

Shield Fields: Modeling and Capturing 3D Occluders

Douglas Lanman^{1,2}

Ramesh Raskar^{1,3}

Amit Agrawal¹

Gabriel Taubin²

¹Mitsubishi Electric Research Laboratories (MERL)

²Brown University

³MIT Media Lab



Figure 1: Our shield field camera allows instantaneous capture and reconstruction of the visual hull of 3D occluders using static illumination. (Left) The shield field camera, containing: an array of point lights, an object for reconstruction, and a large-format, mask-based light field camera. (Middle) Combined shadowgrams imaged on a diffusing screen after passing through a tiled-broadband attenuator. (Right) Four views of the visual hull reconstructed using only a single photograph captured by the proposed system.

Abstract

We describe a unified representation of occluders in light transport and photography using shield fields: the 4D attenuation function which acts on any light field incident on an occluder. Our key theoretical result is that shield fields can be used to decouple the effects of occluders and incident illumination. We first describe the properties of shield fields in the frequency-domain and briefly analyze the “forward” problem of efficiently computing cast shadows. Afterwards, we apply the shield field signal-processing framework to make several new observations regarding the “inverse” problem of reconstructing 3D occluders from cast shadows – extending previous work on shape-from-silhouette and visual hull methods. From this analysis we develop the first single-camera, single-shot approach to capture visual hulls without requiring moving or programmable illumination. We analyze several competing camera designs, ultimately leading to the development of a new large-format, mask-based light field camera that exploits optimal tiled-broadband codes for light-efficient shield field capture. We conclude by presenting a detailed experimental analysis of shield field capture and 3D occluder reconstruction.

Keywords: Computational Photography, Light Fields, Coded Aperture Imaging, Visual Hull, Cast Shadows, Light Transport

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Digitizing and scanning

1 Introduction

Image formation, for example the shadows due to a complex illumination source, depends on the interaction of light and objects within the scene. In this paper, we present a general theory of shield fields: the 4D attenuation function describing the attenuation of incident light fields by occluders. We apply this theory to develop a new single-shot method for acquiring shield fields and visual hull reconstructions using light field cameras. We also present a new class of optimally light-efficient tiled-broadband codes which can be used to develop large-format, heterodyne light field cameras for shield field capture inspired by the design of Veeraraghavan et al. [2007]. Our analysis unifies and extends the applications of several previous ideas on light transport, cast shadows, and light field capture, in particular those due to Chai et al. [2000], Isaksen et al. [2000], and Durand et al. [2005].

At first glance the concept of a light field [Gortler et al. 1996; Levoy and Hanrahan 1996] due to an object that attenuates light might seem unusual. Previously, Debevec et al. [2000] used 8D reflectance fields to describe the object-dependent transformation from incident to reflected 4D light fields. We observe that a general 3D occluder also attenuates the incident 4D light field according to its shape and absorption properties. We refer to this 4D attenuation function as the shield field, which can also be viewed as a 4D slice of a more general 8D reflectance field (where the incident and reflected rays are coincident). Although cast shadows are a complex function of the incident illumination and the occluder, we show that we can *decouple* these effects using shield fields. We analyze two types of problems for shield fields. First, we examine the “forward” problem: given an occluder, using its shield field to efficiently compute cast shadows. Second, we examine the “inverse” problem: reconstructing an occluder given a measured shield field – analogous to the existing shape-from-silhouette problem. The primary focus of this paper is on the later.

In the inverse problem, we begin by analyzing a variety of light field cameras which could be used for capturing, in real-time, the shield fields of objects and reconstructing their visual hulls. We show, using a frequency-domain analysis, how both a simple pinhole array and the more recent heterodyne masks could be used in our applica-

tion. This analysis ultimately leads to a new family of heterodyne patterns that can utilize any tiled-broadband code – significantly contributing to the understanding of mask-based light field cameras. We demonstrate the utility of these masks by constructing a prototype shield field camera and analyze its performance compared to existing visual hull systems. Ultimately, we hope that shield fields provide new insights for both the forward and the inverse problems and that they will stimulate further application of our analysis in higher dimensions and in the frequency domain.

1.1 Contributions

We present a set of techniques to analyze and capture the effects of volumetric occluders using methods inspired by light field photography. Specific technical contributions are as follows.

- We show that 3D occluders can be represented using 4D shield fields and this representation can handle multiple occluders with different shapes. We also describe the spectral properties of shield fields using a frequency-domain analysis and unify occluder analysis across both the forward problem of predicting cast shadows, as well as the inverse problem of occluder reconstruction from light fields.
- We propose a new combination of a light field camera and an area illuminator which allows real-time capture of the visual hull of 3D occluders.
- We develop a new class of optimally light-efficient attenuation masks for light field and shield field capture and show that the Sum-of-Sinusoids mask proposed in [Veeraraghavan et al. 2007] is a special member of this class.
- We apply a signal-processing analysis to shield fields to provide a detailed understanding of the sampling requirements and reconstruction limitations for visual hull methods using light field cameras.

1.2 Related Work

Light Field Capture and Analysis: Light fields were first proposed by Levoy and Hanrahan [1996] and Gortler et al. [1996] to characterize the set of geometric light rays within a volume. Several camera designs for light field capture have been developed within the last century. In the early 20th century, Ives [1928] and Lippman [1908] used pinhole arrays. Camera arrays replaced pinhole arrays in [Wilburn et al. 2005]. Georgiev et al. [2006] used a novel combination of prisms and lenses in front of the camera for light field capture. The plenoptic cameras of Adelson and Wang [1992] and Ng et al. [2005] used a single lens to form an image on a lenslet array. In the heterodyne camera [Veeraraghavan et al. 2007], the lenslet array was replaced by a sinusoidal attenuating mask. We show that shield fields explain these designs in a common framework. Beyond capture systems, frequency domain methods are used in [Durand et al. 2005] to explain occlusion and shading, as well as for analyzing diffraction and interference [Goodman 1996]. However, in this paper we limit our scope to geometric optics.

Coded Aperture Imaging: Coded aperture imaging [Fenimore and Cannon 1978; Accorsia et al. 2001] has been used in astronomical and medical imaging to capture X-rays and gamma rays. Recent applications of high-frequency attenuating patterns in computer vision and graphics include: Nayar et al. [2006] for separating the effects of global and direct illumination, Levin et al. [2007] and Farid [1997] for estimating intensity and depth from defocused images, and [Talvala et al. 2007] to minimize the effects of glare. In addition, parallax barriers have been previously applied to create automultiscopic displays [Zwicker et al. 2006]. We show that




Figure 2: (Left) Two-plane light field parametrization (x, θ) . (Right) A single occluder plane parallel to the receiver plane. Incident ray (x, θ) is attenuated by the occluder at $x - z\theta$.

such attenuating masks can be analyzed in terms of shield fields as general scene occluders.

Shadows: Analysis of cast shadows is important for several computer vision applications. Similarly, efficient methods for the computation of cast shadows is a vital topic in computer graphics. Thornber and Jacobs [2001] use the integral involving the illumination and reflectance properties of surfaces and the visibility constraint, whereas Soler and Sillion [1998] evaluate soft shadow textures using convolution. Ramamoorthi et al. [2005] analyze cast shadows using Fourier basis functions. Ramamoorthi et al. [2007] extend the light transport frequency-domain analysis of Durand et al. [2005] to understand the effect of cast shadows and visibility for arbitrary curved occluders. In a closely-related work, Zhou et al. [2005] use precomputed shadow fields for efficient dynamic soft shadow rendering. We observe that shield fields also allow efficient shadow computation via convolution of the incident light field with the shield field in the frequency domain.

Visual Hulls: Shape-from-silhouette and visual hull approximation algorithms are appealing because the structure of an object can be passively-estimated from multiple images without feature-matching. It has been shown, however, that these algorithms are highly-sensitive to camera calibration errors [Yamazaki et al. 2007]. This sensitivity becomes increasingly apparent as the number of views grows, resulting in poor-quality models. Savarese et al. [2001] proposed a system that avoids silhouette matting. In this system, silhouettes from multiple views were obtained by rotating an object on a turntable and observing shadows cast on a diffusing screen by a single point light source. Several recent systems exploit these planar shadowgrams, including [Yamazaki et al. 2007]. Unlike existing systems, our design allows instantaneous (i.e., single exposure) capture of coplanar shadowgrams using the shield field camera described in Section 3. As a result, our approach allows reconstruction of real-time deforming occluders.

2 Shield Fields

In this section we provide a comprehensive analysis of volumetric occlusion via shield field analysis in ray-space and in the frequency-domain. We show that the effects of illumination can be decoupled from occluder shape when computing shadows using shield fields.

2.1 Basic Concepts

For simplicity, we consider a 2D light field incident on a 1D receiver plane (since 4D light fields and 2D planar receivers are straightforward extensions). We restrict our analysis to a single wavelength of light and ignore the effects of indirect reflection or refraction from the occluder to the sensor. Let the light field be described using the two-plane parametrization [Chai et al. 2000; Durand et al. 2005] shown in Figure 2, where x denotes the spatial dimension and θ denotes the position of intersection (in the local frame of x) of an




Figure 3: (Left) The occlusion function $o(\xi)$ is a square wave. At the occluder plane the shield field will have lines parallel to θ , since the attenuation depends only on the spatial location. (Middle) At the receiver plane the lines remain parallel, but shear depending on the distance z . (Right) The Fourier transform of the shield field concentrates on a line in 2D.

incident ray with a second plane that is a unit distance away from, and parallel to, the first plane. In the absence of any occluders, the incident light field at a receiver plane is defined as $l_{\text{receiver}}(x, \theta)$.

Now consider an occluder placed in front of the receiver. In the absence of an occluder, the receiver plane will record the incident light field $l_{\text{incident}}(x, \theta)$. Assume that the occluder $o(\xi)$, possibly a non-binary attenuator, is located in another parallel plane separated by a distance z from the receiver plane. By tracing the ray (x, θ) backwards, we find it intersects the occluder plane at $\xi = x - z\theta$. As a result, the received light field $l_{\text{receiver}}(x, \theta)$, in the presence of the occluder $o(\xi)$, is given by the multiplication of incident light field by $o(x - z\theta)$. In general, we define the *shield field* $s(x, \theta)$ as the attenuation function applied to the incident light field as follows.

$$l_{\text{receiver}}(x, \theta) = s(x, \theta) l_{\text{incident}}(x, \theta) \quad (1)$$

For the case of equal attenuation as a function of incidence angle, the shield field for a planar occluder, measured in the occluder plane, will be given by $s(x, \theta) = o(x)$. Such an attenuator can be described as a ‘‘Lambertian occluder’’ in analogy to the well-known case of Lambertian reflectors.

As defined, the shield field quantifies the attenuation of each ray due to the occluding surfaces encountered along its path from the emitter to the receiver plane. Physically, the shield field can be understood as the resultant light field due to an occluder when the incident illumination is uniform (i.e., all rays have equal unit radiance). Since the shield field describes the 4D attenuation under uniform illumination, we find that it only depends on the spatial attenuation of the occluder – leading to the key result that shield fields allow the effects of occluders and illumination to be decoupled. Thus, knowing the shield field, shadows on the receiver plane could be obtained by computing the incident illumination light field on the receiver plane and then multiplying by the known shield field.

2.2 Frequency Domain Analysis

Now we turn our attention to the spectral properties of shield fields through a frequency-domain analysis. In the following discussion, let $S(f_x, f_\theta)$ denote the 2D Fourier transform of $s(x, \theta)$, where f_x and f_θ are the frequencies in x and θ , respectively.

Planar Occluders: The shield field due to a planar occluder with an attenuation pattern $o(\xi)$ at the occluder plane is given by $s(x, \theta) = o(x)$. This means that the shield field at the occluder plane only depends on the spatial dimension x and is independent of the angular dimension θ . From known properties of the Fourier transform, all the energy in the spectrum $O(f_x)$ will be concentrated along the \hat{f}_x -axis, such that

$$S(f_x, f_\theta) = O(f_x) \delta(f_\theta), \quad (2)$$

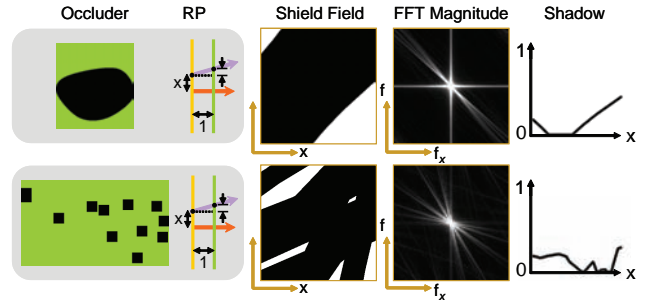


Figure 4: Shield fields of general occluders and resulting Fourier transforms and shadows cast on the receiver plane (RP), assuming uniform incident illumination. Note that each shield field only depends on the occluder shape and its distance from RP.

which vanishes for $\theta \neq 0$.

At the receiver plane, $s(x, \theta) = o(x - z\theta)$. Taking its 2D Fourier transform we find

$$S(f_x, f_\theta) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} o(x - z\theta) e^{-jf_x x} e^{-jf_\theta \theta} dx d\theta. \quad (3)$$

Substituting $u = x - z\theta$ and $v = \theta$, we have $x = u + zv$. By a change of variables, the above integration yields

$$S(f_x, f_\theta) = O(f_x) \delta(f_\theta + f_x z). \quad (4)$$

Thus, at the receiver plane, all the energy of the shield field spectrum lies along the line $f_\theta + f_x z = 0$, as shown in Figure 3. The slope of this line depends on the distance of the occluder from the receiver plane. If $z = 0$, the line coincides with the \hat{f}_x -axis. At $z = \infty$, this line coincides with the \hat{f}_θ -axis.

General Occluders: For a general scene, occluders could be composed of a collection of semi-transparent and opaque objects. One can approximate a general occluder (to any required degree of accuracy) using a set of planes parallel to the receiver. The effect of each of these planar slices can be analytically computed using the shield field equation for a single plane. The overall shield field can then be found as the product of the individual shield fields.

In ray-space, the occluder can be approximated as a combination of k parallel planes at distances $\{z_1, \dots, z_k\}$. The combined shield field $s(x, \theta)$, in the receiver plane coordinate system, is then given by the product of the individual shield fields such that

$$s(x, \theta) = \prod_{i=1}^k o_i(x - z_i \theta). \quad (5)$$

In the frequency-domain, the combined Fourier transform of the shield field can be computed as a convolution of the Fourier transforms of the individual shield fields for each parallel plane in the approximation as follows.

$$S(f_x, f_\theta) = S_1(f_x, f_\theta) * S_2(f_x, f_\theta) \dots * S_k(f_x, f_\theta) \quad (6)$$

2.3 Modeling the Receiver Plane

When the receiver surface is planar, the cast shadows (for any illumination source) are simply a projection of the received light field along the θ direction. This can be efficiently computed using frequency-domain techniques. Recall that the Fourier slice theorem [Ng 2005] shows that the 1D Fourier transform of a projection




Figure 5: Illustration of light field capture and subview-decoding using a 6×6 array of point sources. (Left) Scene as viewed from the emitter plane. (Top right) From left to right: (1) diffuser-plane image recorded using a tiled-pinhole mask and a 30 second exposure, (2) inset region of pinhole diffuser-plane image showing spatially-multiplexed angular subviews, and (3) recovered shadowgrams in central 4×4 region. Note that the presence/absence of a peak in each 6×6 pinhole image indicates if the corresponding light is visible from that sensor position. (Bottom right) From left to right: (1) diffuser-plane image recorded using a tiled-MURA mask and a 0.25 second exposure, (2) inset region showing linear superposition of projected patterns, and (3) recovered shadowgrams. Note that the tiled-MURA pattern adds some decoding noise to the recovered shadowgrams, but requires an exposure time which is two orders of magnitude shorter than pinholes.

of the 2D light field is equivalent to a 1D slice of its 2D Fourier transform. As a result, the cast shadows can be computed by evaluating a slice of the Fourier transform of the incident light field and computing its inverse Fourier transform. Figure 4 shows several shield fields and their corresponding cast shadows for various occluder configurations.

3 Inverse Problem: Shield Field Capture

In the previous section, we showed how shield fields can be used to decouple the effects of occluders and incident illumination. In this section, we describe capturing shield fields using a *shield field camera*: a light field camera and associated illumination system optimized for measuring the cast shadows produced by real-world objects. After measuring the shield field for a given object, we show how to reconstruct its visual hull. Since shield fields can describe the cast shadows for each light source, existing shadowgram-based visual hull methods can be applied to estimate the occluder shape [Savarese et al. 2001; Yamazaki et al. 2007]. Note that, in this paper, we focus our attention on measuring and reconstructing only opaque objects. In addition, we emphasize that our system records the shield field of an object in a single measurement without moving or programmable illumination, and is the *first* shadowgram-based architecture to allow real-time visual hull reconstruction.

3.1 Optical Design

Given our stated goal of measuring, in a “single shot”, the shield field of an opaque object, the simplest possible design is one that directly mirrors our prior analysis: a single light field camera observing a scene composed of opaque objects and lit from behind by a large area light source. In such a system, the object’s shield field $s_{\text{occluder}}(x, \theta)$ can be recovered using an additional calibration image taken without an occluder present in the scene. In this case the camera will directly record the incident light field $l_{\text{incident}}(x, \theta)$.

From Equation 1, we find that the shield field can be recovered by dividing the occluder’s light field $l_{\text{occluder}}(x, \theta)$ by the incident light field as follows.

$$s_{\text{occluder}}(x, \theta) = \frac{l_{\text{occluder}}(x, \theta)}{l_{\text{incident}}(x, \theta)} \quad (7)$$

Thus, the incident light field should be non-zero for all sampled rays (x, θ) . While not strictly required, we propose using a simple area light source which fills the light field camera’s field of view. As we’ll see later, one could also use a point light source array with one source per angular-sampling bin.

Before we examine specific light field camera designs, let’s sketch the basic architecture of our desired system. Our shield field camera will possess two primary elements. First, it will have a large-format light field camera serving as the receiving surface; for our prototype, we propose a receiver baseline of $l_{\text{receiver}} \approx 1\text{m}$. Second, our system should have a large area light source wider than the receiving surface (in order to fill its field of view). We will use a light box or a uniform array of point sources $l_{\text{emitter}} \approx 2\text{m}$ in width. If these surfaces are separated by $d_{\text{emitter}} \approx 1\text{m}$, then we expect to capture the shield field of objects of about half a meter in diameter or smaller (e.g., small sculptures and many household items).

3.2 Design Choices

As discussed in the introduction, a wide-variety of light field cameras have been previously published. Designs we have considered include: camera arrays [Wilburn et al. 2005], pinhole arrays [Lippmann 1908; Ives 1928], heterodyne cameras [Veeraraghavan et al. 2007], lenslet arrays [Ng et al. 2005], and multiple-element combinations of lenses and prisms [Georgiev et al. 2006]. For our application, we recognize that the key criterion is the ability to achieve a very large receiver baseline. In order to study the discrete sampling of shield fields, we also desire a system with easily-controllable spatial and angular sampling rates.

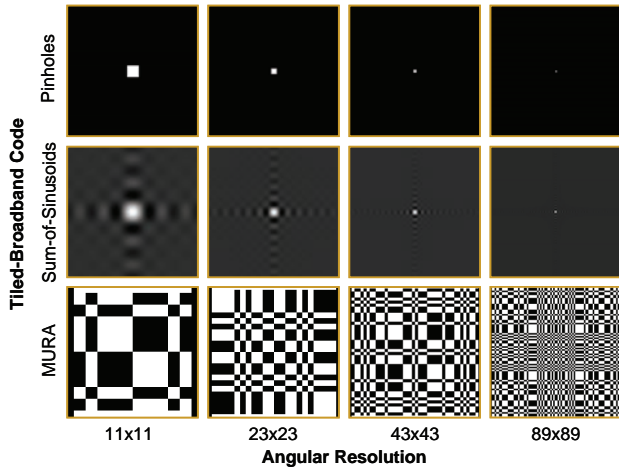


Figure 6: Comparison of several tiled-broadband patterns for use in any heterodyne light field camera. Each row, from left to right, shows a single broadband tile for several angular-sampling rates, including: pinholes, Sum-of-Sinusoids, and MURA. Note that the Sum-of-Sinusoids tiles converge to pinholes for larger angular resolutions, whereas MURA tiles continue to have a 50% duty cycle.

Calibration and Synchronization Issues: While camera arrays have been used to record large-baseline light fields [Wilburn et al. 2005] and reconstruct visual hull models [Matusik et al. 2001], we recognize that a single-sensor camera can also achieve similar baselines – eliminating the calibration and synchronization issues inherent in multiple-camera systems. For instance, Savarese et al. [2001] and Yamazaki et al. [2007] recovered shadowgrams of real-world objects using a single high-resolution camera to photograph the shadow produced by a point light on a diffusing screen. We decided to follow their approach in our design.

Scalability and Cost Issues: Several competing technologies could be used to record a light field using a diffusing screen as the imaging “sensor”. First, a lenslet array could be placed in front of the diffusing screen to form a uniform array of images. Unfortunately, in this configuration, the lenslet array would need to be a similar size as the diffuser and would be difficult to scale to large sizes. Thus, we have elected to use mask-based systems, since they are both low-cost and are capable of being printed at large scales. As shown in Figure 1, our final design includes an additional printed mask in front of the diffusing screen. One can imagine that this architecture could be easily scaled to larger baselines, since masks can be printed or tiled to achieve arbitrary dimensions.

4 Optimal Masks for Light Field Cameras

The heterodyne light field camera introduced by Veeraraghavan et al. [2007] uses a sinusoidal mask placed close to the sensor in order to capture the incident light field. In that work, the key concept was that the Fourier transform of the mask should consist of a sum of equally-spaced impulses (which cause spectral replication of the incident light field). However, in [Veeraraghavan et al. 2007], an ad-hoc Sum-of-Sinusoids (SoS) mask was used to obtain the necessary frequency-domain impulse train. A theoretical question remains: is there an equivalent family of masks which also produce frequency-domain impulse trains? In this paper, we prove that such patterns can be generalized to form a broader class of equivalent *tiled-broadband codes*, and that the SoS mask is a special member of this family. In fact, a uniform array of pinholes could also be used. From this analysis, we derive the *optimal* binary pattern for maximizing light throughput. This optimal pattern is required for

two reasons. First, we need a mask that transmits as much light as possible to reduce the capture time. Second, binary masks are easier to print and do not suffer from quantization noise compared to continuous SoS masks.

Consider the following intuitive explanation. The SoS pattern used in [Veeraraghavan et al. 2007] has sinusoids of *equal* phase. Changing the phase of a cosine does not change the magnitude of its Fourier transform. Thus, we immediately realize that a family of masks exists with sinusoids of varying phases – each member consisting of impulse trains in the frequency-domain. To obtain a mask with maximum light transmission, one could optimize over the phase of the individual sinusoids. However, this still leads to continuous masks and one wonders if there exists equivalent binary masks. Our theoretical contribution is based on the fact that the tiling of *any* pattern leads to impulse trains in the frequency-domain. The magnitude and phase of these impulses depends on the pattern which is tiled. Since, for making spectral replicas, we would like impulses with approximately-equal magnitudes, we conclude that the tiled pattern should be broadband. Hence, tiled-broadband patterns constitute a general class of such masks. Below, we compare these masks in detail.

4.1 Pinhole Arrays

Although not realized by [Veeraraghavan et al. 2007], a pinhole array mask also leads to impulses in frequency domain. Consider a single planar occluder placed at a small distance d_{pinhole} from the sensor, which is opaque except for a series of uniformly-spaced pinholes. The pinhole array occluder $o_{\text{pinhole}}(\xi)$ is given by

$$o_{\text{pinhole}}(\xi) = \sum_{k=-\infty}^{\infty} \delta(\xi - ka_0), \quad (8)$$

where $a_0 = (d_{\text{pinhole}} l_{\text{emitter}})/d_{\text{emitter}}$ is the spacing between pinholes and is selected to ensure that images from adjacent pinholes cannot overlap. The corresponding shield field $s_{\text{pinhole}}(x, \theta)$ is given by

$$\begin{aligned} s_{\text{pinhole}}(x, \theta) &= o_{\text{pinhole}}(x - d_{\text{pinhole}} \theta) \\ &= \sum_{k=-\infty}^{\infty} \delta(x - d_{\text{pinhole}} \theta - ka_0). \end{aligned} \quad (9)$$

Thus, the Fourier transform of the pinhole array shield field is

$$S_{\text{pinhole}}(f_x, f_\theta) = \omega_0 \sum_{k=-\infty}^{\infty} \delta(f_x - k\omega_0) \delta(f_\theta + f_x d_{\text{pinhole}}), \quad (10)$$

where $\omega_0 = 2\pi/a_0$. The shield field spectrum consists of a series of impulses along the line $f_\theta + f_x d_{\text{pinhole}} = 0$. The overall effect of this shield field is to modulate the incoming light field by creating spectral replicas at the center of each impulse.

4.2 Sum-of-Sinusoids

Since pinhole arrays consist of an impulse train in the frequency-domain, they are sufficient for our application. Unfortunately, they severely attenuate the incident light and require either very bright sources or long exposures (prohibiting real-time applications). Recall that an ideal pinhole array leads to an infinitely long impulse train in the Fourier transform. However, depending on angular resolution and sensor bandwidth, only a few of the impulses may be required. Consider the *truncated* Fourier transform of $S_{\text{pinhole}}(f_x, f_\theta)$

$$S_{\text{SoS}}(f_x, f_\theta) = \omega_0 \sum_{k=-(N_\theta-1)/2}^{(N_\theta-1)/2} \delta(f_x - k\omega_0) \delta(f_\theta + f_x d_{\text{pinhole}}), \quad (11)$$

where N_θ is the number of desired angular samples in the captured light field. To achieve this modulation function, we could use one cosine for every pair of impulses, resulting in an occluder given by

$$o_{\text{SoS}}(\xi) = 1 + \sum_{k=1}^{(N_\theta-1)/2} 2 \cos(2\pi k \omega_0 \xi). \quad (12)$$

Note that this attenuation pattern is a summation of *equal-phase* sinusoidal functions with fundamental frequency ω_0 and $(N_\theta - 1)/2$ harmonics. This is exactly the design previously proposed in [Veeraraghavan et al. 2007]. Interestingly, the authors appear to have arrived at their design without realizing the similarity to a pinhole array. Thus, our shield field analysis unifies previous light field capture methods and shows that the SoS mask is a natural extension of pinhole arrays. As shown in Figures 6 and 7, this pattern is significantly more efficient in terms of total transmission. In general, 2D SoS masks (for 4D light field capture) transmit about $\approx 18\%$ of the incident light for angular resolutions of 11×11 or greater.

4.3 General Tiled-Broadband Patterns

While SoS patterns allow more light compared to pinholes, they are continuous-valued functions. We considered printing such patterns using continuous-tone film recorders, such as the Light Valve Technology (LVT) printing process. Unfortunately, commercial LVT printers typically only provide prints up to approximately $25\text{cm} \times 20\text{cm}$ at 1,524 DPI [BowHaus, Inc. 2007]. While one could tile several small printed SoS patterns, this approach leads to seams between masks and non-linearities in the subsequent demodulation. We instead focus on finding equivalent heterodyne patterns with two primary properties: (1) minimal light loss and (2) associated commercial printing processes capable of producing seamless masks with widths in excess of one meter. In addition, we note that SoS patterns require a continuous-tone printing process which introduces additional artifacts in the recovered light field due to printer quantization.

Notice the primary commonality between pinhole arrays and SoS masks: they are both periodic functions. At this point we recall a fundamental result for Fourier transforms: the spectrum of a continuous periodic function is composed of a set of discrete values given by its Fourier series. If we assume that the occlusion function for a single tile, defined by a periodic function of period T , is given by $o_{\text{tile}}(\xi; T)$. The Fourier transform $O_{\text{tile}}(f_\xi; T)$ is then

$$\begin{aligned} O_{\text{tile}}(f_\xi; T) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} o_{\text{tile}}(\xi; T) e^{-j f_\xi \xi} d\xi \\ &= \sum_{k=-\infty}^{\infty} O_{\text{tile}}[k; T] \delta(f_\xi - k f_{\xi_0}), \end{aligned} \quad (13)$$

where $f_{\xi_0} = 2\pi/T$ and the coefficients of the discrete Fourier series $O_{\text{tile}}[k; T]$ are given by

$$O_{\text{tile}}[k; T] = \frac{1}{T} \int_{-T/2}^{T/2} o_{\text{tile}}(\xi; T) e^{-j k f_{\xi_0} \xi} d\xi. \quad (14)$$

Thus, the spectrum of *any* periodic function is composed of a weighted combination of impulses. As discussed earlier, we would like the impulses to have approximately the same magnitude. Thus, an individual tile of the periodic function should be *broadband*. In addition, to achieve the same effect as a SoS mask, the tiled-broadband mask should be placed at the same distance d_{mask} from the sensor and should have an equal period. As derived in [Veeraraghavan et al. 2007], this distance is given by

$$d_{\text{mask}} = \left(\frac{f_{\theta_R}}{2f_{x_0} + f_{\theta_R}} \right) d_{\text{emitter}}, \quad (15)$$




Figure 7: Mean transmission for proposed two-dimensional tiled-broadband codes. Notice that the Sum-of-Sinusoids tile converges to $\approx 18\%$ transmission for large angular resolutions. In contrast, tiled-MURA codes remain near 50% transmission for any angular resolution desired. As a result, exposure times with MURA tiles will be 2.7 times less than with equivalent Sum-of-Sinusoids masks.

where $f_{x_0} = N_x/(2l_{\text{receiver}})$, $f_{\theta_R} = 1/l_{\text{emitter}}$, and N_x is the desired number of spatial samples in the captured light field.

Now we only need to search for an optimal binary broadband code to be used as the tile in our mask. Fortunately, Modified Uniformly Redundant Arrays (MURA) are a known set of optimal binary broadband codes which have been previously applied for astronomical and medical imaging [Fenimore and Cannon 1978]. In general, such binary patterns are easier to print; for instance, commercial printers from Heidelberg Druckmaschinen AG are capable of producing 5,080 DPI transparencies up to $70\text{cm} \times 50\text{cm}$. In addition, MURA patterns transmit approximately 50% of incident light – reducing exposures by a factor of about 2.7 and eliminating printing quantization artifacts when compared to SoS patterns (see Figure 7). As described by Fenimore and Cannon, these codes are known to be equivalent to a pinhole aperture in X-ray imaging.

For completeness, we describe the specific tiled-MURA attenuation pattern used in our prototype. The two-dimensional MURA occlusion function of prime-dimensions $p \times p$ is given by

$$o_{\text{MURA}}[n, m] = \begin{cases} 0 & \text{if } n = 0, \\ 1 & \text{if } n \neq 0 \text{ and } m = 0, \\ 1 & \text{if } C_p[n]C_p[m] = 1, \\ 0 & \text{otherwise,} \end{cases} \quad (16)$$

where (n, m) are the orthogonal pixel coordinates in the mask plane and $C_p[k]$ is the Jacobi symbol given by

$$C_p[k] = \begin{cases} 1 & \text{if } \exists x, 1 \leq x < k, \text{ s.t. } k = x^2 \pmod{p}, \\ -1 & \text{otherwise.} \end{cases} \quad (17)$$

Our primary contribution is to recognize that a *tiled* array of broadband codes also creates impulse trains in the frequency domain, similar to pinhole arrays and SoS masks. We emphasize that our mask is significantly different from the traditional MURA pattern used in medical and astronomical imaging. In those fields, the goal is to obtain a mask with a *broadband* Fourier transform, while we seek a mask with *impulses* in the frequency-domain. In addition, MURA codes are only defined for specific prime numbers, while our analysis could be used to obtain codes of any length. For example, one could search for the best binary code, of arbitrary length, which minimizes the variance of the magnitude of its Fourier coefficients (similar to [Raskar et al. 2006]). We believe our theoretical analysis of optimal tiled-broadband masks will lead to further research in light field capture and its frequency-domain analysis.

5 Visual Hulls from Shield Fields

To reconstruct the 3D shape of an occluder, we begin by decoding the diffuser-plane image, using the heterodyne decoding

method [Veeraraghavan et al. 2007] to estimate the incident light field. For sufficiently-large angular sampling rates, each sample along the emitter surface will correspond to a small area source. Each 2D slice of the 4D light field, for a fixed angular resolution element, will then contain the shadowgram produced by each emitter element. Unfortunately, our diffuser limited the effective angular resolution to approximately 11×11 . As a solution, we propose an alternate emitter configuration comprised of a uniform array of point sources (see Figure 1). A single light field photograph then contains the shadowgrams produced by each point light with minimal crosstalk between neighboring angular samples (see Figure 5). This effectively solves a key limitation of previous shape-from-silhouette systems [Savarese et al. 2001; Yamazaki et al. 2007] which could only use a single point source at any given instant.

In principle, the 3D shape of the occluder can be recovered from the individual shadowgrams using any visual hull algorithm [Matusik et al. 2001]. As shown in Figure 5, the detected light field can be split into N individual subviews $\{I_1(u_1), \dots, I_N(u_N)\}$, where u_j is a pixel in the j^{th} image and $I_j(u_j)$ is the normalized image intensity. Each of these shadowgrams has a projection equation $u_j = \pi_j(q)$ which, for a 3D point q within the reconstruction volume, defines its mapping q to a position in the j^{th} subview. For reconstruction, each intensity image $I_j(u_j)$ is thresholded to obtain a binary image $p_j(u_j)$ so that, for each point q contained within an occluder, $p_j(\pi_j(q)) = 1$. Notice that, if $p_j(u_j) = 0$, then none of the 3D points q for which $\pi_j(q) = u_j$ are in the object. Since this is true for $j = \{1, \dots, N\}$, we find $p(q) = 1$ for every point q in the object, where

$$p(q) = \prod_{j=1}^N p_j(\pi_j(q)). \quad (18)$$

A sufficient condition for a point q to be outside of the object is that one of the factors in this product equals zero. Since our individual subviews are low resolution, the thresholding operation eliminates useful information about the 3D shape contained in the normalized image intensity. Instead, we set $p_j(u_j) = 1 - I_j(u_j)$ and regard $p(q)$ as a probability density function: if $p(q)$ is small it is very likely that q is outside the object.

As shown in Figure 5, our images exhibit a widely-varying signal-to-noise ratio (SNR) due to the diffuser. If not addressed, these effects prevent the reconstruction of an accurate 3D shape. To solve this problem we estimate confidence images $c_j(u_j)$, using the calibration light field taken without an occluder, so that $c_j(u_j) \approx 1$ for high-SNR pixels and $c_j(u_j) \approx 0$ for low-SNR pixels. We then form a confidence-weighted probability density function given by

$$p'(q) = \prod_{j=1}^N [c_j(u_j) p_j(u_j) + (1 - c_j(u_j))], \text{ for } u_j = \pi_j(q). \quad (19)$$

The visual hull of the occluder is approximated by an isosurface of this probability density. Typical results achieved by our reconstruction method are shown in Figures 1 and 10.

6 Analysis

In this section, we analyze our visual hull capture system with respect to discrete sampling and demodulation noise issues.

Discrete Sampling of Shield Fields: We briefly analyze the effects due to discretely-sampling shield fields. For this analysis we will now consider the shield field $s'(x, y)$ as a function of two coordinates belonging to parallel planes, with the y -coordinate on the emitter plane and the x -coordinate on the receiver plane.

The change of variables $s(x, \theta) = s'(x, y)$ is given by the relation $d_{\text{emitter}} \theta = x - y$. Consider an occluder plane, placed parallel to and between the emitter and receiver planes at a distance z from the receiver plane, with a sinusoidal attenuation pattern $o(\xi) = \cos(\omega \xi)$ of angular frequency $\omega = 2\pi f$. First, we consider impulse sampling for a *discrete shield field* defined as $s[n, m] = s'(n \Delta x, m \Delta y)$, where: n and m are integers, Δx is the receiver plane sampling period, and Δy is the sampling period in the emitter plane. In the continuous domain we have

$$s'(x, y) = o(x - \sigma(x - y)) = \cos(\omega(1 - \sigma)x + \omega\sigma y), \quad (20)$$

where $\sigma = z/d_{\text{emitter}}$ is the normalized depth. Thus the discrete shield field is

$$s[n, m] = \cos([\omega(1 - \sigma)\Delta x]n + [\omega\sigma\Delta y]m). \quad (21)$$

The Nyquist sampling theorem tell us that, in order to prevent aliasing, we need at least two samples per cycle – leading to two inequalities: $\omega(1 - \sigma)\Delta x < \pi$ and $\omega\sigma\Delta y < \pi$. From these constraints we find that the minimum spatial wavelength that can be recovered, at a depth z , is given by

$$T_{\min} = 2 \max\{(1 - \sigma)\Delta x, \sigma\Delta y\}. \quad (22)$$

A more realistic model for the sampling process can be achieved using integral sampling, where

$$s'[n, m] = \frac{1}{\Delta x \Delta y} \int_{(n-\frac{1}{2})\Delta x}^{(n+\frac{1}{2})\Delta x} \int_{(m-\frac{1}{2})\Delta y}^{(m+\frac{1}{2})\Delta y} s(x, y) dx dy. \quad (23)$$

In this case, a straightforward derivation yields

$$s'[n, m] = s[n, m] \frac{\sin(\omega(1 - \sigma)\Delta x/2)}{(\omega(1 - \sigma)\Delta x/2)} \frac{\sin(\omega\sigma\Delta y/2)}{(\omega\sigma\Delta y/2)}, \quad (24)$$

where $s[n, m]$ is the expression for the impulse sampling derived above. Note that, if the angular frequency ω satisfies the impulse sampling constraints, these two additional factors can be compensated for since they are always non-zero, resulting in the same constraint on the minimum wavelength as in the impulse sampling case.

Demodulation Noise: Although pinholes require longer exposures, they do not require demodulation to recover the spectral replicas of the incident light field. Instead, each ray is assumed to be captured at a different pixel (as with lenslets). However, heterodyning using SoS or tile-broadband codes requires frequency-domain demodulation, which will reduce the SNR of recovered shield field. This fact was ignored in [Veeraraghavan et al. 2007]. We analyze the decrease in SNR due to demodulation using simulations. For our prototype system, we consider adding 5% Gaussian noise to the captured image. On average, we find that heterodyne masks, including SoS and tiled-MURA patterns, decrease the received SNR by about 0.5 dB (corresponding to the increased noise in Figure 5).

7 Implementation and Performance

Now we describe our implementation in detail.

7.1 Implementation

Following the design in Section 3, we built the prototype shown in Figure 9. The imaging system is composed of four primary elements: (1) an 8.0 megapixel Canon EOS Digital Rebel XT camera, (2) a $75\text{cm} \times 55\text{cm}$ diffusing screen made of Grafix GFX clear velum paper, (3) three sheets of 3mm thick laminated safety glass, and

(4) a set of interchangeable masks printed at 5,080 DPI on $100\ \mu\text{m}$ polyester base using a Heidelberg Herkules emulsion printer. The diffuser was placed between the first two panes of glass closest to the camera. The various masks were then placed between the second and third glass panes. The diffuser/glass/mask assembly was rigidly mounted to an aluminum frame. The illumination system is composed of three primary elements: (1) a 6×6 array of Philips Luxeon Rebel LEDs, distributed uniformly over a $1.2\text{m}\times 1.2\text{m}$ region, with each 3mm-wide LED producing 180 lumens at 700mA, (2) a regulated power supply, and (3) an aluminum scaffold. The illumination plane was separated from the mask by 65cm.

Since our camera records a 3456×2304 image, we printed pinhole arrays and tiled-MURA masks with $N_x = 151$ and $N_\theta = 11$. That is, the camera oversampled both the spatial and angular dimensions by a factor of two. Substituting these parameters into Equation 15 indicates that the masks should be separated from the diffuser by 2.7mm to recover the shadowgrams produced by each LED.

7.2 Performance

Shadowgram Recovery: After correcting for lens distortion, the primary system limitations arise from the diffusing screen. In practice, diffusers exhibit some amount of subsurface scattering which increases the point spread function (PSF) of the imaging system. Examining the photographs produced using a pinhole array, we estimate that this PSF has a half-width of approximately $300\ \mu\text{m}$ at the diffuser plane, corresponding to a 1.4×1.4 pixel region. In addition, since the sheets of glass cannot be made perfectly flat, the distance between the mask and diffuser varies slowly across the image. Furthermore, refraction due to the multiple glass sheets will result in a per-pinhole barrel distortion (see Figure 5). In the supplementary material, we describe how to compensate for these effects and recover higher-quality shadowgrams, when using pinhole arrays, by recording a calibration image for each point light source.

Visual Hull Reconstruction: A variety of sculptures and household objects were scanned, with typical results shown in Figure 10. We observe that these results primarily exhibit the inherent limitations of using limited-baseline, low-resolution shadowgrams. Since, in our system, we can consider each point light as a camera, then the maximum baseline is the length of the emitter plane. As a result, the reconstructions are elongated perpendicular to the sensor plane, since there are no views to carve away these regions.

Dynamic Scenes: The results included in Figures 1 and 10 indicate that single-frame visual hull reconstruction was achieved. As shown in Figure 5, our system is capable of recovering an array of shadowgrams from a single image. In that figure, the pinhole array required a 30 second exposure, whereas the tiled-MURA mask required only 0.25 seconds. While the resolution of current video camera technology limits the degree to which we could demonstrate real-time capture, we have implemented a preliminary video-rate prototype using a Point Grey Grasshopper video camera capable of recording 1600×1200 8-bit RGB images at 15 fps. Due to the reduced resolution of the video camera, the recovered light field consists of a 6×6 array of 75×50 pixel images. Results for several dynamic scenes are included in the supplementary material.

7.3 Validation

A user of our system will be concerned with the reconstruction accuracy; fortunately, Equation 22 gives one theoretical limit on our system’s resolution (in a plane parallel to the sensor). To confirm this equation, we printed a linear square-wave chirp and recovered its shield field when placed at several depths. As shown in Figure 8, the shadowgrams show aliasing approximately where predicted by




Figure 8: Validation of spatial sampling rate as a function of planar occluder depth. A linear square-wave chirp, over the interval $[0,1,2,0]$ cycles/cm, was placed in the system at several depths. (Top left) Shadowgram when placed 2cm from sensor and (Top right) 16cm from sensor. (Bottom) Theoretical spatial resolution (dashed black line) given by Equation 22. As predicted, features below $\approx 0.5\text{cm}$ cannot be resolved near the sensor. Note that the red and green arrows in the top figures indicate the positions where aliasing should begin, corresponding to the intersection of the red and green lines with the dashed black line in the lower figure.

theory. Note that the spatial resolution is maximized when the object is placed near the center of the volume slightly on the emitter side. This is expected, since the narrow LEDs cause the projected image to enlarge as the object moves towards the emitter (up to the point where the finite width of the LEDs causes blurred shadows).

8 Benefits and Limitations

Benefits: Through shield fields, we have made three primary contributions: (1) provided a deeper understanding of 4D light field modulation by 3D occluders, (2) demonstrated optimal methods to achieve frequency-domain impulse trains using tiled-broadband codes, and (3) shown how these results can be applied to construct a novel visual hull system that is the first to allow dynamic scene capture using shadowgrams. While previous papers, such as Chai et al. [2000] and Durand et al. [2005], presented frequency-domain analyses for the forward rendering problem, this work is the first to focus on the inverse problem of 3D occluder reconstruction. In addition, our tiled-MURA code significantly outperforms the prior Sum-of-Sinusoids pattern [Veeraraghavan et al. 2007] – transmitting almost three times as much light, requiring shorter exposures, and replacing a continuous mask with a binary pattern that uses a lower-cost, larger-format, and higher-accuracy printing process.

Comparison to Existing Visual Hull Systems: Our design is the first single-camera, shadowgram-based system to compete with complex multi-camera systems in the arena of motion capture. Compared to prior shadowgram-based designs, such as [Savarese et al. 2001] and [Yamazaki et al. 2007], ours is the first to allow single-shot dynamic scene capture without temporal or wavelength-domain multiplexing. As diffuser and sensor technology improves, our design will allow hundreds of shadowgrams to be acquired, while the cost and complexity of calibrating and synchronizing similarly large camera arrays becomes prohibitive. Unlike camera arrays or lenslet-based systems, our design supports wavelengths that are difficult to image using lenses (e.g., X-rays). Furthermore, Ng et al. [2005] have shown that lenslet-based designs require precise manufacturing tolerances and are significantly more expensive to produce in large formats. In contrast, our system can be built at low cost by any consumer who owns a single high-resolution camera.




Figure 9: Prototype shield field capture system. An object is placed in the center and simultaneously illuminated from behind by a 6×6 array of LEDs. The shadows are cast onto a thin diffuser after passing through a tiled-MURA attenuation pattern. A single high-resolution camera photographs the diffuser plane. Afterwards, the light field can be recovered from this image, yielding the individual shadowgrams for each LED. The visual hull is estimated using the known calibration and acquired object silhouettes.

Limitations: The shield field representation is similar to the two-plane parametrization of light fields and shares the corresponding limitations of non-uniform sampling. Shield fields cannot represent opacity in **concavities** of objects. A natural extension will be surface shield fields that represent occlusions in a local coordinate system [Wood et al. 2000]. Similarly, visual hull reconstruction cannot recover concave regions. In its current form, our system suffers from several implementation issues, some of which could be overcome with better engineering using higher-quality components. It allows single-shot capture, but at the cost of **reduced spatial resolution**. The diffuser limits the sensor resolution by increasing the PSF by about $300 \mu\text{m}$. One solution is to scale our design to the size of the Blue-C system [Gross et al. 2003], since the **diffuser PSF** will stay constant, but the size of a diffuser-plane pixel will increase. We assume minimal reflections from the object to the diffuser plane. In practice, we found this was not a significant issue, yet **specular objects** could produce artifacts. The mask introduces minor **diffraction** and one must carefully select the mask resolution. Our prototype uses $400 \mu\text{m}$ MURA cells with 3mm between the mask and diffuser, causing diffraction of $8 \mu\text{m}$. The mask and glass panes may introduce artifacts including scattering, refraction, or reflections. Because we record a linear combination of projected patterns, we must radiometrically-calibrate our camera and prevent saturated/underexposed pixels. These non-linearities may introduce errors in our linear inversion. Finally, **heterodyne demodulation** decreases the received SNR due to the decoding process.

9 Future Directions

Shield fields, like light fields, are defined in 4D ray-space. Extensions of light fields can be applied to shield fields. With shield fields, we have the opportunity to develop application-specific data structures similar to opacity hulls of alpha mattes [Vlasic et al. 2003] and surface light fields [Wood et al. 2000]. Further research is required on reconstruction filters, relighting schemes, and efficient rendering strategies using shield fields. Masks also play an important role in imaging outside the visible spectrum in astronomy, scientific imaging, and medical scanning. Shield fields will prove useful for developing advanced imaging methods based on attenu-




Figure 10: Visual hulls from shield fields. (Top row) A 5mm thick spring (shown with a tiled-MURA pattern) and its visual hull reconstruction. (Bottom row) A 2cm thick sculpture (shown with a tiled-pinhole pattern) and its reconstruction. The right column shows four views of visual hull (i.e., the 50%-confidence isosurface).

ating masks. In the future we expect to see volumetric masks with complex, angularly-dependent attenuation functions to find many applications. For example, single-shot decomposition of multiple light sources is useful in digital relighting. Beyond spatial encoding, we can further apply coding in illumination and time.

Our single-shot system is ideally-suited for creating large real-time scanning environments, such as 3D tele-immersion systems [Gross et al. 2003], or for gesture-based interfaces. In addition, emerging technologies support large area emitters and sensors. The recent commercialization of Light Emitting Capacitor (LEC) technology by CeeLite, provides an ideal flat, high-lumen area source. For large area sensors, Sharp has recently started selling flat panel LCDs where each pixel is also a photosensor. New high-resolution consumer-grade cameras, such as the Casio Exilim EX-F1, support 5 megapixel image capture at 60 frames per second. Paired together, one can imagine a new light-efficient shape capture system composed of a cylinder of light sources and detectors surrounding an object – significantly increasing the reconstruction baseline.

10 Conclusions

Occluders are becoming valuable for a range of Computational Photography research problems. We have described a unified representation of occluders in light transport and photography using shield fields. We have provided a complete analysis of their application to the inverse problem of occluder reconstruction, leading to the development of new tiled-broadband codes. Our tiled-coding leads to a simple scanning system. We believe we have presented the first single-camera, single-shot approach to capture visual hulls without moving or programmable illumination. Shield fields can be used for several interesting applications in the future, including efficient shadow computations, volumetric/holographic masks, and for modeling general ray-attenuation effects as 4D light field transformations. By creating a signal processing framework and practical

light-efficient solutions, we hope to inspire research in novel capture devices based on large area emitters and sensors.

Acknowledgements

We thank the anonymous reviewers for their useful suggestions, John Barnwell for assisting in the construction of the prototype system, Takafumi Aoki for preparing the video, and Jay Thornton and Joseph Katz for their support. We would also like to thank Ashok Veeraraghavan, Ankit Mohan, Cyrus Wilson, and Daniel Crispell for helpful discussions. This work has been partially funded by the National Defense Science and Engineering Graduate (NDSEG) Fellowship and by Mitsubishi Electric Research Laboratories.

References

- ACCORSIA, R., GASPARINI, F., AND LANZA, R. C. 2001. Optimal coded aperture patterns for improved SNR in nuclear medicine imaging. *Nuclear Instruments and Methods in Physics Research A* 474, 3, 273–284.
- ADELSON, T., AND WANG, J. 1992. Single lens stereo with a plenoptic camera. *IEEE TPAMI* 14, 2, 99–106.
- BOWHAUS, INC., 2007. BowHaus LVT specs and resolutions. <http://www.bowhaus.com/services/lvtspecs.php4>.
- CHAI, J.-X., TONG, X., CHAN, S.-C., AND SHUM, H.-Y. 2000. Plenoptic sampling. In *SIGGRAPH*, 307–318.
- DEBEVEC, P., HAWKINS, T., TCHOU, C., DUKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *SIGGRAPH*, 145–156.
- DURAND, F., HOLZSCHUCH, N., SOLER, C., CHAN, E., AND SILLION, F. X. 2005. A frequency analysis of light transport. *ACM Trans. Graph.* 24, 3, 1115–1126.
- FARID, H. 1997. *Range Estimation by Optical Differentiation*. PhD thesis, University of Pennsylvania.
- FENIMORE, E., AND CANNON, T. 1978. Coded aperture imaging with uniformly redundant arrays. *Appl. Optics* 17, 3, 337–347.
- GEORGIEV, T., ZHENG, K. C., CURLESS, B., SALESIN, D., NAYAR, S., AND INTWALA, C. 2006. Spatio-angular resolution tradeoffs in integral photography. In *EGSR*, 263–272.
- GOODMAN, J. W. 1996. *Introduction to Fourier Optics*.
- GORTLER, S. J., GRZESZCZUK, R., SZELISKI, R., AND COHEN, M. F. 1996. The lumigraph. In *SIGGRAPH*, 43–54.
- GROSS, M., WÜRMLIN, S., NAEF, M., LAMBORAY, E., SPAGNO, C., KUNZ, A., KOLLER-MEIER, E., SVOBODA, T., GOOL, L. V., LANG, S., STREHLKE, K., MOERE, A. V., AND STAADT, O. 2003. blue-c: a spatially immersive display and 3d video portal for telepresence. *ACM Trans. Graph.* 22, 3, 819–827.
- ISAKSEN, A., McMILLAN, L., AND GORTLER, S. J. 2000. Dynamically reparameterized light fields. In *SIGGRAPH*, 297–306.
- IVES, H. E. 1928. Camera for making parallax panoramagrams. *J. Opt. Soc. of America* 17, 435–439.
- LEVIN, A., FERGUS, R., DURAND, F., AND FREEMAN, W. T. 2007. Image and depth from a conventional camera with a coded aperture. *ACM Trans. Graph.* 26, 3, 70.
- LEVOY, M., AND HANRAHAN, P. 1996. Light field rendering. In *SIGGRAPH*, 31–42.
- LIPPMANN, G. 1908. Epreuves reversible donnant la sensation du relief. *Journal of Physics* 7, 4, 821–825.
- MATUSIK, W., BUEHLER, C., AND McMILLAN, L. 2001. Polyhedral visual hulls for real-time rendering. In *EGSR*, 115–126.
- NAYAR, S. K., KRISHNAN, G., GROSSBERG, M. D., AND RASKAR, R. 2006. Fast separation of direct and global components of a scene using high frequency illumination. *ACM Trans. Graph.* 25, 3, 935–944.
- NG, R., LEVOY, M., BRÉDIF, M., DUVAL, G., HOROWITZ, M., AND HANRAHAN, P. 2005. Light field photography with a hand-held plenoptic camera. Tech. rep., Stanford University.
- NG, R. 2005. Fourier slice photography. *ACM Trans. Graph.* 24, 735–744.
- RAMAMOORTHY, R., KOUDELKA, M., AND BELHUMEUR, P. 2005. A fourier theory for cast shadows. *IEEE TPAMI* 27, 2, 288–295.
- RAMAMOORTHY, R., MAHAJAN, D., AND BELHUMEUR, P. 2007. A first-order analysis of lighting, shading, and shadows. *ACM Trans. Graph.* 26, 1, 2.
- RASKAR, R., AGRAWAL, A., AND TUMBLIN, J. 2006. Coded exposure photography: motion deblurring using fluttered shutter. *ACM Trans. Graph.* 25, 3, 795–804.
- SAVARESE, S., RUSHMEIER, H., BERNARDINI, F., AND PERONA, P. 2001. Shadow carving. In *ICCV*, 190–197.
- SOLER, C., AND SILLION, F. X. 1998. Fast calculation of soft shadow textures using convolution. In *SIGGRAPH*, 321–332.
- TALVALA, E.-V., ADAMS, A., HOROWITZ, M., AND LEVOY, M. 2007. Veiling glare in high dynamic range imaging. *ACM Trans. Graph.* 26, 3, 37.
- THORNER, K., AND JACOBS, D. 2001. Cast shadows and linear subspaces. Tech. rep. TR-2001-100, NEC.
- VEERARAGHAVAN, A., RASKAR, R., AGRAWAL, A., MOHAN, A., AND TUMBLIN, J. 2007. Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocusing. *ACM Trans. Graph.* 26, 3, 69.
- VLASIC, D., PFISTER, H., MOLINOV, S., GRZESZCZUK, R., AND MATUSIK, W. 2003. Opacity light fields: Interactive rendering of surface light fields with view-dependent opacity. In *Symposium on Interactive 3D Graphics (i3D)*, 65–74.
- WILBURN, B., JOSHI, N., VAISH, V., TALVALA, E.-V., ANTUNEZ, E., BARTH, A., ADAMS, A., HOROWITZ, M., AND LEVOY, M. 2005. High performance imaging using large camera arrays. *ACM Trans. Graph.* 24, 3, 765–776.
- WOOD, D., AZUMA, D., ALDINGER, K., CURLESS, B., DUCHAMP, T., SALESIN, D., AND STUETZLE, W. 2000. Surface light fields for 3d photography. In *SIGGRAPH*, 287–296.
- YAMAZAKI, S., NARASIMHAN, S. G., BAKER, S., AND KANADE, T. 2007. Coplanar shadowgrams for acquiring visual hulls of intricate objects. In *ICCV*, 1–8.
- ZHOU, K., HU, Y., LIN, S., GUO, B., AND SHUM, H.-Y. 2005. Precomputed shadow fields for dynamic scenes. *ACM Trans. Graph.* 24, 3, 1196–1201.
- ZWICKER, M., MATUSIK, W., DURAND, F., AND PFISTER, H. 2006. Antialiasing for automultiscopic displays. In *EGSR*, 1–10.