figure 4-1: In single-patch settings, our approach can achieve near-perfect reconstruction with relatively low noise-levels. Here, the BRDF of Copper was picked from Ngan et al.[18]. Column (a) shows streak images; (b) shows rendered spheres under environment lighting using the ground truth (top) and 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%.

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 [57, 58] 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.

46

(a)   (b)   (c)
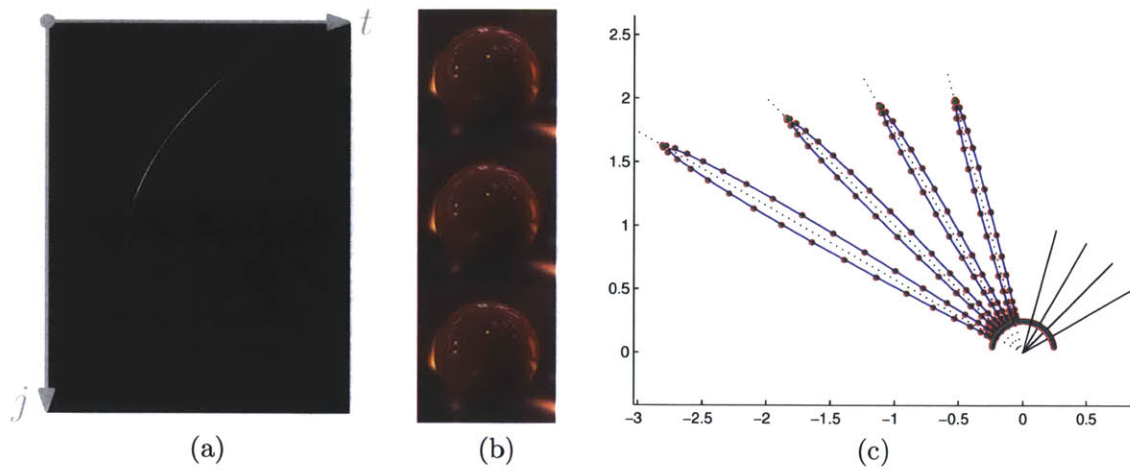
Figure 4-2: Single patch: BRDF of Red Plastic picked from Ngan et al.[18]. Column
(a) shows streak images; (b) shows rendered spheres under environment lighting using the
ground truth (top) and the recovered BRDFs (middle with 1% of noise and bottom with
10%), respectively; (c) shows plots of the BRDFs 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%.

(a)   (b)

(c)   (d)
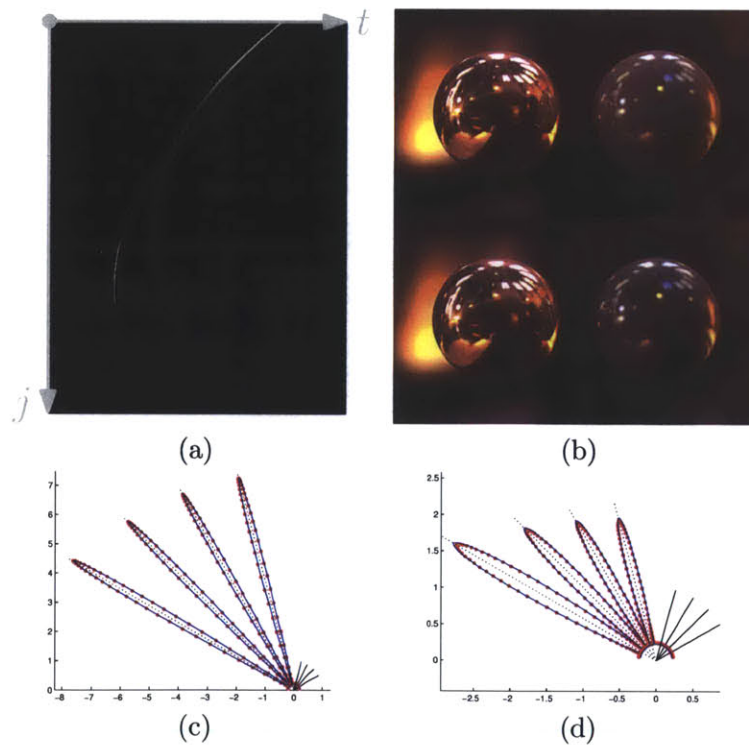
Figure 4-3: Synthetic multiple-patch results using copper and pink jasper. Column (a)
shows the streak image; (b) shows rendered spheres under environment lighting using the
ground truth (top) and the recovered BRDF (bottom), respectively; (c) and (d) show plots
of the BRDFs. The reconstruction was performed with 5% noise added to the input.
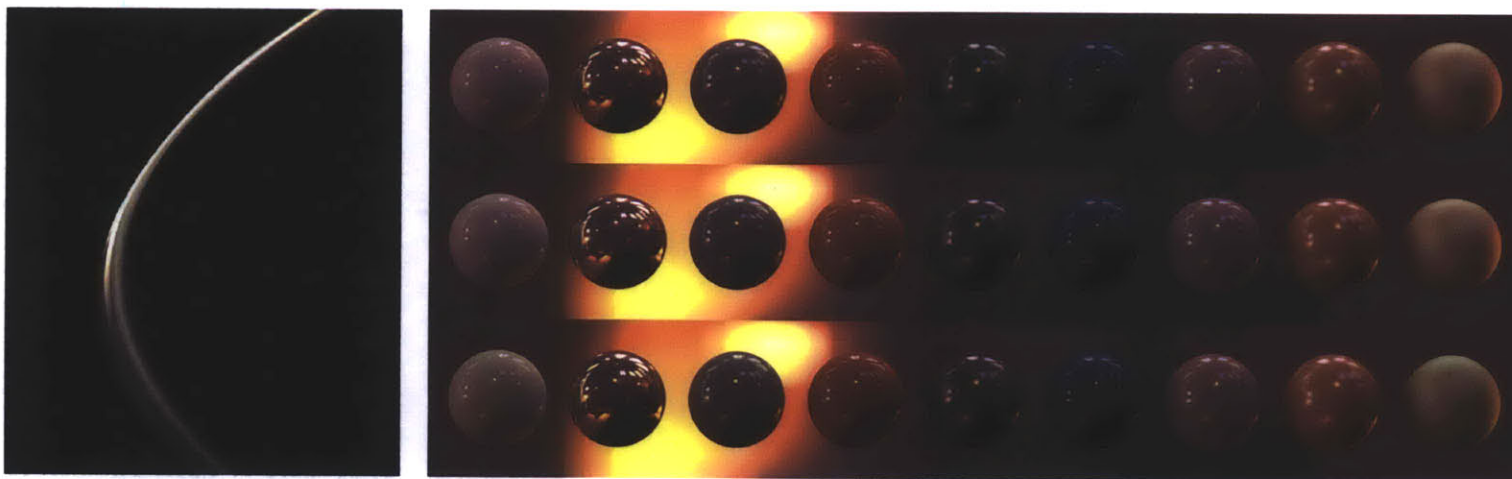
Figure 4-4: With nine patches with different BRDFs picked from Ngan et al.[18], our approach is able to obtain high-quality reconstructions: (left) the streak image; (right) rendered spheres under environment lighting using the ground truth (top) and the recovered BRDFs (middle with 1% of noise and bottom with 10%), respectively. See the supplementary material for BRDF parameter values.

Material	Ground Truth			Recovered (5% noise)		
	k_d	k_s	N	k_d	k_s	N
Copper	0.076	1.040	40800	0.072	1.038	40803
	0.041	0.609	40800	0.044	0.609	40798
	0.029	0.266	40800	0.029	0.268	40800
Pink jasper	0.202	0.054	2590	0.140	0.057	2589.2
	0.109	0.062	2590	0.091	0.060	2601.1
	0.060	0.043	2590	0.050	0.044	2613.9

Table 4.1: BRDF parameters for Figure 4-3.

4.1.2 Results

We simulate two scenarios using this setup: single patch and multiple patch. We choose nine materials from the database by Ngan et al [18] 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 **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 3.8.3, for three different materials with different percentage of additive noise: $\eta = 0\%$, 1% and 10%. 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 **P**. Again for this case, we recover the parameters for three different percentages of additive noise: $\eta = 0\%$, 1% and 10%. Figure 4-4 shows rendered spheres using the ground truth, and the recovered BRDFs with $\eta = 1\%$, and $\eta = 10\%$. Figure 4-4-(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.
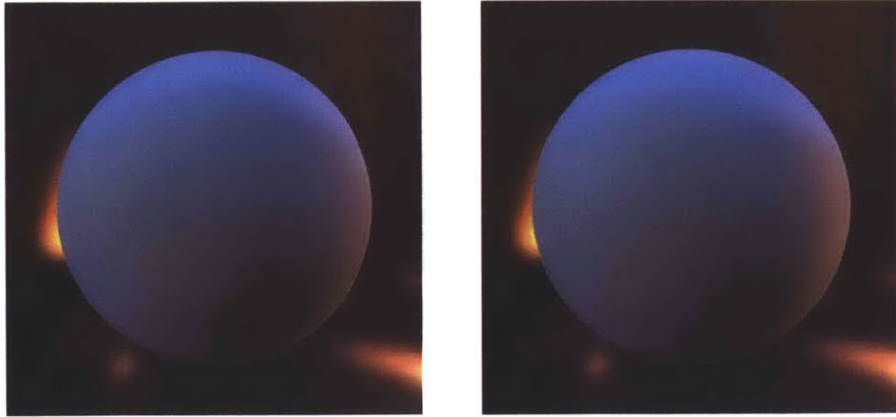
Figure 4-5: Validation using published gonioreflectometer data from the Cornell Reflectance Database: (left) measured data, (right) the recovered BRDF.

4.2 Validation using Published Data

We further evaluate our proposed method for the "around the corner" viewing mode using published gonioreflectometer data from the Cornell Reflectance Database. We employ a simulation setup with the geometry from Figure 3-6-(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 4-5).

4.3 Experimental Setup

Modifying existing ToF (ToF) cameras to report a full time profile with sufficient SNR for our purpose is quite challenging. Our emphasis in this thesis is on the computational and algorithmic aspects. Instead of modifying existing high precision and high speed ToF cameras, we set out to build an ambitious electro-optic hardware that can show a proof-of-concept.

Even when using well understood and established technologies, transitioning from simulated to actual results is by no means trivial. We run into this problem by trying to repurpose the lasers and fast cameras that originally designed for centimeter sized imaging experiments and not for free-space.
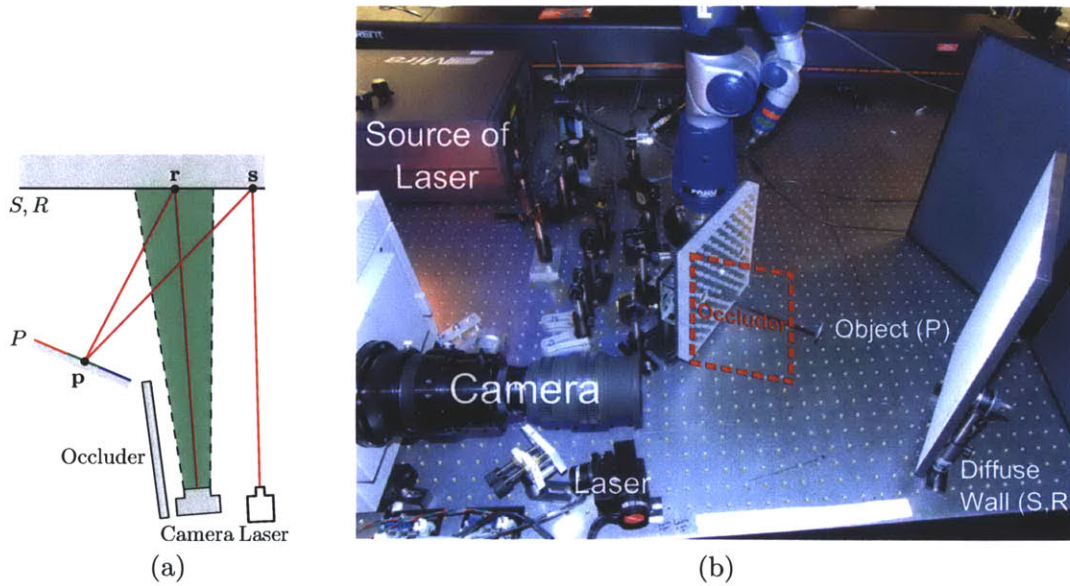
Figure 4-6: The ToF camera setup involves a laser projector via a Ti:Sapphire laser with steering mirrors and a pico-second accurate camera. Reflectance of the object patch P is recovered by aiming the camera and the laser projector at the diffuse wall.

Characterizing in detail the behavior of the laser, camera and synchronization, reliably determining the scene geometry and accounting for day to day or hour to hour drift in crucial parameters like laser intensity is a challenging task. The simulations consider a perfect experiment that exploits the full capabilities of the hardware in every aspect. While this is certainly possible with a considerable amount of research and investment, it is an effort comparable or larger in scale than finding the best possible algorithm for BRDF reconstruction.

Nevertheless we show that an experiment is possible and argue that it will be come much easier in the future both by engineering a system that reaches todays limits and through new scientific breakthroughs in this very active field. Following experimental setup was used for our proof-of-concept implementation.

4.3.1 Details of Experiment

The experimental setups for both the systems described in this thesis consist of a streak camera and an ultrafast pulsed laser. The camera is a Hamamatsu C5680 streak camera that captures one spatial dimension, i.e., a line segment in the scene, with an effective time resolution of 15 ps and a quantum efficiency of about 10%. The position and viewing

direction of the camera are fixed.

The light source is a Kerr lens mode-locked Ti:Sapphire laser. It delivers pulses of about 50 fs length at a repetition rate of 75 MHz. Dispersion in the optical path of the pulse does not stretch the pulse beyond the resolution of the camera of 2 ps and therefore can be neglected. The laser wavelength is centered at 795 nm.

The main laser beam is focused on the diffuser wall with a 1 m focal length lens. The spot created on the wall is about 1 mm in diameter and is scanned across the diffuser wall via a system of two galvanometer actuated mirrors. A small portion of the laser beam is split off with a glass plate and is used to synchronize the laser and streak camera. The diffuser wall serves as both S and R is placed 62 cm from the camera. It is covered with Edmund Optics NT83 diffuse white paint.

For time jitter correction, another portion of the beam is split off, attenuated and directed at the wall as the calibration spot. The calibration spot is in the direct field of view of the camera and serves as a time and intensity reference to compensate for drifts in the synchronization between laser and camera as well as changes in laser output power.

The calibration spot also helps in detecting occasional shifts in the laser direction due to, for example, beam pointing instabilities in the laser. If a positional shift is detected, the data is discarded and the system is re-calibrated. The streak camera's photocathode tube, much like an oscilloscope, has time decayed burn-out and local gain variations. We use a reference background photo to divide and compensate.

Right now, the equipment is quite costly. The laser and camera each cost about $\$150000$ each and triggering plus synchronization electronics, and rest of the electronic and optical hardware brings the cost to over $\$600000$. However it will not be too optimistic to say that the ultrafast lasers and detectors will be more practical in coming decades. The following section gives a brief overview of the history and future scope for progress in this area.

4.3.2   Results from Experiements

Single patch data: For a single material, we image a small material patch (of size 1.5 × 1.5 cm^2) using our acquisition setup. Figure 4-7-(b) shows the rendered spheres using an environment map with the recovered parameters for the two materials copper and plastic, with the error plots (c) and (d). 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 principle RGB
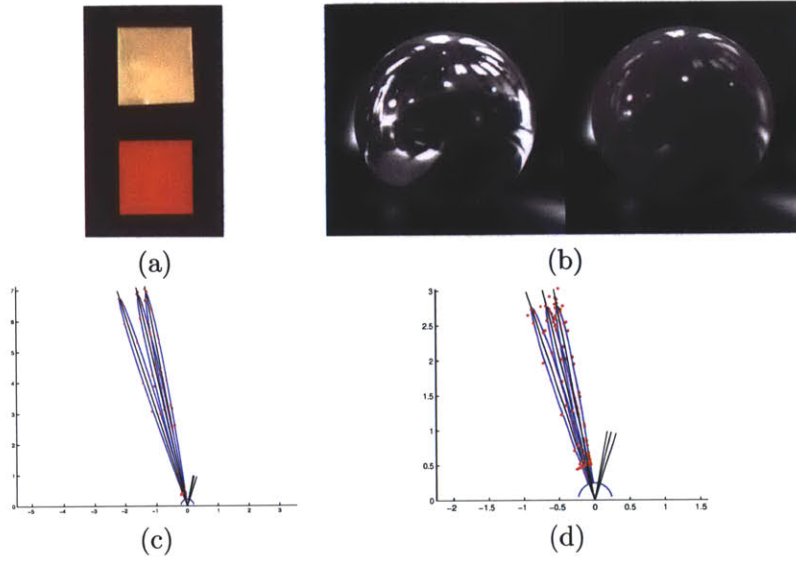
Figure 4-7: For data using our experimental device our method obtains good results. Top row: single patch; Bottom row: two patches. (a)Actual patches of copper and plastic used for acquisition; (b) Spheres using the recovered BRDFs (copper – top, plastic – bottom); (c) and (d) Plots for the BRDFs: red dots indicate measured data points and blue curves are the recovered BRDFs.

color channels could be realized, for example by using an optical parametric oscillator [59] 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. [18].

Multi patch data: We place two patches of copper and red plastic side-by-side and measure them using our experimental device. Figure 4-8-(a) shows the streak image for both simultaneously. We recover the parametric BRDFs for them simultaneously, shown as rendered spheres in Figure 4-8-(b), with error plots in (c) and (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. [18].

(a)　　　　　　　　　(b)

(c)　　　　　　　　　(d)
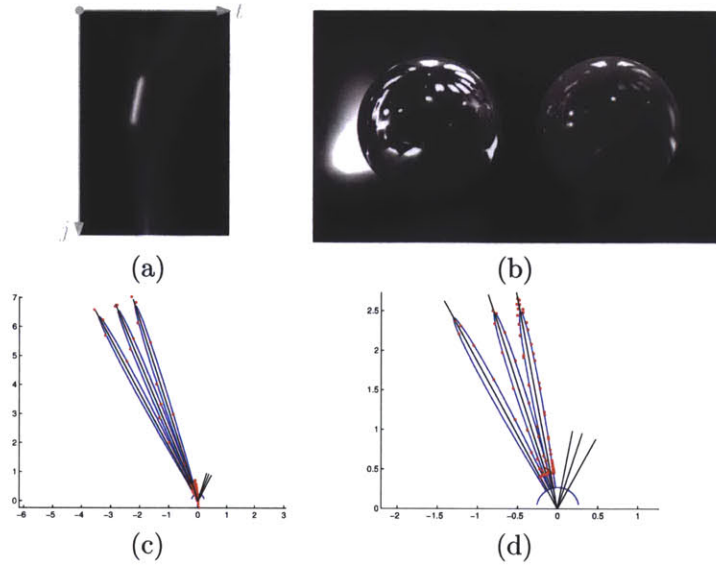
Figure 4-8: For data using our experimental device our method obtains good results. Two patch results. (a) Streak image for both materials together. (b) Spheres rendered using recovered BRDFs (copper on left, plastic on right). (c) and (d): Error plots for the BRDFs.

4.4 Discussion and Limitations

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

Our single-view point approach exploits the temporal dimension, but it introduces traditional problems of ToF cameras in space-time resolution, signal to noise ratio and dynamic range. We summarize these points of discussion below.

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.

Capture dimensionality: Currently our camera only captures one spatial dimension, however, we can scan the laser in two dimensions over the scene to increase sampling range of angles, and further sweep the camera.

The fact that the captured data is only one dimensional reduces the diversity of our collected data but is no serious limitation. The direction missing from the camera can be probed by the laser that can be scanned in 2 dimensions across the scene.

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 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.

Time Resolution: A more challenging problem is timing jitter. Electronic circuits do not operate on timescales shorter than several hundred picoseconds. While the concept of the streak camera circumvents this problem almost entirely, it still requires an electronic signal to synchronize the laser. Timing instabilities in this signal limit our time resolution to a few picoseconds. There is also a trade off between the blur in time and the amount of light collected imposed by the slit in the streak camera.

Dynamic range: When collecting specular materials the brightness of the specular peak can hinder the detection of the remaining BRDF due to limited dynamic range. The diffuse light reflected off a mirror for example will be overpowered by the specular peak of the mirror unless a laser and camera position is used where the specular peak is outside the captured image. Our approach also has a limited resolution for small features with varying BRDF or surface normals.

As a consequence, 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 ToF 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. The average visible light power from a 100 W light bulb arriving at a 1×1 mm^2 camera pixel, at a distance of three meters, is only about 2 photons per picosecond. Therefore, 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 [60, 61].

Intensity Limiting: Higher output power generally increases the intensity of the light striking the scene and can cause safety issues and in extreme cases damage the scene material. The intensity can be lowered by increasing the laser beam diameter and with coded illumination techniques.

Portability: While ultrafast 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 speckle 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 ultrafast optics with emerging topics in computer graphics. Our computational approach has been validated, but our

56

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.

Chapter 5

Imaging Through Scattering Media

Imaging through scattering media is a problem of great interest to study various natural phenomena at a variety of scales from microscopic imaging of biological tissues to astronomical imaging through many miles of atmosphere. It also has interesting applications in machine vision, including looking through fog and smoke and even trying to image what lies behind a curtain.

We develop a method for reconstructing albedo and depth of planar surfaces imaged through a scattering medium using ultrafast imaging. We image the light scattered off this surfaces through this medium and use multibounce light transport analysis to solve for the scene parameters. We specifically use a diffuser as the scattering medium and analyse the reflection mode scattering through the diffuser.

5.1 Geometry of Acquisition

Consider a scene shown in Figure 5-1. The femtosecond laser (L) emits a short light pulse and is pointed towards the diffuser (D) to form a laser spot D_L. Light transmitted through the diffuser which is incident on a scene point W gets reflected and a part of it comes back to the diffuser. A streak camera focused on the diffuser captures the time varying image on the wall at a time resolution of $2\,ps$. The streak camera has one spatial dimension. For each laser position D_L, a $2D$ streak image (1 spatial and 1 temporal dimension) is recorded. A generic point on the plane of focus of camera on the diffuser is denoted by d. We pick out one of these points, denoted by D_C to demonstrate the light transport in the scene in Figure 5-2.

58
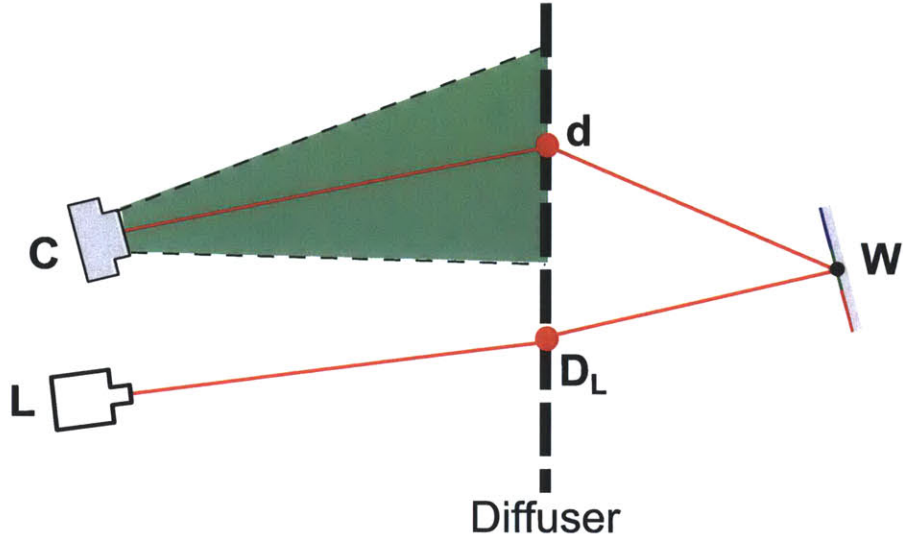
Figure 5-1: Our imaging setup consists of a femtosecond pulsed laser (L), a streak camera with picosecond resolution (C) and a diffuser. The scene(W) is illuminated by a laser pulse on the diffuser and reflected light off the diffuser is captured by the streak camera.

The camera captures the light reflected along the path: $D_L \rightarrow W \rightarrow D_C$. As shown in Figure 5-4, given a point W with normal $\mathbf{1}_W$, γ is the angle made by D_L at W with respect to $\mathbf{1}_W$, namely $\cos\gamma = \mathbf{1}_W \cdot \frac{(D_L - W)}{\|(D_L - W)\|}$. Similarly, we define α for each D_C as the angle made by W at D_C with respect to $\mathbf{1}_D$. We also define β for each D_C as the angle made by D_C at W with respect to $\mathbf{1}_W$

For a single infinitesimally small scene point W, the reflected spherical wavefront propagating from the patch reaches different points on the wall at different times creating a hyperbolic curve in the streak image. We now describe the mathematical formulation of the imaging process.

5.2 Streak Image Formation

5.2.1 Modeling the Diffuser

In our experimental setup, we use a ground-glass diffuser with a Gaussian Diffusion Function. The Gaussian Diffusion Function has a $\mu = 0.0$ and $\sigma = 8.8$ both in degrees. The diffusion function was measured using a calibration surface which is perfectly diffuse and at a known distance behind the diffuser. Figure 5-3 shows the actual Gaussian diffuser Profile which is included in our model. It is represented using the following equation -
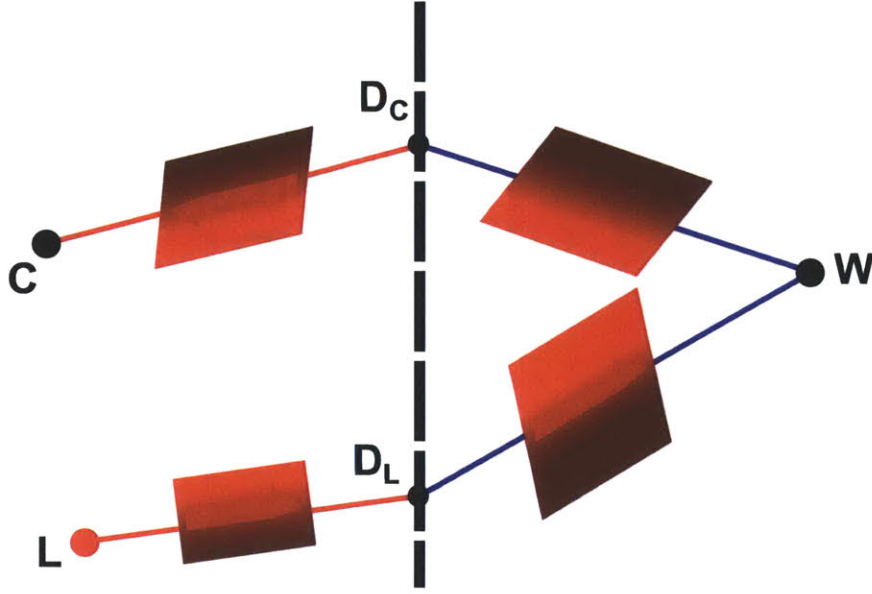
Figure 5-2: The light propagation in this imaging setup follows the path $L \to D_L \to W \to D_C \to C$

$$N(x) = \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) \qquad (5.1)$$

The Gaussian profile is used to model the power transmitted from the diffuser to the scene i.e. along the ray from D_L to W in Figure 5-4. The power transmitted along the ray $D_L \to W$ is a function $N(\phi)$ of the angle ϕ made by $\overrightarrow{D_L W}$ with $\overrightarrow{LD_L}$.

Similarly, the Gaussian profile is also used to model the intensity collected by the camera C focused on the diffuser plane. For example, the power transmitted along the ray from D_C to C in Figure 5-4 is a function a function $N(\theta)$ of the angle θ made by $\overrightarrow{CD_C}$ with $\overrightarrow{WD_C}$.

5.2.2 Bounce Reduction

The process of bounce reduction is similar to that described in Section 3.5.1. If the camera records the streak image $I_C(p, t)$, where p is the pixel coordinate and t is time. In I_C, $t = 0$ corresponds to the instant the laser pulse is emitted from L. Then I_C is related to the intensity $I_D(d, t)$ of light incident on the diffuser plane by the transformation

$$I_D(d, t) = I_C(H(d), t - ||D_L - L|| - ||C - d||). \qquad (5.2)$$
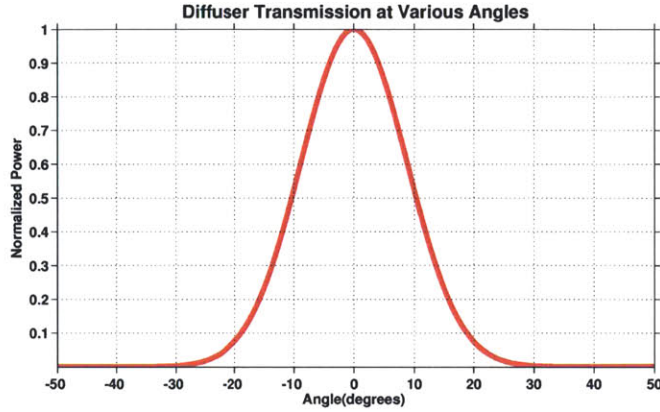
Figure 5-3: Gaussian diffuser profile for the ground glass diffuser used in experiments

In summary, Equation 5.2 reduces the problem to a single scattering problem, with an unfocused point source at D_L emitting a pulse at $t = 0$ and an unfocused virtual array of receivers on D recording the intensity of the reflected wavefront, $I_D(d, t)$.

5.2.3 Image Formation Model
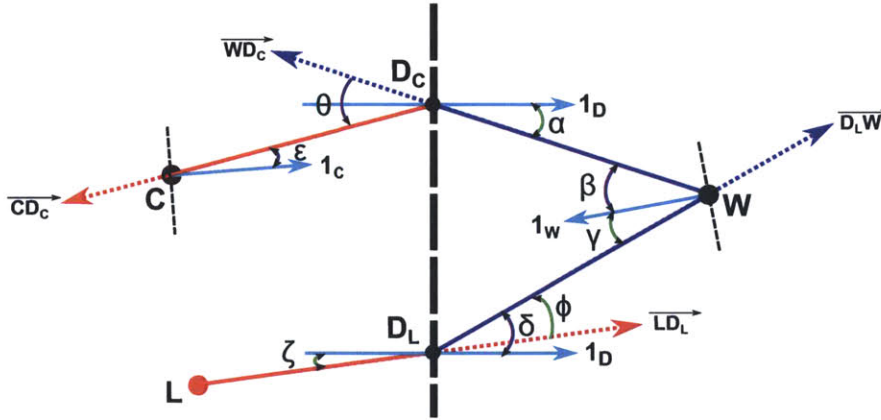
Figure 5-4: We model the streak image formation using different angles and pathlengths in the scene. The diffuser has a Gaussian intensity response.

The homography correction from Equation 5.2 allows us to consider a simplified scenario. The diffuser D is illuminated by the laser at D_L. It can further be assumed to host a virtual array of ultrafast photo-detectors. The virtual photo-detectors create an image $I_D(d, t)$ which is an intensity pattern of the incoming light as a function of time t, and position d.

Hence the image, $I_D(d, t)$, is the intensity observed at $d \in D$ at time t. We introduce an

additional term $N(\theta)$ in the image model to account for the fact that the intensity measured by the camera at a point d is multiplied by a factor which depends on the Gaussian diffuser Profile, as explained in 5.2.1.

Ignoring time, a regular camera viewing R would measure:

$$\frac{I_0\, N(\theta)\, N(\phi)}{\pi} f_a(W) \frac{\cos\gamma\cos\beta\cos\alpha}{\|D_L - W\|^2\, \|D_C - W\|^2} \tag{5.3}$$

where $f_a(W)$ is the albedo of W, I_0 is the laser intensity, and $N(\theta)$ and $N(\phi)$ are the Gaussian diffuser Profile values along the paths $D_C \to C$ and $D_L \to W$ respectively for a given light path $L \to D_L \to W \to D_C \to C$. We introduce the geometry term $g(L, D_L, W, D_C, C)$, which entirely depends on the geometry of the scene.

$$g(L, D_L, W, D_C, C) = \frac{I_0\, N(\theta)\, N(\phi)}{\pi} f_a(W) \frac{\cos\gamma\cos\beta\cos\alpha}{\|D_L - W\|^2\, \|D_C - W\|^2} \tag{5.4}$$

Given a scene S with finite area, a regular camera imaging the patch at location D_C would capture:

$$\int_{W \in S} g(L, D_L, W, D_C, C) f_a(W)\, \mathrm{d}W \tag{5.5}$$

Given the laser illumination D_L, a ToF camera focused at the point D_C at a time t, measures the following (assuming no inter-reflections):

$$I_R(D_C, t) = \int_{W \in S'} g(L, D_L, W, D_C, C) f_a(W)\, \mathrm{d}W \tag{5.6}$$

where, $S' \subseteq S$ consists of all points W such that the path length x for the laser light to arrive at D_C is the same. That is, at a given instant t light arrives at D_C from all points W which have an equal path length along $D_L \to W \to D_C$:

$$x = \|D_L - W\| + \|D_C - W\| = c \cdot t$$

where c is the speed of light.

5.3   Reconstructing Scene Parameters

In this section we describe how we recover albedo and geometry of the unknown scene using the acquired data.

5.3.1 Discretizing the Problem Domain

We start by discretizing the problem domain by dividing the 3D Cartesian space behind the diffuser into a voxel grid of a fixed resolution. The voxels which have a non-zero albedo and a finite size will contribute to the formation of a streak image when light is reflected off them towards the diffuser being imaged by the streak camera.

Given a number of acquired streak images using different laser positions, the idea is to search through this voxel grid to identify voxel locations and albedo which can reconstruct these images. We use the forward model described in Section 5.2.3 to generate streak images using different voxel locations and albedos. We set up an optimization problem to reconstruct a scene such that an error norm between the streak images generated from this reconstructed scene and the acquired streak images is minimized. We now describe the discrete problem formulation in detail.

5.3.2 Linear System Formulation

Let us divide the voxel grid into K discrete voxels. The value of $f_a(k)$ is set to zero for all the voxels that do not contain a physical object. The value of $f_a(k)$ for voxels that lie on the surface of the objects is the albedo of the corresponding surface point. Voxels inside the surface of a physical object are occluded by the voxels on the surface. Hence the voxels are set to zero as well.

Now consider the streak image I_R recorded with the laser at position L_1. Vectorize the streak image pixels into a single vector $b_{R,1} \in \mathbb{R}^M$, where M is the total number of spatio-temporal pixels present. The pixel values will depend linearly on the albedos in f_A and hence satisfy a linear equation of the form

$$b_{R,1} = G_1 \cdot f_A \tag{5.7}$$

The matrix G_1 is created using individual entries which are derived from the values of $g(L_1, D_{L_1}, k, D_C, C)$ from Equation 5.4, where k is an individual voxel. The formation of this matrix is exactly the same as the matrix formulation in the BRDF acquisition problem described in Section 3.8.2.

If multiple streak images $1, \ldots, n$ are recorded corresponding to different locations of the laser, then those different streak images are stacked on top of each other in a vector b,

63

which satisfies the linear equation

$$b = \begin{bmatrix} b_{R,1} \\ b_{R,2} \\ \vdots \\ b_{R,n} \end{bmatrix} = \begin{bmatrix} G_1 \\ G_2 \\ \vdots \\ G_n \end{bmatrix} f_A = G \cdot f_A \tag{5.8}$$

5.3.3   Optimization for Scene Reconstruction

Now we describe an optimization problem which recovers the scene geometry and albedo using the linear system formulation described above. Let us consider a dataset of n acquired images from an unknown scene. A streak image vector consisting of all these images is denoted by b_r. The goal is to solve for f_A, which is then rearranged to reconstruct the scene.

Using Scene Priors: In order to solve for an unknown scene, we use some scene priors which help to impose limits on the search space. Specifically, we use the acquired streak images to constrain the depth of the voxel grid along the z-plane. We use the fact that given fixed x and y constraints on the voxel grid i.e. the height and width of the diffuser and known locations of camera, laser and diffuser, it is possible to determine the minimum and maximum depth of the unknown scene.

For this purpose, we pick out the first and last onset in time, which are first and last pixels with non-zero intensity along the time-axis. The total path-length required for a scene point to contribute to the first offset is used to derive the depth at which the scene point must exist. This value corresponds to the minimum depth of the scene. Similarly the last offset is used to compute the maximum depth of the scene. This process helps us to constrain the search-space and speed-up the optimization process.

The Optimization Problem: We start the optimization by a random initialization of f_A and generate the streak image vector described by Equation 5.8. We denote this simulated vector by b_s. Now we setup an optimization problem as follows -

$$\underset{f_A}{\arg\min} \ ||b_r - b_s|| \, , \text{ for } 0 \leq f_A \leq 1 \tag{5.9}$$

We solve this problem using constrained minimization over the voxel grid defined using geometric constraints derived from scene priors.

5.4 Real Data Calibration

In this section, we describe how we calibrate our physical model for this setup to recover unknown albedo and geometry of real-world scenes. This calibration is necessary due to many physical factors such as laser jitter, camera vignetting and imperfections in the diffuser design are unaccounted for in our physical model.

In order to recover the albedo and geometry of an unknown scene from streak images we setup a linear system as explained in Section 5.3.2. When using this linear system approach with a real dataset, we perform following calibration steps. These steps are in addition to the optimization procedure used to solve for the scene parameters -

1. **Initial streak simulation:** We assume that the geometry of the scene is known, though this assumption is dropped when solving the optimization in Section 5.3.3. Using Equation 5.6, we render the streak image for each laser position. We set the factors $N(\theta)$ and $N(\phi)$ to 1 for the moment. Figure 5-5-(a) shows a streak image rendered at this step for a single white patch of $1\ cm^2$ area.

2. **Correction for sensor response:** An ideal camera sensor has a perfectly linear response to the light incident on it. However the streak camera used in our experiments has a non-linear response to the incident light due to aging. We record this sensor response by capturing an image of a white lightbox. The camera response is irregular with a sharp fall-off along the edges as it can be seen in Figure 5-5-(b). This camera response is fairly stable across time. We capture a reference image every time a new dataset is collected in order to make sure that an accurate reference image is used.

   In order to match the rendered streaks with the real data, we now multiply the rendered streak image by a correction image (Figure 5-5-(b)), which is a smoothened version of the original sensor response along space-dimension. We denote the vectorized version of this image by $I_{correct}$. Figure 5-5-(c) shows the multiplication of the original rendered image and the correction image. We denote the vectorized rendered streak images at this stage by b_s.
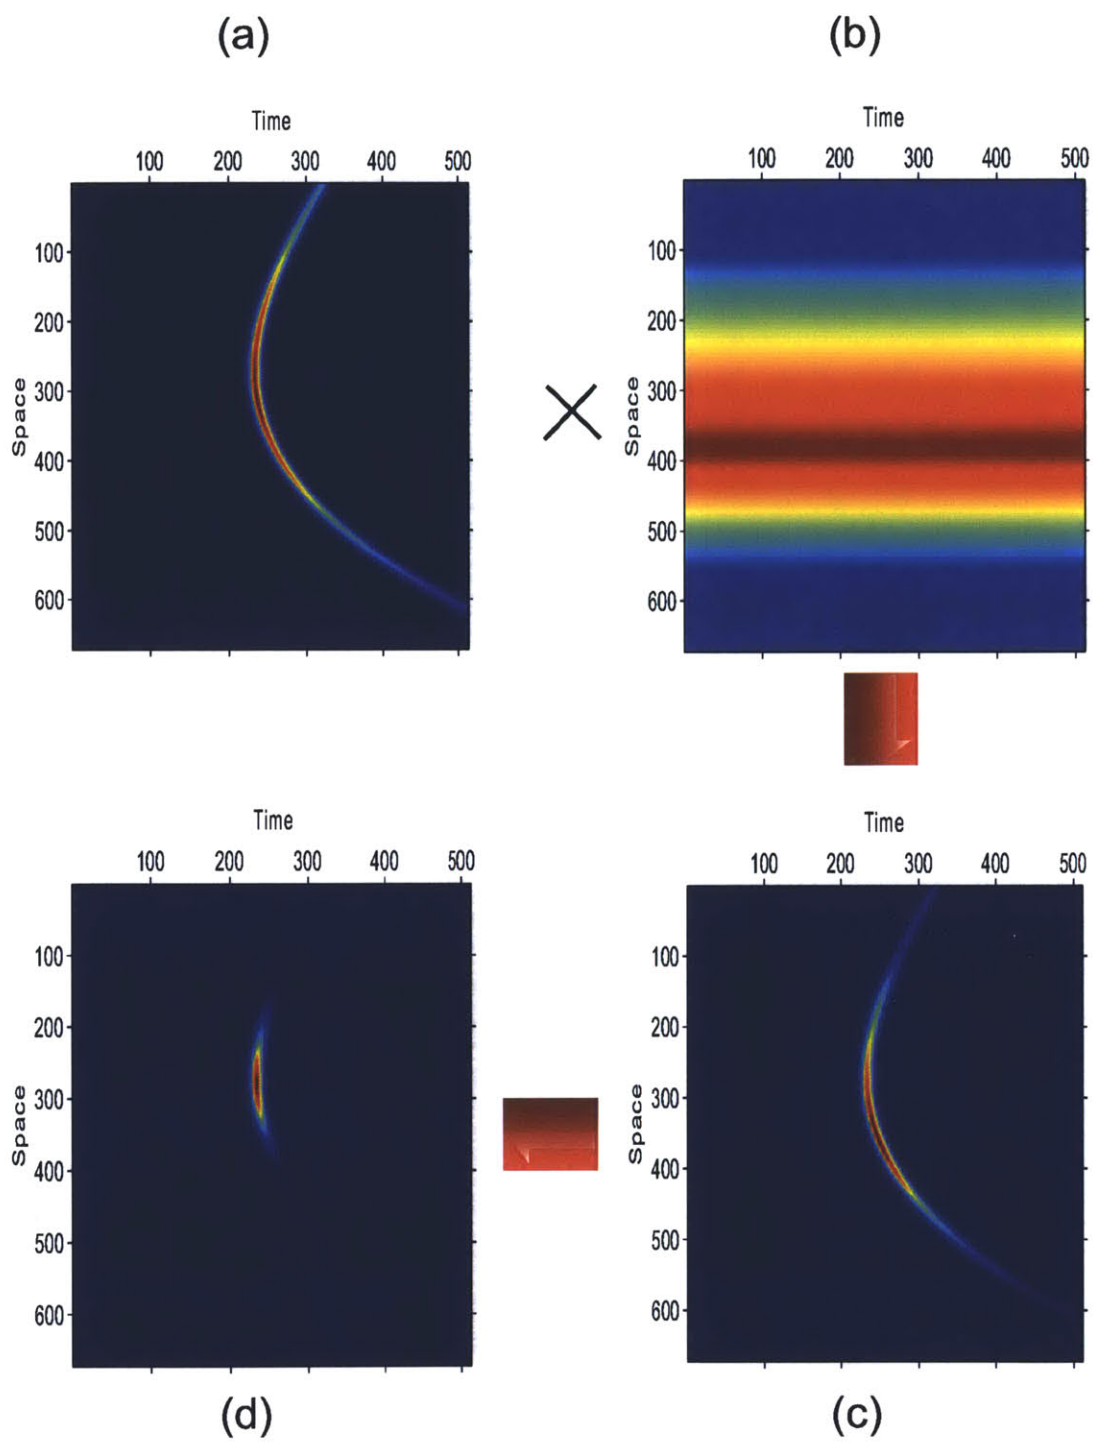
Figure 5-5: Real data calibration process: The rendered image (a) is multiplied by a correction image (b) which compensates for non-linear sensor response. The resultant image (c) is multiplied by the Gaussian diffuser profile obtained using optimization to obtain the final rendered image (d).
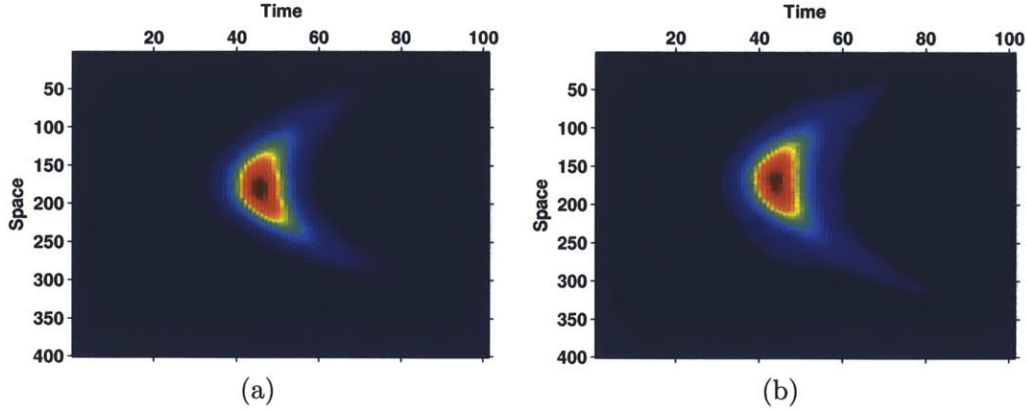
Figure 5-6: The a calibrated simulated image (a) for a small white patch is compared with the real streak image (b) .The two images look fairly identical except for experimental noise. The error norm $\zeta = 0.1012$

3. **Solving for the diffuser profile:** We use a ground glass diffuser which has a Gaussian intensity profile as explained in Section 5.2.1. We set an optimization problem which solves for μ and σ of a Gaussian diffuser Profile (Equation 5.1). The optimization tries to minimize $||b_s - b_r||$, where b_s and b_r are vectorized simulated and real streaks respectively. The values of μ and σ remain extremely stable (within 0.5% variation) for our diffuser across datasets. Figure 5-5-(d) shows a streak image rendered at this step for a single white patch of $1\ cm^2$ area.

4. **Matching intensity profiles:** After multiplying the rendered image with an appropriate diffuser profile, we match the intensities of the rendered and real streak images. It is hard to find the actual laser intensity I_0 as mentioned in Section 5.6, because of the laser jitter. Therefore we instead find a scaling factor κ which includes laser intensity and any other global scaling factors introduced in the physical setup.

$$\kappa = ||b_s - b_r||$$

We now generate the final calibrated rendered streak image $\kappa \cdot b_s$. Figure 5-6-(a) shows the final rendered streak image calibrated using real data. Figure 5-6-(b) is the corresponding real streak image. It can be clearly seen that these two images match visually. An error norm ζ between the two images is calculated by using the following equation -

$$\zeta = \frac{||b_r - b_s||}{||b_r||} \tag{5.10}$$

Inclusion in optimization problem: While solving an optimization problem to recover the parameters of a real scene, the aforementioned calibration steps are included in the process as follows. The Gaussian diffuser Profile is derived using a small white calibration patch. Therefore the generated G matrix includes these factors. After rendering streak images, the vector b_s is multiplied by a vectorized version of $I_{correct}$ to compensate for the non-linear sensor profile. After this step, the final optimization is performed where the minimization problem is modified to include the global scaling factor κ as follows-

$$\underset{f_A}{\arg\min} \; ||b_r - \kappa \cdot b_s||, \; \text{for } 0 \leq f_A \leq 1 \tag{5.11}$$

Chapter 6

Imaging Through Scattering Media: Results

We now present the results of imaging through scattering media using ultrafast imaging. 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 this section we describe our simulation setup and different scenarios used to verify our approach. We present and analyze results obtained using this simulation setup.

6.1.1 Simulation Setup

We simulate our physical setup from Figure 5-1. D is a diffuser which is illuminated on the left at various positions by a pulsed femtosecond laser source L. The light is reflected off a scene placed to the right of the diffuser. A streak camera images a horizontal line on the diffuser from the left. The camera measures $T = 512$ time slots, where each time slot corresponds to $1.6678\,ps$. The diffuser is assumed to have a Gaussian Diffuser Profile with $\mu = 0$ and $\sigma = 8.8$, both in degrees.

We simulate different measurement scenarios by introducing various scenes. Using the geometry information and the supplied albedos, we generate the streak images. The resolution of voxel grid is $1\,\mu m \times 1\,\mu m$. In order to simulate a finite area of the scene point, we

Noise Level	Recovered Depth (cm)
0%	23.9976
1%	24.0092
5%	23.9845

Table 6.1: Mean recovered depths for different noise levels for the scene from Figure 6-1. The ground truth depth is 24.00 cm

blur the rendered streak image in both space and time dimensions using a Gaussian Kernel obtained from calibration data from a small white patch.

Our simulation provides us with the measurement vector b_s. 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. We use the same noise model as used in the reflectance capture process. Please refer to Section 4.1.1 for details about the noise model.

6.1.2 Results

Recovering a single plane at fixed unknown depth: We reconstruct an image of a planar object at a fixed unknown distance behind the diffuser. The object consists of a 2D grid of 100 × 100 patches. It is placed at 24 cm behind the diffuser. We use 60 streak images from 60 different laser positions in total. We add Gaussian noise to the rendered streak images using the previously described noise model.

In order to visualize the performance better, the object is constructed using a picture. We repeat the experiment in three channels using three different frequencies of lasers. We run the optimization process with random initialization in a voxel grid around the object. We perform reconstructions in three channels to obtain an RGB image. Figure 6-1 show reconstructions obtained from images with no noise, 1% noise and 5% noise respectively. The recovered depths are summarized in Table 6.1.2.

Recovering two planes at unknown depths: We reconstruct two planar objects placed at two unknown depths behind the diffuser. The objects are two 2D grids of 100 × 100 patches. They are placed at 13 cm and 13.2 cm behind the diffuser respectively. We use 60 streak images from 60 different laser positions in total. We add Gaussian noise to the rendered streak images using the previously described noise model.

We run the optimization process with random initialization in a voxel grid around the object. Figure 6-2 shows albedo and depth reconstructions obtained from streak images

Original Image

(R,G,B)	R	G	B

Reconstructed Images

η=0.0

η=0.01

η=0.05

Figure 6-1: We reconstruct an image of a planar object at a fixed unknown distance behind the diffuser. The object consists of a 2D grid of 100×100 patches. We reconstruct the scene from streak images with additive Gaussian noise. The figures show reconstructions obtained from images with no noise, 1% noise and 5% noise respectively. We perform reconstructions in three channels to obtain an RGB image.

| (a) | (b) |

Figure 6-2: We recover two planar objects placed at two unknown depths from 60 streak images containing 5% noise in simulation. sub-figure (a) shows the recovered albedo, (b) shows the recovered depth map.
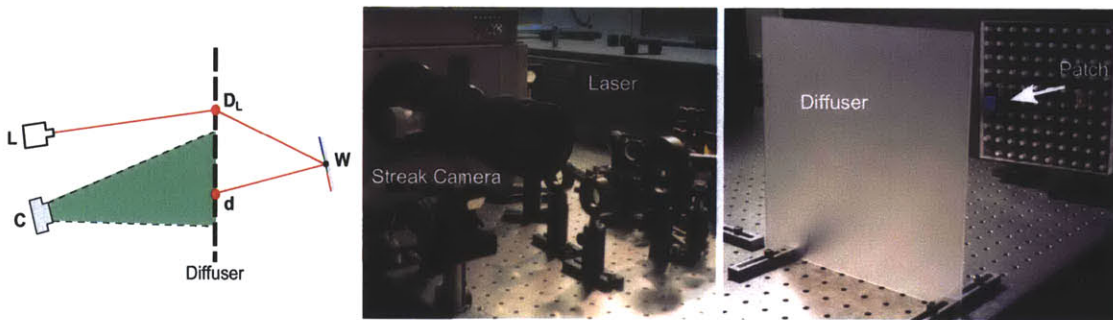

Figure 6-3: The ToF camera setup involves a laser projector via a Ti:Sapphire laser with steering mirrors and a pico-second accurate camera. The scene parameters of an unknown scene patch W are recovered from the images acquired by the streak camera.

with 5% noise. The average recovered depths are 13.0124 cm and 13.1945 cm respectively.

6.2 Experimental Setup and Results

As mentioned in Section 4.3, modifying existing streak cameras for large scale experiments with natural scenes is quite challenging. Nonetheless, we repurpose an ultrafast imaging system used for femtochemistry experiments to obtain various proof-of-concept results.

6.2.1 Details of Experiment

The camera and laser assembly used for these experiments is the same as the reflectance capture setup. Please refer to Section 4.3.1 for details about the hardware setup. We now describe the specific arrangements of the setup for imaging through scattering media.

We use a ground glass diffuser with a Gaussian diffuser profile. It is placed at a distance of 35 cm from the camera. The position and viewing direction of the camera are fixed. The

Figure 6-4: Left - A picture of the diffuser with a scene behind it as seen with a naked eye. Our method can successfully reconstruct the scene using ultrafast imaging. Note that the black vertical lines are for calibration purposes. Right - A picture of the diffuser with the laser illuminating its surface.
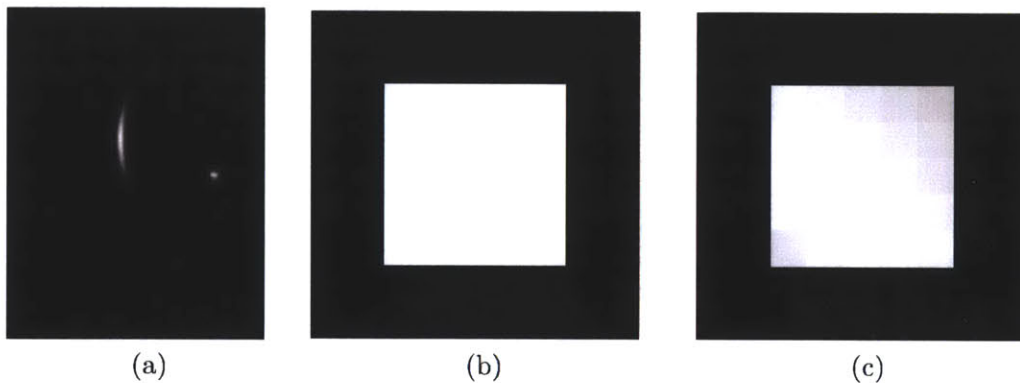

(a)   (b)   (c)

Figure 6-5: In single-patch settings, our approach can achieve near-perfect reconstruction for both albedo and depth. Figure (a) shows an acquired streak image, (b) shows the actual patch and background. Albedo of the patch = 1, Depth = 13 cm. Figure (c) shows the recovered patch with mean albedo = 0.9509 and Mean depth = 12.9 cm.

main laser beam is focused on the diffuser with a $1\,m$ focal length lens. Both the camera and laser are on one side of the diffuser while the scene being imaged is on the other side. The scene being imaged is placed at various depths behind the diffuser. Figure 6-3 shows the experimental setup with a sample scene.

6.2.2   Results

Recovering a single patch at fixed unknown depth: We reconstruct an image of a single white patch of size $2\,cm \times 2\,cm$ at a fixed unknown distance ($13\,cm$) behind the diffuser. We use a voxel grid with each voxel having dimensions of $4\,mm \times 4\,mm \times 1\,mm$. Figure 6-5 shows the recovered patch and background using 16 different laser positions in
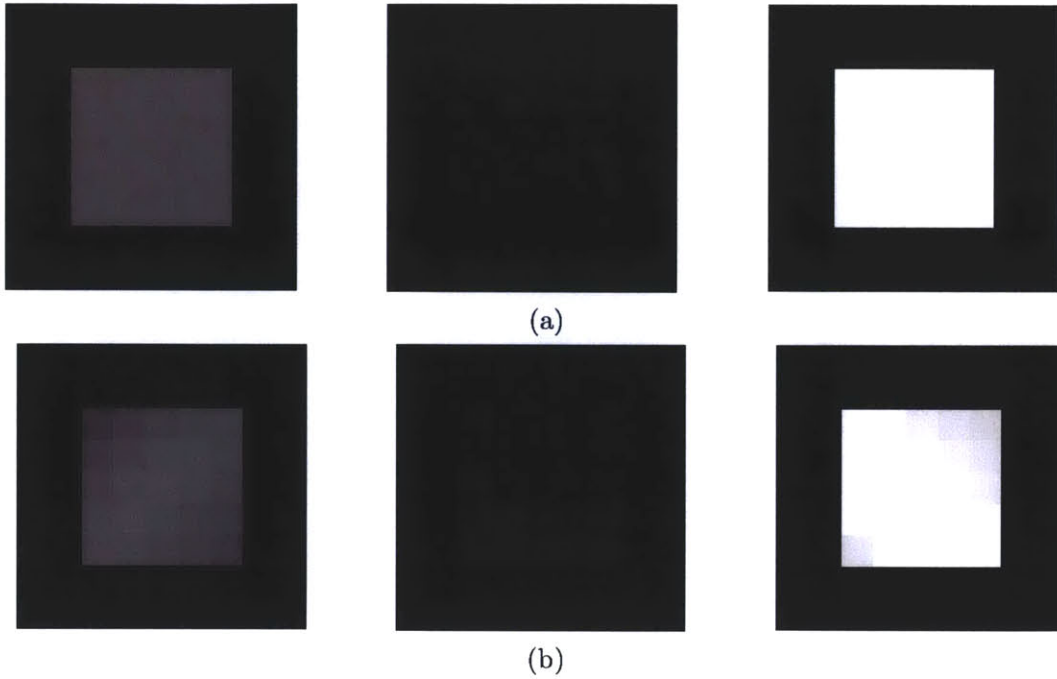
Figure 6-6: We recover three patches at three different depths from real streak images. sub-figure (a) shows the actual patches and background. sub-figure (b) shows the recovered patches using real streak images. The average albedo and depths of recovered patches can be found in Table 6.2

total. The recovered depth is 12.9 cm. The imperfections in recovery are due to intensity variations and temporal jittering of the laser along with vignetting and non-linear temporal distortions in the streak camera.

Recovering multiple patches from different depths We reconstruct an image of three patches of $2\ cm \times 2\ cm$, each having a different albedo, placed at different depths behind the diffuser. We use a voxel grid of resolution $4\ mm \times 4\ mm \times 1\ mm$. Figure 6-6 shows the rendered recovered patches from 16 laser positions. Table 6.2 shows the average

Albedo - True	Albedo - Recovered	Depth - True	Depth - Recovered
0.4000	0.3755	13.5	13.4
0.0500	0.0707	14.0	14.2
1.0000	0.9876	14.5	14.5

Table 6.2: Recovered average albedo values and depth values (in centimeters) for the 'three patch' dataset along with the ground truth. Figure 6-6 shows the rendered scene. Note that the error is higher for the middle patch with albedo 0.0500 due to the fact that SNR is lower for a dark patch, as it reflects less light.

74

depth and albedo recovered along with the ground truth values. As it can be seen, the recovered values are pretty accurate.

6.3 Limitations

There are several limitations of the current work. We assume that the scattering medium has a known scattering profile and transmission model and include this model in our forward model. In reality, scattering media can have more complex profiles which can only be approximately modeled using probabilistic approaches. Furthermore, we use the scattered component of light to only recover planar patches at discrete depth levels. We further assume that the planar patches are parallel to the surface of the diffuser. In reality, many objects of interest can have complex 3D geometry as well as inter-reflections.

In theory, our algorithm can solve for continuous non-planar geometry when no inter-reflections are present. We intend to analyze this case in future work. Moreover, our approach also suffers from traditional problems of ToF cameras in space-time resolution, signal to noise ratio and dynamic range. Please refer to Section 4.4 for a detailed discussion on these limitations.

However, this work proposes a new approach for imaging through scattering media by utilizing the diffuse component of light transport and providing an elegants solution for recovering albedo and depth of objects imaged through scattering media.

Chapter 7

Applications and Conclusion

In this chapter, we describe some potential applications for ultrafast reflectance capture and imaging through scattering media and conclude by summarizing the contributions of the thesis.

7.1 Applications

Our novel approaches for reflectance acquisition and imaging through scattering media can enable many exciting applications. In this section, we explore some of these applications.
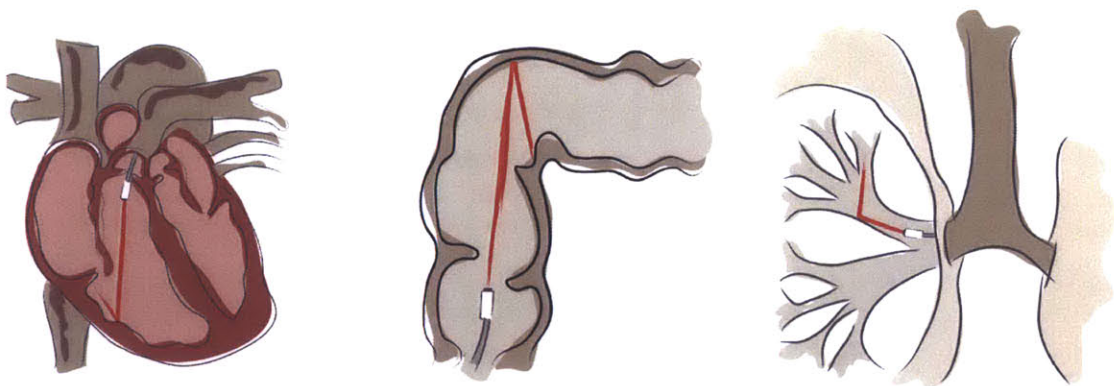
7.1.1 Ultrafast Reflectance Acquisition

Figure 7-1: Ultrafast reflectance acquisition can enable novel applications in medical imaging including recovering malignant growth in endoscopy beyond the reach of the camera. (Sketches by artist Tiago Allen)

We believe that our method for ultrafast reflectance acquisition will inspire uninstru-

mented 'in the wild' acquisition techniques. 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. We can envision novel real-world applications in industrial and commercial photography for industries such as clothing, furniture and automobile manufacturing.

In addition, around the corner recovery of material properties can enable radical applications,especially in medical imaging, e.g., recovering malignant growth in endoscopy beyond the reach of a camera. Our approach also fits in the general spirit 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. Figure 7-1 demonstrates some of these novel application scenarios.

7.1.2 Imaging through Scattering Media

We hope that the theoretical framework developed in this thesis for imaging through scattering media will inspire research to essentially develop an 'ultrasound with light'. It can lead to the development of fast, non-invasive and portable medical imaging tools by better ways of analyzing light transport through these media. This approach can also enable new ways of imaging through non-biological scattering media such as atmosphere and different fluids for applications in computer vision and atmospheric imaging.

7.2 Conclusion

In this thesis, we have developed a theoretical framework for analyzing multibounce light transport using ultrafast imaging. We use this framework for developing solutions for two important problems in computer graphics and medical imaging, namely, reflectance acquisition and imaging through scattering media.

For reflectance capture, we describe a method for fast,simultaneous acquisition of segmented scene reflectances using indirect reflections. We identify 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 by fitting to low dimensional parametric reflectance models. We validate our method using several simulated and real

experiments.

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.

We also describe our novel approach for imaging through scattering media by acquiring albedo and depth information of planar scenes imaged through a diffuser using an ultrafast camera. We describe the linear system approach to analyze indirectly scattered light and set up an optimization problem to recover scene parameters. We validate our method using simulations and real experiments.

In summary, this thesis describes a novel theoretical framework for multibounce light transport analysis and its applications for ultrafast reflectance acquisition and imaging through scattering media. These methods have the potential to spur novel research in computer graphics and medical imaging using advanced optoelectronic devices and new mathematical tools. We hope that they lead to exciting applications in diverse fields including industrial photography, medical imaging and animation.

Appendix A

Ultrafast Imaging: State of the Art

The ultrafast imaging system used for this project is expensive and cumbersome to use. However the limitations of our particular setup do not necessarily reflect the state of the art in advanced optoelectronic devices. In this appendix, we discuss the development of devices and current hardware advances in the industry.

Since their initial conception in 1960, lasers have become rapidly cheaper, smaller and more robust. While the first lasers tended to be large room filling devices or could only operate at small powers and for fractions of a second to avoid overheating, today's diode lasers deliver superior performance with microscopic sizes at the cost of less than a dollar. Wall plug efficiencies of continuously operating laser diodes can be above 60%. Blue lasers, for example where first demonstrated as dye lasers in the early 70s. The first practical blue diode laser was built in 1996 and only seven years later, the first blue laser was available in a commercial product.

Ultrafast lasers have not yet reached this state of maturity. The first Ti:Sapphire lasers were built in the early nineties and the technology matured shortly thereafter to produce femtosecond duration pulses. While the size, price and power consumption have been significantly improved over the past two decades, Ti:Sapphires today still have the size of a large desk top computer, cost about 100,000-200,000 dollars, and have wall plug efficiencies of well below 1%. There is no fundamental reason for this low efficiency and new technologies will continuously improve it.

The basic problem with the ultrafast detection is that ultimately the information is processed in the form of electronic signals which are too slow to represent information

on picosecond or femtosecond timescales. Optical gathering, processing, amplification and storage of the information becomes necessary. Developing methods to do these tasks is one of the primary goals of today's research in computer processors, telecommunication and integrated optics. Significant and rapid progress is being made in these areas and will benefit efforts in transient imaging.

In our experiments, the hardware is standard laboratory equipment available to most well equipped optics laboratories. While both the laser and camera are suitable for a demonstration of our concept they can hardly be called an ideal imaging system for reflectance or free-space.

The wavelength emitted by the laser (795 nm) does not match the wavelength where the camera is most sensitive (about 530 nm). Output power of our laser is $500 - 800\,mW$, less than that of a light bulb and the laser emits pulses with low pulse energy at a high repetition rate. LIDAR applications with similar imaging requirements usually use amplified laser light sources with much lower repetition rates in the kilohertz range with the same or higher average output power than our laser and much higher pulse energies, dramatically improving SNR.

Therefore, even in a present day system, we can expect dramatic improvements in results from a system selected or designed for the applications described in this thesis. Further improvement is expected by future developments in ultrafast optics.

Appendix B

Ultrafast Reflectance Acquisition: Additional Results

In this appendix, we document additional results obtained in simulation for reflectance acquisition.

B.1 Synthetic Single-Patch Results

In this section, we present results obtained for nine materials in a single patch scenario. In each figure, we show renderings of the ground truth, and renderings using recovered Ashikhmin parameters at 1% and 10% additive noise to the streak images. We also show 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%.

B.1.1 Acrylic white
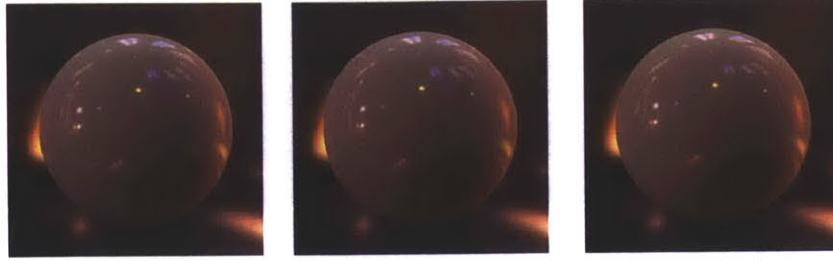
ground truth   recovered (1% noise)   recovered (10% noise)
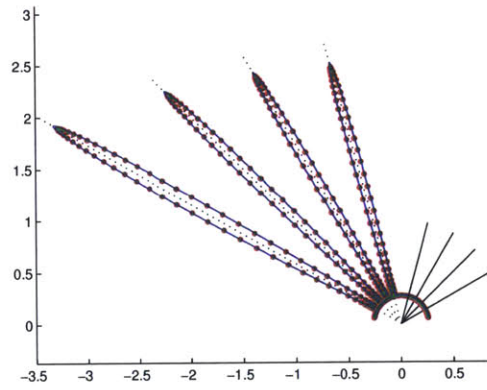
B.1.2 Copper

ground truth   recovered (1% noise)   recovered (10% noise)
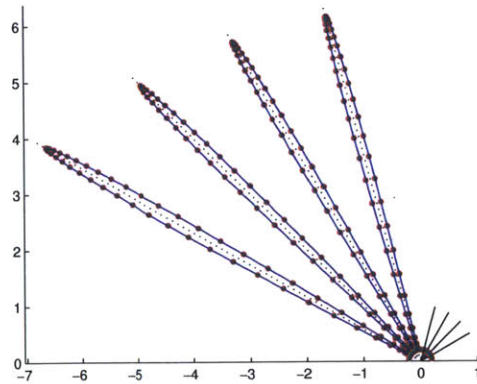
B.1.3 Gray plastic
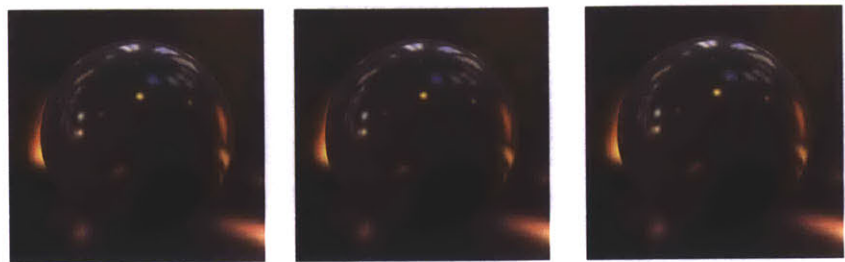
ground truth   recovered (1% noise)   recovered (10% noise)
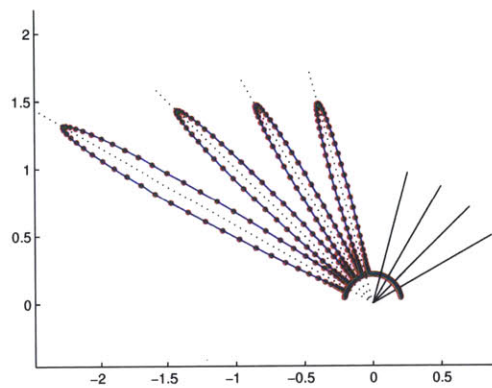
B.1.4 Red plastic specular
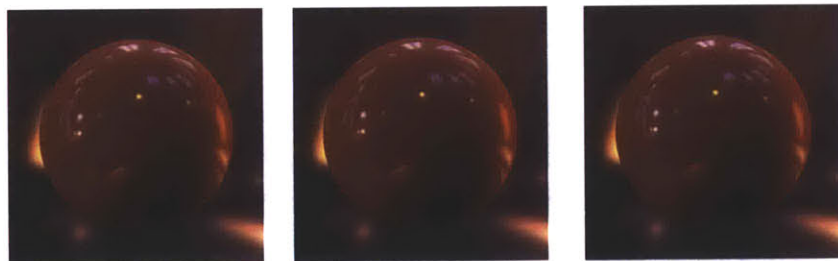
ground truth   recovered (1% noise)   recovered (10% noise)
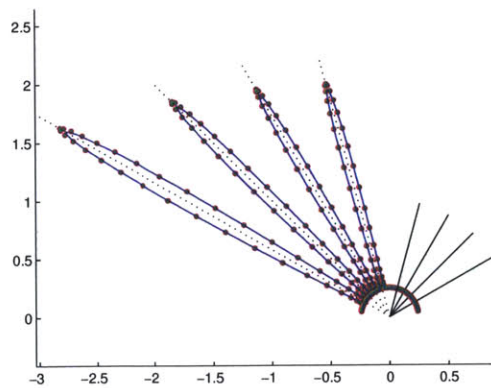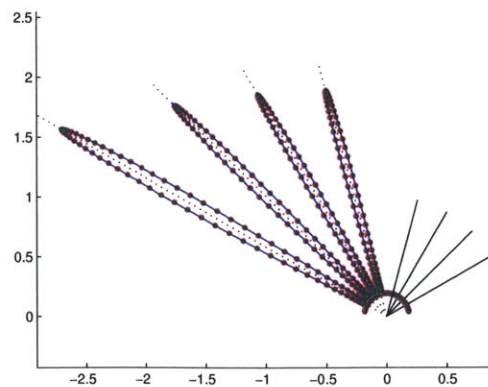
B.1.5   Nickel

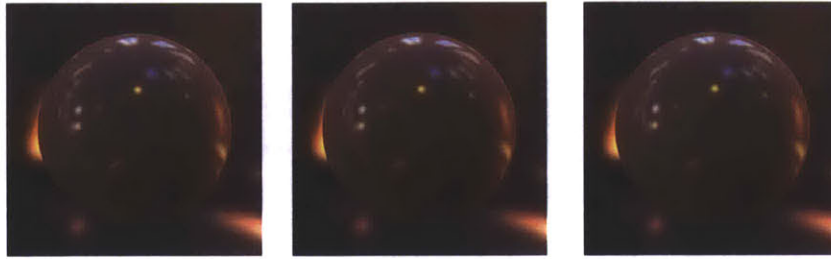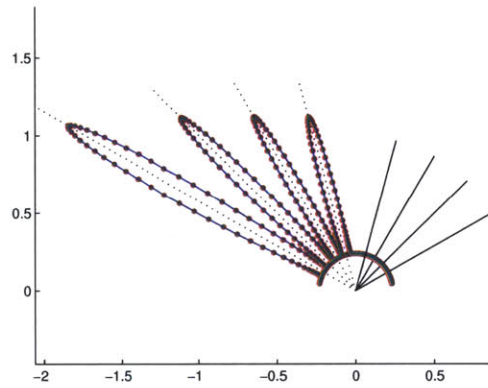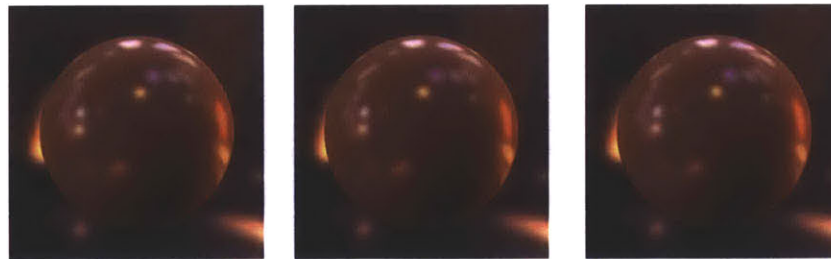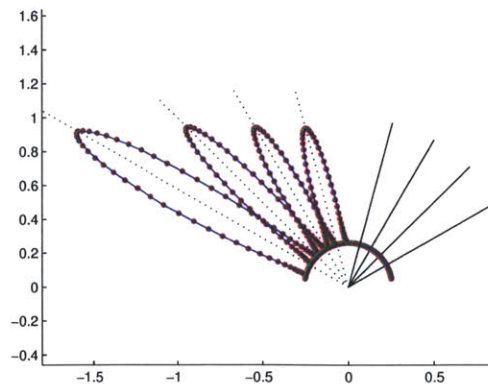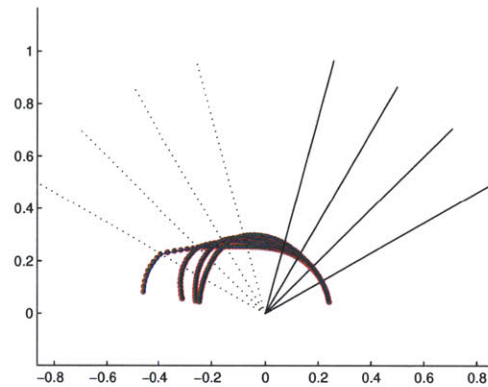ground truth   recovered (1% noise)   recovered (10% noise)
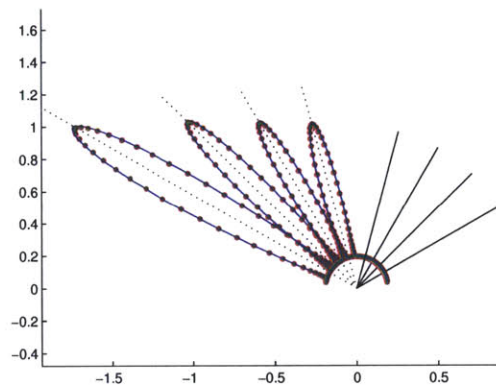
B.1.6   Acrylic blue

ground truth   recovered (1% noise)   recovered (10% noise)

B.1.7 Pink jasper

ground truth   recovered (1% noise)   recovered (10% noise)
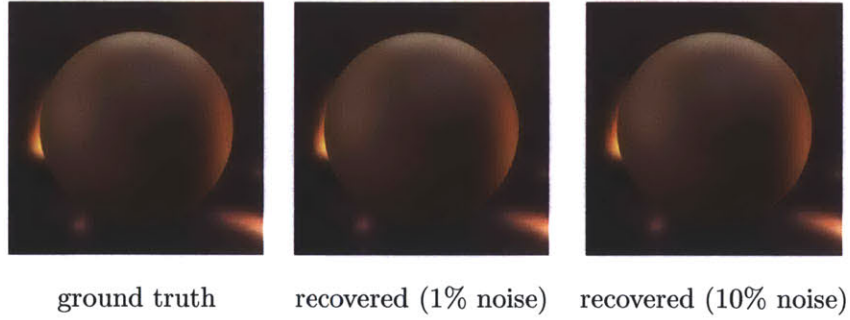

B.1.8 Yellow matte plastic

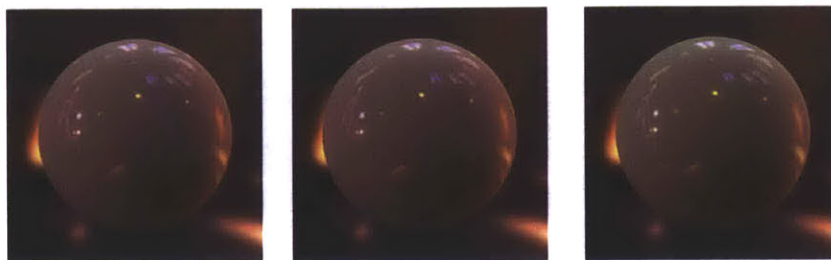ground truth   recovered (1% noise)   recovered (10% noise)

B.1.9 Yellow plastic

ground truth   recovered (1% noise)   recovered (10% noise)

B.2 Synthetic Multiple-Patch Results

In this section, we present results obtained for nine materials simultaneously by arranging them in a 3×3 two-dimensional grid. In each figure, we show renderings of the ground truth, and renderings using recovered Ashikhmin parameters at 1% and 10% additive noise to the streak images. We also show 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%.
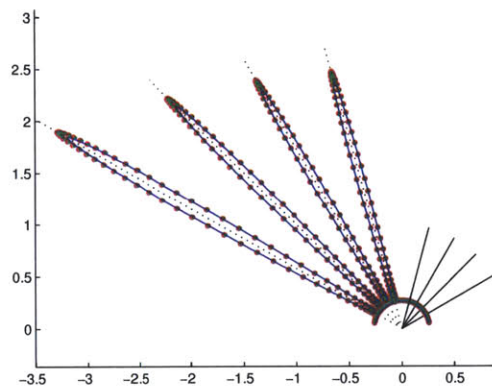
B.2.1   Acrylic white

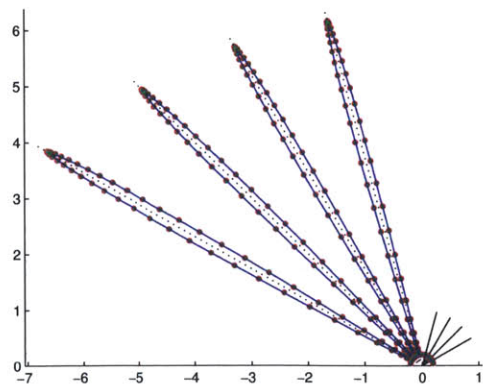ground truth   recovered (1% noise)   recovered (10% noise)
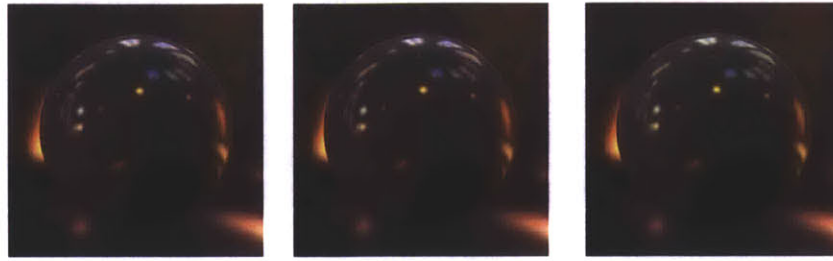

B.2.2   Copper

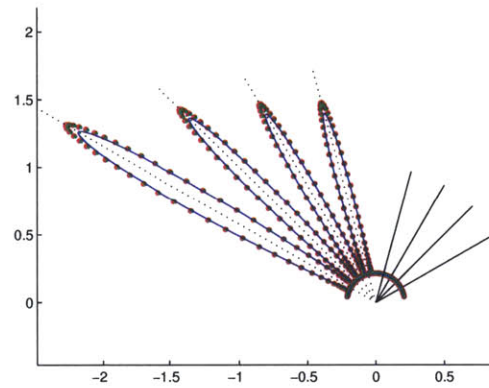ground truth   recovered (1% noise)   recovered (10% noise)
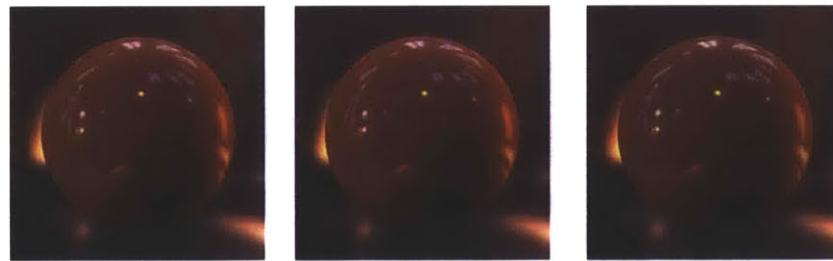
B.2.3 Gray plastic
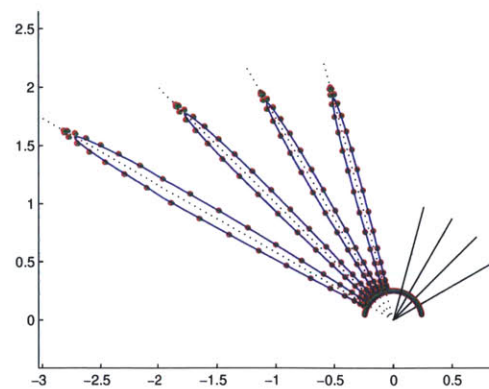

ground truth   recovered (1% noise)   recovered (10% noise)

B.2.4 Red plastic specular

ground truth   recovered (1% noise)   recovered (10% noise)

B.2.5 Nickel

ground truth   recovered (1% noise)   recovered (10% noise)

B.2.6 Acrylic blue

ground truth   recovered (1% noise)   recovered (10% noise)

B.2.7 Pink jasper

ground truth   recovered (1% noise)   recovered (10% noise)

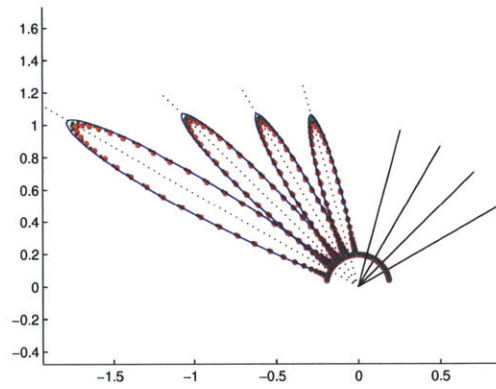## B.2.8 Yellow matte plastic
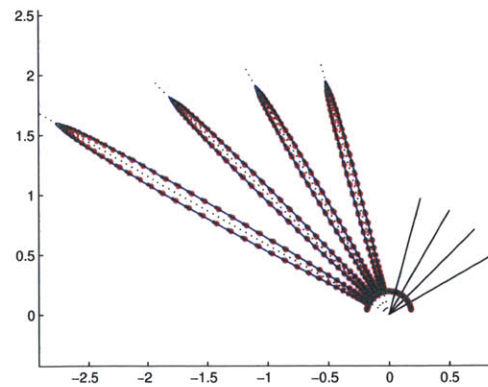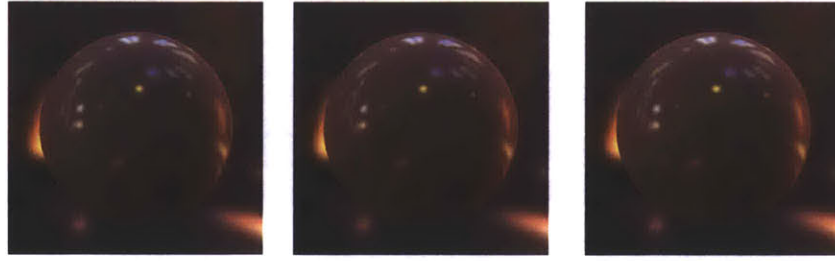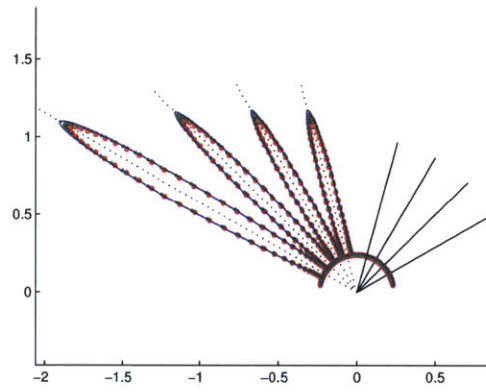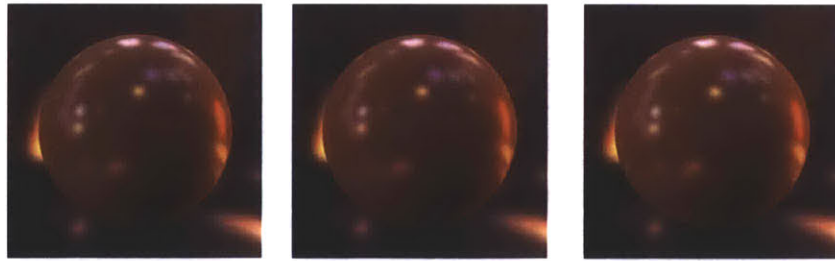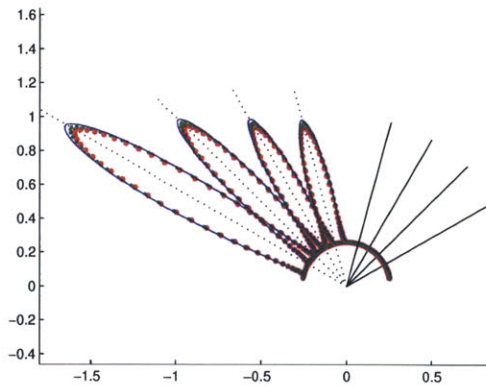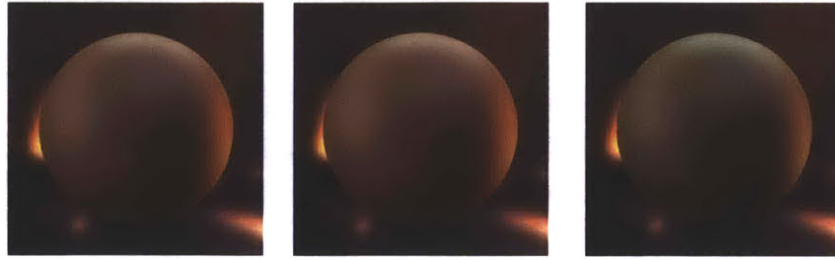

ground truth   recovered (1% noise)   recovered (10% noise)
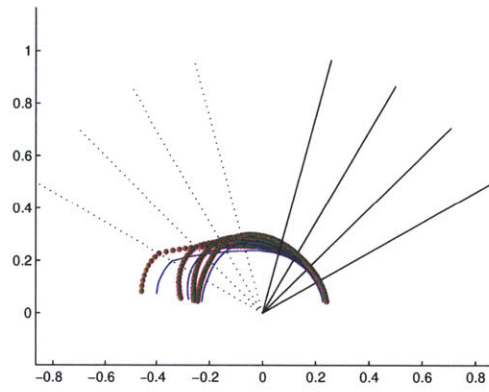
B.2.9 Yellow plastic

ground truth   recovered (1% noise)   recovered (10% noise)

Bibliography

[1] Nikhil Naik, Shuang Zhao, Andreas Velten, Ramesh Raskar, and Kavita Bala. Single view reflectance capture using multiplexed scattering and time-of-flight imaging. *ACM Trans. Graph.*, 30(6):171:1–171:10, December 2011.

[2] A. Campillo and S. Shapiro. Picosecond streak camera fluorometry - a review. *Quantum Electronics, IEEE Journal of*, 19(4):585 –603, april 1983.

[3] J.M. Schmitt. Optical coherence tomography (oct): a review. *Selected Topics in Quantum Electronics, IEEE Journal of*, 5(4):1205 –1215, jul/aug 1999.

[4] W Denk, JH Strickler, and WW Webb. Two-photon laser scanning fluorescence microscopy. *Science*, 248(4951):73–76, 1990.

[5] Ramesh Raskar and James Davis. 5d time-light transport matrix: What can we reason about scene properties? *MIT Technical Report*, 2008.

[6] Ahmed Kirmani, Tyler Hutchison, James Davis, and Ramesh Raskar. Looking around the corner using transient imaging. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 159 –166, 29 2009-oct. 2 2009.

[7] R. Pandharkar, A. Velten, A. Bardagjy, E. Lawson, M. Bawendi, and R. Raskar. Estimating motion and size of moving non-line-of-sight objects in cluttered environments. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 265 –272, june 2011.

[8] Andreas Velten, Thomas Willwacher, Otkrist Gupta, Ashok Veeraraghavan, Moungi Bawendi, and Ramesh Raskar. Recovering three-dimensional shape around a corner using ultra-fast time-of-flight imaging. *Nature Communications*, 2011.

[9] Andreas Velten, Everett Lawson, Andrew Bardagjy, Moungi Bawendi, and Ramesh Raskar. Slow art with a trillion frames per second camera. In *ACM SIGGRAPH 2011 Posters*, SIGGRAPH '11, pages 13:1–13:1, New York, NY, USA, 2011. ACM.

[10] D. Wu, M. O'Toole, A. Velten, Agrawal A., and R. Raskar. Decomposing global light transport using time of flight imaging. In *Computer Vision and Pattern Recognition, 2012. CVPR 2012. IEEE Conference on*, pages 1–8. Ieee, 2012.

[11] F. E. Nicodemus, J. C. Richmond, and J. J. Hsia. Geometrical considerations and reflectance. *National Bureau of Standards*, (NBS Monograph 160), October 1977.

[12] Gregory J. Ward. Measuring and modeling anisotropic reflection. In *ACM SIGGRAPH 1992 papers*, pages 265–272, 1992.

[13] S. Marschner, S. Westin, E. Lafortune, K. Torrance, and D. Greenberg. Image-Based BRDF measurement including human skin. In *Eurographics Workshop on Rendering*, pages 139–152, 1999.

[14] Kristin Dana, Bram van Ginneken, Shree Nayar, and Jan Koenderink. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics*, 18(1):1–34, 1999.

[15] Hendrik P. A. Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel. Image-based reconstruction of spatially varying materials. In *Eurographics Workshop on Rendering*, pages 63–70, 2001.

[16] Hendrik P. A. Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel. Image-based reconstruction of spatial appearance and geometric detail. *ACM Transactions on Graphics*, 22(2):234–257, 2003.

[17] Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. A data-driven reflectance model. In *ACM SIGGRAPH 2003 Papers*, pages 759–769, 2003.

[18] Addy Ngan, Frédo Durand, and Wojciech Matusik. Experimental Analysis of BRDF Models. In *Eurographics Symposium on Rendering*, pages 117–226, 2005.

[19] Jason Lawrence, Aner Ben-Artzi, Christopher DeCoro, Wojciech Matusik, Hanspeter Pfister, Ravi Ramamoorthi, and Szymon Rusinkiewicz. Inverse shade trees for non-parametric material representation and editing. In *ACM SIGGRAPH 2006 Papers*, pages 735–745, 2006.

[20] Abhijeet Ghosh, Tongbo Chen, Pieter Peers, Cyrus A. Wilson, and Paul Debevec. Estimating specular roughness and anisotropy from second order spherical gradient illumination. *Computer Graphics Forum*, 28(4):1161–1170, 2009.

[21] Abhijeet Ghosh, Tongbo Chen, Pieter Peers, Cyrus A. Wilson, and Paul Debevec. Circularly polarized spherical illumination reflectometry. In *ACM SIGGRAPH Asia 2010 papers*, pages 162:1–162:12, 2010.

[22] Yue Dong, Jiaping Wang, Xin Tong, John Snyder, Yanxiang Lan, Moshe Ben-Ezra, and Baining Guo. Manifold bootstrapping for SVBRDF capture. In *ACM SIGGRAPH 2010 papers*, pages 98:1–98:10, 2010.

[23] Jiaping Wang, Yue Dong, Xin Tong, Zhouchen Lin, and Baining Guo. Kernel nystrom method for light transport. In *ACM SIGGRAPH 2009 papers*, pages 29:1–29:10, 2009.

[24] Tim Weyrich, Jason Lawrence, Hendrik P. A. Lensch, Szymon Rusinkiewicz, and Todd Zickler. Principles of appearance acquisition and representation. *Foundations and Trends in Computer Graphics and Vision*, 4(2):75–191, 2009.

[25] Y. Sato, M.D. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *ACM SIGGRAPH 1997 papers*, pages 379–387, 1997.

[26] Yizhou Yu, Paul Debevec, Jitendra Malik, and Tim Hawkins. Inverse global illumination: recovering reflectance models of real scenes from photographs. In *ACM SIGGRAPH 1999 papers*, pages 215–224, 1999.

[27] David McAllister. *A Generalized Surface Appearance Representation for Computer Graphics*. PhD thesis, UNC Chapel Hill, 2002.

[28] Andrew Gardner, Chris Tchou, Tim Hawkins, and Paul Debevec. Linear light source reflectometry. In *ACM SIGGRAPH 2003 Papers*, pages 749–758, 2003.

[29] D.B. Goldman, B. Curless, A. Hertzmann, and S.M. Seitz. Shape and spatially-varying BRDFs from photometric stereo. *IEEE PAMI*, 32(6):1060 –1071, june 2010.

[30] Todd Zickler, Sebastian Enrique, Ravi Ramamoorthi, and Peter Belhumeur. Reflectance sharing: Image-based rendering from a sparse set of images. In *Eurographics Symposium on Rendering*, pages 253–264, 2005.

[31] Pradeep Sen, Billy Chen, Gaurav Garg, Stephen R. Marschner, Mark Horowitz, Marc Levoy, and Hendrik P. A. Lensch. Dual photography. In *ACM SIGGRAPH 2005 Papers*, pages 745–755, 2005.

[32] Gaurav Garg, Eino-Ville Talvala, Marc Levoy, and Hendrik P. A. Lensch. Symmetric photography: Exploiting data-sparseness in reflectance fields. In *Rendering Techniques 2006*, pages 251–262, June 2006.

[33] Jiaping Wang, Shuang Zhao, Xin Tong, John Snyder, and Baining Guo. Modeling anisotropic surface reflectance with example-based microfacet synthesis. In *ACM SIGGRAPH 2008 papers*, pages 41:1–41:9, 2008.

[34] Andreas Wenger, Andrew Gardner, Chris Tchou, Jonas Unger, Tim Hawkins, and Paul Debevec. Performance relighting and reflectance transformation with time-multiplexed illumination. In *ACM SIGGRAPH 2005 Papers*, pages 756–764, 2005.

[35] Wan-Chun Ma, Tim Hawkins, Pieter Peers, Charles-Felix Chabert, Malte Weiss, and Paul Debevec. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Eurographics Symposium on Rendering*, pages 183–194, 2007.

[36] Abhijeet Ghosh, Wolfgang Heidrich, Shruthi Achutha, and Matthew O'Toole. A basis illumination approach to BRDF measurement. *Int. J. Comput. Vision*, 90:183–197, November 2010.

[37] Jefferson Y. Han and Ken Perlin. Measuring bidirectional texture reflectance with a kaleidoscope. In *ACM SIGGRAPH 2003 Papers*, pages 741–748, 2003.

[38] Tim Hawkins, Per Einarsson, and Paul E. Debevec. A dual light stage. In Kavita Bala and Phil Dutre, editors, *Rendering Techniques*, pages 91–98, 2005.

[39] Sujit Kuthirummal and Shree K. Nayar. Multiview radial catadioptric imaging for scene capture. In *ACM SIGGRAPH 2006 Papers*, pages 916–923, 2006.

[40] D. Forsyth and A. Zisserman. Shape from shading in the light of mutual illumination. *Image Vision Comput.*, 8:42–49, February 1990.

[41] Shree K. Nayar, Katsushi Ikeuchi, and Takeo Kanade. Shape from interreflections. *Int. J. Comput. Vision*, 6:173–195, August 1991.

[42] Shree K. Nayar, Gurunandan Krishnan, Michael D. Grossberg, and Ramesh Raskar. Fast separation of direct and global components of a scene using high frequency illumination. In *ACM SIGGRAPH 2006 Papers*, pages 935–944, 2006.

[43] Steven M. Seitz, Yasuyuki Matsushita, and Kiriakos N. Kutulakos. A theory of inverse light transport. In *ICCV 2005*, pages 1440–1447, 2005.

[44] S Liu, TT Ng, and Y Matsushita. Shape from second-bounce of light transport. In *ECCV 2010*, pages 280–293, 2010.

[45] Jiamin Bai, Manmohan Chandraker, Tian-Tsong Ng, and Ravi Ramamoorthi. A dual theory of inverse and forward light transport. In *ECCV 2010*, pages 1–8, 2010.

[46] Michael Holroyd, Jason Lawrence, Greg Humphreys, and Todd Zickler. A photometric approach for estimating normals and tangents. In *ACM SIGGRAPH Asia 2008 papers*, pages 133:1–133:9, 2008.

[47] K. M. Yoo, Qirong Xing, and R. R. Alfano. Imaging objects hidden in highly scattering media using femtosecond second-harmonic-generation cross-correlation time gating. *Opt. Lett.*, 16(13):1019–1021, Jul 1991.

[48] L. WANG, P. P. HO, C. LIU, G. ZHANG, and R. R. ALFANO. Ballistic 2-d imaging through scattering walls using an ultra-fast optical kerr gate. *Science*, 253(5021):769–771, 1991.

[49] K. M. Yoo, Zhi-Wei Zang, S. A. Ahmed, and R. R. Alfano. Imaging objects hidden in scattering media using a fluorescence-absorption technique. *Opt. Lett.*, 16(16):1252–1254, Aug 1991.

[50] J.M. Schmitt. Optical coherence tomography (oct): a review. *Selected Topics in Quantum Electronics, IEEE Journal of*, 5(4):1205 –1215, jul/aug 1999.

[51] D.A. Boas, D.H. Brooks, E.L. Miller, C.A. DiMarzio, M. Kilmer, R.J. Gaudette, and Quan Zhang. Imaging the body with diffuse optical tomography. *Signal Processing Magazine, IEEE*, 18(6):57 –75, nov 2001.

[52] Ahmed Kirmani, Tyler Hutchison, James Davis, and Ramesh Raskar. Looking around the corner using transient imaging. In *ICCV 2009*, pages 159–166, 2009.

[53] Szymon Rusinkiewicz. A new change of variables for efficient BRDF representation. In *Eurographics Workshop on Rendering*, pages 11–22, 1998.

[54] Michael Ashikhmin. Distribution-based BRDFs. *University of Utah – Technical Report*, 2007.

[55] Michael Ashikmin, Simon Premože, and Peter Shirley. A microfacet-based BRDF generator. In *ACM SIGGRAPH 2000 Papers*, pages 65–74, 2000.

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

[57] S.W. Hasinoff, F. Durand, and W.T. Freeman. Noise-optimal capture for high dynamic range photography. In *CVPR 2010*, pages 553 –560, june 2010.

[58] Y.Y. Schechner, S.K. Nayar, and P.N. Belhumeur. Multiplexing for optimal lighting. *IEEE PAMI*, 29(8):1339 –1354, aug. 2007.

[59] J. M. Dudley, D. T. Reid, M. Ebrahimzadeh, and W. Sibbett. Characterisics of a non-critically phasematched Ti:sapphire pumped femtosecond optical parametric oscillator. *Optics Communications*, 104(4,5,6):419–430, January 1994.

[60] John J. Degnan. Asynchronous laser transponders for precise interplanetary ranging and time transfer. *Journal of Geodynamics*, 34(3-4):551 – 594, 2002.

[61] Ryan E. Warburton, Aongus McCarthy, Andrew M. Wallace, Sergio Hernandez-Marin, Robert H. Hadfield, Sae Woo Nam, and Gerald S. Buller. Subcentimeter depth resolution using a single-photon counting time-of-flight laser ranging system at 1550 nm wavelength. *Opt. Lett.*, 32(15):2266–2268, Aug 2007.