Rendering Wave Effects with Augmented Light Field

Se Baek Oh,[1] Sriram Kashyap,[2] Rohit Garg,[2] Sharat Chandran,[2] and Ramesh Raskar[1]

[1]MIT, USA,   [2]IIT Bombay, India

Abstract

Ray–based representations can model complex light transport but are limited in modeling diffraction effects that require the simulation of wavefront propagation. This paper provides a new paradigm that has the simplicity of light path tracing and yet provides an accurate characterization of both Fresnel and Fraunhofer diffraction. We introduce the concept of a light field transformer at the interface of transmissive occluders. This generates mathematically sound, virtual, and possibly negative-valued light sources after the occluder. From a rendering perspective the only simple change is that radiance can be temporarily negative. We demonstrate the correctness of our approach both analytically, as well by comparing values with standard experiments in physics such as the Young's double slit. Our implementation is a shader program in OpenGL that can generate wave effects on arbitrary surfaces.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three–Dimensional Graphics and Realism—Raytracing I.3.6 [Computer Graphics]: Methodology and Techniques—Graphics Data Structures and Data Types

1. Introduction

Ray–based methods are strong contenders for efficiently and accurately capturing the myriad ways light interacts with matter. Starting from the eighties, backward ray tracing, with recursive secondary rays, enabled the efficient depiction of shadows, reflection, and refraction. Morphing ideas from radiosity and signal processing enabled effects such as soft shadows, lens blurring, and motion blur. Yet another significant jump is the realization that caustic effects are best obtained – demonstrated by photon mapping and bidirectional ray tracing – by source level tracing from the light source, rather than backward tracing from the eye. Meanwhile the Light Field (LF) model for understanding light transport had led to a range of new techniques and applications such as digital refocusing, depth estimation, synthetic apertures, and glare reduction.

Given this background, it would be more complete to also envelope another frontier of light interaction with matter – diffraction. The goal of this paper is to provide a source level description of an augmented light field that enables the unified treatment of light effects – the traditional effects mentioned above, and the less explored wave effects. Prior researchers have implemented some of the large variety of wave effects; however these are hybrid methods that keep track of the phase of the wavefront while propagating rays [Mor81]. They are cumbersome and are not plug–and–play compatible with well understood ray tracing paradigms.

1.1. Contributions

We introduce the Augmented Light Field (ALF) framework, which is a new theoretical way of modeling wave phenomena [OBR08, ROAZ09]. The ALF is an upgraded ray–based representation that, put simply, mimics the Wigner Distribution Function (WDF) analysis popular in the optics community. Since the ALF is still essentially ray–based (but with phase information of wavefronts already encoded) the ALF can be integrated easily into existing ray–based methods. Specific technical contributions are as follows:

- We show that the ALF can model diffraction and interference rigorously despite being a pure ray–based formulation.
- We show that the proposed method can predict both Fresnel and Fraunhofer diffraction without any change in the representation.
- We show the compatibility of ALF in ray–based implementations and provide an OpenGL implementation.
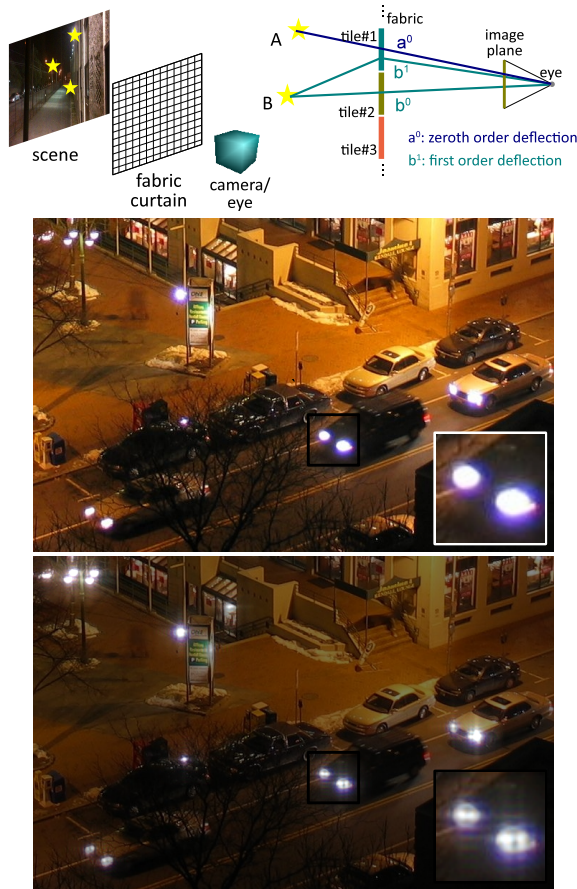
Figure 1: Effects created by a window fabric due to diffraction. (Top left) Looking at a scene from behind a fabric with fine tiled pattern. (Top right) Our framework predicts zeroth, first, and higher diffraction orders for intricate patterns. We model the fabric as tiled gratings and render diffraction from each tile independently. (Middle and Bottom) Rendering without and with the fabric. Notice the effects around bright sources.

1.2. Scope and Limitations

The goal of the paper is three fold. First, we present the theoretical foundations of the ALF. Next, we demonstrate how the upgraded ray representation can handle wave effects. Finally, we demonstrate the benefits using instructive rendering examples. Building a complete rendering system is beyond the scope of this paper. Our formulation is explained in flat land (*i.e.*, in the plane of the paper) for notational simplicity, where the LF, ALF, and WDF are 2D functions. However, the analysis applies to the real 3D world and corresponding 4D functions in a straightforward manner. All our rendering examples are realized in the 3D world and on arbitrary surfaces as shown in Sec. 5.3. The examples are cho-
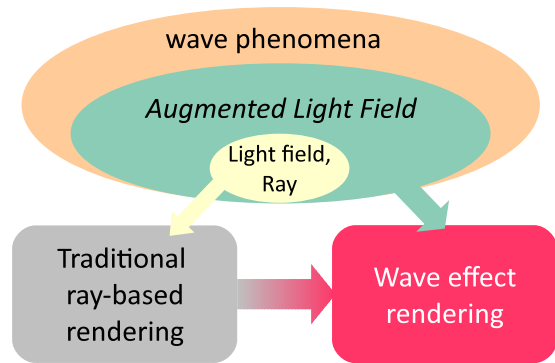

Figure 2: The Augmented Light Field framework supports diffraction and interference. It shares the simplicity of ray–based representation. Modeling wave–phenomenon typically require Fourier Optics but we exploit the ALF and implement such effects in GLSL.

sen to be simple but representative cases commonly used in wave optics, enabling us to validate our results. Further, the diffractive occluders are chosen to be symmetric, enabling us to easily move from 2D to 3D using separability. All the diffractive occluders are transmissive but our framework also can be applied to reflective occluders by simple coordinate inversions.

Paraxial approximation: We only consider light propagation in the paraxial region, and neglect any non–linear effect including singularities. This assumption is essential because the radiance defined by the WDF is preserved only in the paraxial region [LW93, Alo01]. To describe the light interaction in the non–paraxial region, a different type of phase–space representation should be used such as the Angle–Impact Wigner Distribution Function [WAF99, PA07].

Polarization: Our current implementation does not address polarization properly because it is valid for linearly polarized light. However, it can be extended to handle different polarization states with the coherency matrix approach [Alo04, BW99].

Coherence: For simplicity, the concept of the ALF framework is described with coherent light. The ALF can be extended since it is based on the WDF, which can be applied to both coherent and incoherent light [Bas09, ZL09].

1.3. Related work

Light Propagation in Optics: Walther proposed the concept of generalized radiance [Wal73] and there has been extensive research on the subject. The notable one is the Wigner Distribution Function [Bas09], where light is described in terms of local spatial frequency. The WDF has been exploited in a variety of analysis and design problems in optics: 3D display [FMCJS06], digital
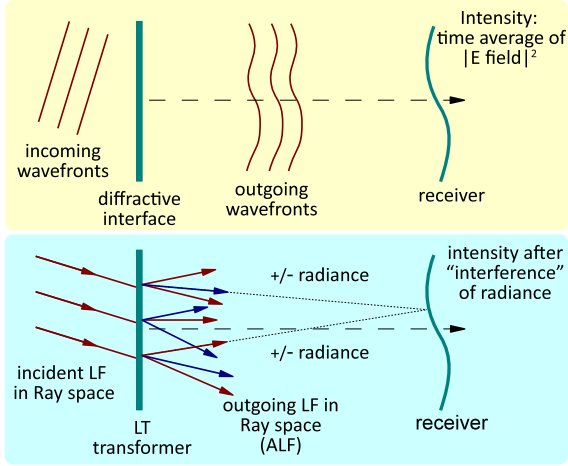
Figure 3: *The ALF can represent challenging wave-phenomena. (Top) Traditional Fourier Optics tools can predict intensities after diffraction, propagation, and interference. (Bottom) ALF can support the same via (i) potentially negative radiance and (ii) lightfield transformer. The constructive or destructive "interference" effect is for free when the radiance is integrated.*

holography [MMH*06, SJ04], generalized sampling problems [SJ04], and superresolution [WMZ98]. Although the WDF provides rigorous models of light propagation, the WDF has not been used for rendering to the best of our knowledge. Recently the important connection between the LF and WDF was made by Zhang and Levoy [ZL09]; the concept of observable light field, which is the ray–representation limited subset of the WDF (the smallest ellipse in Fig. 2), was introduced. Indirectly, they have described the space of effects spanned by today's photon and ray-based renderings. Our work goes beyond because we can represent wave–like effects (the middle ellipse in Fig. 2). The key idea is to allow these enhanced abilities while preserving the simplicity of ray-like operations. However, our framework currently does not support all wave-phenomena such as vectorial diffraction and evanescent waves.

Wave-based Rendering: Moravec proposed a wave model to render complex light transport efficiently [Mor81]. Ziegler et al. developed a new wave–based framework [ZCG08]. As light propagates, the phase change due to the optical path length is tracked and generates constructive and destructive interference. Also, Stam presented a technique to model reflective diffraction using shaders [Sta99].

In principle, keeping track of the optical path length difference (OPD) of all possible rays should be sufficient to model diffraction rigorously. Our intention is to show that the proposed method also provides physically correct results *without computing optical path difference*. A further conceptual difference between these two methods should be noted.

The OPD method "knows" the effect of the rays only when it computes the (difference in) path lengths created by occluders. In contrast, the ALF "preprocesses" the occluders and produces source level information of the light field after the occluder, "waiting to be sampled by the world." Radiance is temporarily negative in our paradigm, but the mathematics ensures that at the end, samples in the real world end up with non–negative intensity [Bas09].

Light Propagation in Graphics: The light field is a parametrized plenoptic function describing radiance of rays in position–angle space [LH96, GGSC96]. Ray–based image synthesis is computationally efficient and produces physically realistic results for many optical elements. This geometric optics approach is also used in optical system design and analysis of various optical phenomena. Light transport, shadow casting, and light field in frequency domain have been studied by Chai et al. [CCST00], Isaksen et al. [IMG00], Veeraraghavan et al. [VRA*07], and Durand et al. [DHS*05].Ray tracing and photon mapping are based on geometric optics theory of propagation.

2. Ray and Wave Effects

Before we introduce the ALF framework, we briefly present the WDF and how it models diffraction because the ALF is formulated based on the WDF.

2.1. Light field and Wigner Distribution Function

In the ray–based LF representation, the plenoptic function is parameterized with four variables (x, y, θ and ϕ) indicating the position and direction of the ray in 3D. Rays propagate straight without bending in a homogeneous medium and exhibit refraction and reflection at an interface between two different media. In wave optics, light is represented as an electromagnetic wave, where both amplitude and phase describe light behavior governed by Maxwell's equations. The wavefront is a surface of a constant phase of the electromagnetic wave, and rays are defined to be perpendicular to the wavefront as in Fig. 4.

The WDF for an input $g(x)$, which can be either electric (or magnetic) field or complex transparencies of thin optical elements, is defined as

$$\mathcal{W}(x, f_l) = \int g(x + \tfrac{x'}{2}) g^*(x - \tfrac{x'}{2}) e^{-i2\pi x' f_l} dx', \quad (1)$$

where x represents the spatial coordinate and f_l indicates the local spatial frequency, $*$ denotes complex conjugate, and the value of the WDF is often called *generalized radiance*. It is known that the local spatial frequency f_l of the wavefront is related to the ray propagation angle θ by $f_l = \theta/\lambda$ in the paraxial region, where λ is the wavelength [Goo05]. Hence, both LF and WDF represent light transport in the position–angle space. Note that as shown in Fig. 5 the LF and WDF
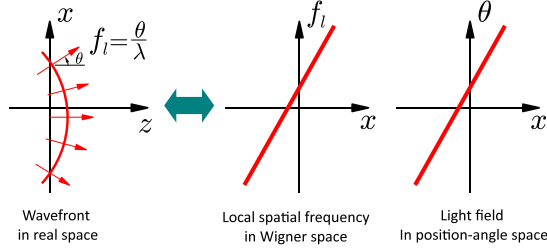
Figure 4: The phase of a wavefront is related to the angle of corresponding rays. The phase is represented as local spatial frequency in the WDF representation. (Left) Illustration of a spherical wavefront and rays. Rays are normal to the wavefront. (Right) The WDF is a local spatial frequency spectrum. The WDF of the spherical wave at a given z in Wigner coordinate space is similar to the ALF in the position-angle coordinate space. Propagation angles of rays are encoded in the local spatial frequency.
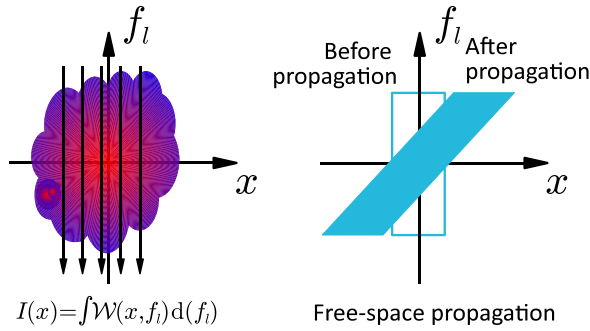

Figure 5: The benefit of the ALF is that image formation and propagation are identical in the WDF, LF and ALF. (Left) Image formation is via intensity computation which is a projection along the angle or spatial frequency dimension, *i.e.* integration at receiving patch along all incoming angles. (Right) The free–space propagation is via the x–shear transform in the position–angle space.

exhibit similar properties: 1) projection of the WDF along the f_l–axis yields the intensity and 2) free–space propagation is illustrated as the x–shear transform.

2.2. Ray Limitations

Since a ray always has a positive radiance, has no concept of phase, and propagates along a straight line in a homogeneous medium, pure ray–based representations cannot model any phase sensitive phenomena, such as light transport associated with diffraction or phase–sensitive optical elements (*i.e.*, phase gratings or holograms). To demonstrate a fundamental limitation of ray–based representations, we revisit the Young's experiment (two pinholes illuminated by a laser)
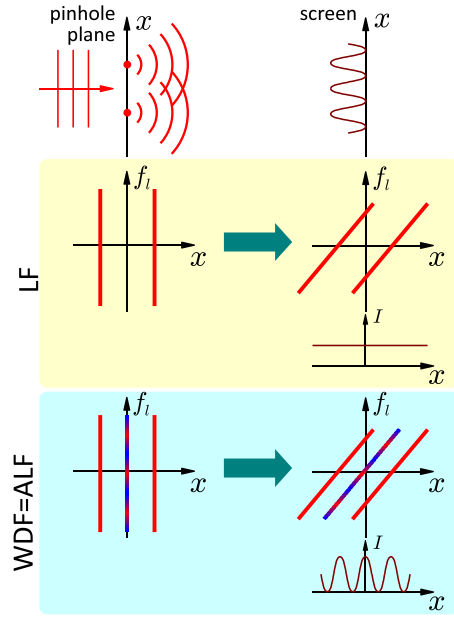

Figure 6: Unlike LF, WDF and ALF can support diffraction and intereference. Consider the double slit experiment in flat land. (Top) Diffracted wavefronts at pinholes propagate to the screen to form a sinusoidal pattern. (Middle–left) Even if we include diffracted rays in the LF representation immediately after the pinholes, (Middle–right) after propagation, we shear the LF along the x–direction and form the image by taking a vertical projection. The computed intensities do not show sinusoidal patterns. (Bottom) The WDF and ALF representations introduce a virtual projector at the mid–point between pin-holes, with positive and negative radiance rays. After shearing and projection, we can successfully produces the sinusoidal interference fringes. The WDF and ALF have three components, two corresponding to the pinholes and the other representing the interference. (Color code; red: positive, blue: negative)

and illustrate the interference via both the WDF and LF in Fig. 6.

If we use the WDF paradigm, we obtain the WDF of the two pinholes from Eq. (1) as

$$\mathcal{W}(x, f_l) = \delta(x - a) + \delta(x - b)$$
$$+ 2\delta\left(x - \frac{a+b}{2}\right)\cos\left(2\pi\left[a - b\right]f_l\right), \quad (2)$$

where infinitesimally small pinholes are located at $x = a$ and b. (As we will describe later, Eq. (2) also represents the WDF of diffracted light when the pinholes are probed by an on–axis plane wave.) Note that three components exist in Eq. (2): the first two terms correspond to the two pinholes, where two point sources emit rays along all different directions at $x = a$ and $x = b$, and the third term represents the

interference. The third term, often called as interference or cross term, has been obtained by the mathematical manipulation of the WDF [Cas08]. In contrast, the LF representation has the first two terms only. The third term is obviously not expected in the LF because there is no physical light source.

The beauty of the WDF representation can be further appreciated by noting that the interference term is oscillatory and can be negative. Since the intensity is the projection of the WDF along the f_l–direction, the intensity of the interference term is zero immediately after the pinhole. However, as light propagates to a screen (we can think of any object as a screen) the cross term plays a significant role in creating interference. In other words, the WDF is sheared along the x–direction by the propagation, where the sheared interference term of the WDF leads to intensity modulation as shown in Fig. 6. The LF model expects no interference fringes because there are only two point sources producing uniform intensity.

3. Augmented Light Field

How can we use the concept of WDF and bring it closer to traditional rendering? The answer is to mimic the WDF and introduce the concept of the ALF where negative radiance is permitted. The key is to figure out how and where to introduce negative radiance.

The ALF is an enhanced representation of ray–space beyond traditional light field. Beyond representation, operations on ALF, such as propagation, scattering (including diffraction), gathering for image formation (including interference), are supported via minor modifications of similar operations on traditional light fields. The ALF formulation is mathematically similar to WDF. However, the ALF representation interprets the WDF with two unique concepts: 1) In representation, the ALF introduces the notion of a virtual projector with potentially negative radiance rays. 2) In operation, the ALF also provides a set of transformer functions. The ALF intends to avoid complicated Fourier optics terminology and makes it compatible with traditional ray–space formulations like photons and lightfields.

3.1. Free Space Propagation

The first task is to think about situations where we need to make little or no modification to traditional light field, and that is in free space. As light propagates in free space, the WDF exhibits the x–shear transform as shown earlier in Fig. 5. Since the x–shear transform is sufficient to describe the free–space propagation in the paraxial region, there is no change; we assert that the ALF is the same as the LF. Note, however, that in the non–paraxial region or an inhomogeneous medium, the x–shear transform needs extra correction factors [Bas09].

3.2. Virtual projector and negative radiance

To model light transport associated with diffraction and interference, we adopt the interference term from the WDF and regard it as a virtual projector. Hence, the virtual projector always emits rays whose radiance is varying along the angle and could be negative. Although radiance can be temporally negative, the projection of the WDF along any arbitrary direction yields non–negative intensity [Bas09].

We can formally derive where to place these light sources and how much radiance they should contribute by computing the WDF of an occluder. In the special case of the two–pinhole experiment, the location of the virtual projector is simply the mid point of the two pinholes and its radiance varies along the θ–axis as Eq. (2). In the general case, the practical technique for locating the virtual projectors is described by *light field transformers* in the next section.

3.3. Light field transformer

As we mentioned earlier, the WDF can be defined for light as well as optical elements – occluders in this paper. The WDFs of occluders describe the light transport before and after the occluders. In the ALF framework, the WDFs of occluders are called Light Field Transformers. Hence, the light field transformer contains the information on where we should put virtual projectors and what radiance we should assign. Visualized this way, the light field transformers for many canonical optical elements can be easily computed as we demonstrate in Table 1.

The natural question of what happens when the occluders are arbitrary and not available in the table is easily answered by doing an explicit computation of the convolution either numerically or symbolically. For example, for a coded aperture whose transmittance is $t(x)$, the LF transformer is given by

$$LF_t(x, \theta) = \int t\left(x + \frac{x'}{2}\right) t^*\left(x - \frac{x'}{2}\right) e^{-i2\pi\frac{\theta}{\lambda}x'} dx' \quad (3)$$

Next we present how to use the light field transformer. In the ALF model as shown in Fig. 7, an input ALF is incident on an occluder and an outgoing ALF $L_2(x_2, \theta_2)$ is produced. For thin occluders, the incident and outgoing rays have identical position in space. Also most thin occluders exhibit angle shift invariance in the paraxial region, where the outgoing ALF is rotated in the same fashion as the incident ALF (see Fig. 7). Thus the incident and outgoing ALF can be related as

$$L_2(x, \theta_2) = \int T(x, \theta_2 - \theta_1) L_1(x, \theta_1) d\theta_1, \quad (4)$$

where T is the light field transformer of an occluder. Equation (4) involves a multiplication along the x–direction but a convolution along the θ–direction. Note that the ALF of an infinitely extended plane wave propagating along the optical axis is $L(x, \theta) = \delta(\theta)$; hence, the light field transformer of an
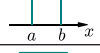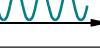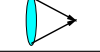
occluder $t(x)$	$T(x,\theta)$
one pinhole $\delta(x-x_0)$	$\delta(x-x_0)$
two pinholes $\delta(x-a)+\delta(x-b)$	$\delta(x-a)+\delta(x-b)+2\delta\left(x-\frac{a+b}{2}\right)\cos\left(\frac{2\pi}{\lambda}(a-b)\theta\right)$
rectangular aperture $\text{rect}\left(\frac{x}{A}\right)$	$2A\Lambda\left(\frac{x}{A/2}\right)\text{sinc}\left([2A-4|x|]\frac{\theta}{\lambda}\right)$
finite two pinholes $\text{rect}\left(\frac{x-a}{A}\right)+\text{rect}\left(\frac{x+a}{A}\right)$	$2A\Lambda\left(\frac{x-a}{A/2}\right)\text{sinc}\left([2A-4|x-a|]\frac{\theta}{\lambda}\right)+$ $2A\Lambda\left(\frac{x+a}{A/2}\right)\text{sinc}\left([2A-4|x+a|]\frac{\theta}{\lambda}\right)+$ $4A\Lambda\left(\frac{x}{A/2}\right)\text{sinc}\left([2A-4|x|]\frac{\theta}{\lambda}\right)\cos\left[\frac{2\pi}{\lambda}\theta(2a)\right]$
sinusoidal amplitude grating $\frac{1}{2}+\frac{m}{2}\cos\left(\frac{2\pi}{p}x\right)$	$\frac{1}{4}\left[\left\{1+\frac{m^2}{2}\cos\left(\frac{2\pi}{p}2x\right)\right\}\delta(\theta)+m\cos\left(\frac{2\pi}{p}x\right)\right.$ $\left.\times\left\{\delta\left(\theta-\frac{\lambda}{2p}\right)+\delta\left(\theta+\frac{\lambda}{2p}\right)\right\}+\frac{m^2}{4}\left\{\delta\left(\theta-\frac{\lambda}{p}\right)+\delta\left(\theta+\frac{\lambda}{p}\right)\right\}\right]$
binary amplitude grating $\sum_{q=-\infty}^{\infty}\text{sinc}(\alpha q)e^{i2\pi\frac{x}{p}q}$	$\alpha^2\sum_{q_1,q_2=-\infty}^{\infty}\text{sinc}(\alpha q_1)\text{sinc}(\alpha q_2)e^{i2\pi\frac{x}{p}(q_1-q_2)}\delta\left(\frac{\theta}{\lambda}-\frac{q_1+q_2}{2p}\right)$
lens (focal length f) $\exp\left(-i\frac{\pi}{\lambda}\frac{x^2}{f}\right)$	$\delta\left(\theta+\frac{x}{f}\right)$

Table 1: LF transformer of canonical occluders are presented in 2D; extension to 3D for many interesting cases is enabled by an appropriate separable product. For a rectangular aperture, $\text{rect}\left(\frac{x}{A}\right)=1$ if $|x|<A$, otherwise $\text{rect}\left(\frac{x}{A}\right)=0$. $\Lambda(x)=1-|x|$ if $|x|\leq 1$, otherwise $\Lambda(x)=0$. For finite two pinholes, A is the size of the pinholes. For gratings, p is a grating pitch. For binary amplitude grating, α represents the duty cycle.
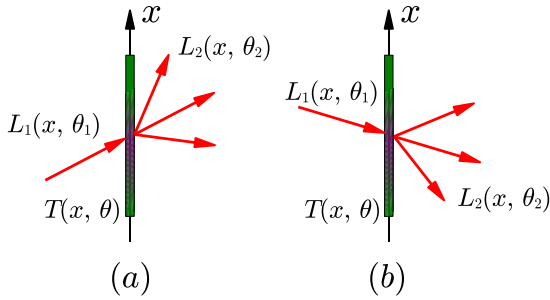

Figure 7: Outgoing ALF as a function of incoming angle. Angle shift invariance indicates that as the incident ray rotates the outgoing ALF simply rotates in the same fashion.
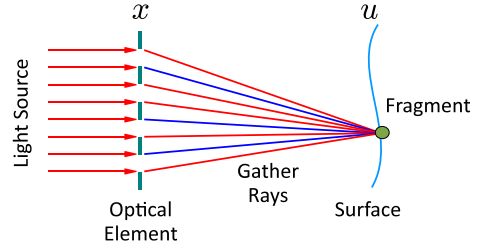

Figure 8: Implementation in OpenGL exploits shader fragments. Rays from a light source diffract at an optical element and the diffracted rays are gathered at a patch on a receiving surface of arbitrary geometry and position.

occluder can also be interpreted as the outgoing ALF when the occluder is probed by an infinitely extended plane wave.

4. Rendering with Augmented Light Field

Because the ALF is simply an instance of light field (albeit with negative radiance), any sampling based strategy can be used to rendering. We first describe an OpenGL shader based method which we implement, and then explore other possible approaches.

4.1. Shader Implementation

For simplicity, we first describe a proof-of-concept implementation limited to a light source wavefront, an occluder that lives in 3D, and a screen of arbitrary shape modeled as polygons in 3D. In this context, we can show the impact of the ALF in a backward manner, *i.e.*, from the eye to screen to the optical element. We render the screen using OpenGL. As polygons get rasterized, the effects of the ALF are accommodated in their fragment shader. For each fragment, we sum the total radiance (positive or negative) emanating from all the samples on the optical element. More specifically, consider the setup shown in Fig. 8. We compute the light field

as a two–plane parameterization of ray–space. In Fig. 8 the two "planes" are the grating and the receiving object, which could be the eye, or any object modeled using polygons. We use Eq. (4) (but now in 3D) for relating the output light field given the light source resulting in the shader pseudocode for the net radiance $L(u_i)$ as seen by the eye for fragment u_i.

> **for** each sample point $p_j = (x, y)$ on occluder **do**
> (θ_1, ϕ_1) = Direction of input ray from source to p_j
> (θ_2, ϕ_2) = Direction of output ray from p_j to u_i
> $L(u_i) + = T(x, y, \theta_1 - \theta_2, \phi_1 - \phi_2) \times$ SourceRadiance
> **end for**

If N fragments are rendered for the screen and we take M samples on the optical element, then the work to be done is $O(MN)$, all of which happens in parallel on the GPU. The function $T(x, y, \theta, \phi)$ described above incorporates Eq. (4). It can be numerically computed offline and stored in a look up table or computed on the fly using the analytical formulae. We avoid the high space requirements of $T(x, y, \theta, \phi)$ because our transformers are separable and can be factorized as $T(x, y, \theta, \phi) = T_1(x, \theta) \times T_2(y, \phi)$. In essence, then, the pseudocode represents the mathematical equation

$$L(u_i) = \sum_x S(x, u_i) = \sum_x \int T(x, \theta_1 - \theta_2) L_1(x, \theta_1) d\theta_1 \quad (5)$$

Note that in practice, some fragments end up with a very small negative intensity. However, with sufficient sampling, we can clamp them to zero without noticeable loss of quality.

4.2. Discussion

The advantage of the shader approach is that it is aligned with existing programmable graphics hardware, which support general gather operations, but not scatter operations. This makes integrating our method with existing renderers less invasive. Also, shaders enable diffraction patterns to appear on arbitrary surfaces. These surfaces can move from near–zone to far–zone, and the rendering code does not change.

On the flip side, the approach so far described is for the case of a single occluder (represented as a plane). As we stack up several occluders that are not necessarily aligned, the backward approach is not computationally feasible since we have to hold in memory the lookup table that can grow in an exponential manner. For multiple scatterings, caustic effects and other global illumination effects, "forward" photon mapping may be more appropriate, and we give a brief description here on how our theory could be useful.

Since the ALF representation is source based, we can express the rendering in the photon mapping step as follows. Recall the regular expression syntax for rays [Hec90] where all possible ray transport phenomenon is of the type `L(S|D)*E` with the alphabet (`S`) for specular, (`D`) for diffuse surfaces, and with `L` and `E` denoting light and eye respectively. For photon mapping with the ALF, we hypothesize new optically active elements with the acronym `I` which

can cause "interference". In the the first stage of classic photon mapping, we exploit the light field transformer to decide the impact in direction (θ_2) for each photon arriving in direction (θ_1) at a point $\vec{x}$ on the element.

- Assign output angle (θ_2) for the deflected photon using stochastic sampling and Russian roulette.
- Assign radiance by using Equation (4).

Notice that this interaction with `I` is no different in spirit, from the other two cases of `D` or `S` in photon tracing. The rest of the rendering remains the same.

In this paper our goal is to show verifiable wave effects in a practical rendering pipeline and not to incur the full cost of global illumination as in photon mapping. The ALF is similar to photon production and is a source based approach where the rays or photons simply propagate, whereas previous methods, *e.g.*, based on keeping track of OPD, are receiver based phenomena.

A source based approach is suitable when the receiving geometry is dynamic and the diffracting geometry undergoes only a rigid transform. For instance, in the wavy mesh example (Fig. 13), the OPD method constantly updates the path differences, whereas in the ALF method, the source light field still remains the same. A small difference in the position in the screen location may make a big difference in the observed phenomena. We note that projective textures are not suitable for diffractive renderings as the pattern after interference can morph dramatically from near field to far field. The ALF can handle rigid transformations in incident light angle and grating more efficiently. If the incident light angle changes, the projected 4D illumination field rotates correspondingly. This insight is appreciated in the OPD method only after a screen is placed and the image is computed.

5. Results

We verified our rendering approach in the context of typical diffractive elements and also on arbitrary dynamic objects.

5.1. Young's double slit experiment

We demonstrate Young's double slit experiment with our implementation, in Fig. 9. The spatial frequency of the interference fringe depends on the separation distance between the slits and the size of the slits decides the overall envelope shape of the fringe. Since diffraction angle is wavelength dependent, we observe color smearing in Fig. 9b.

5.2. Near–zone and Far–zone diffraction patterns from a square aperture

We render the diffraction pattern from a square aperture in both near and far–zone. In Fourier optics, the near–zone and far–zone are often referred to Frensel and Fraunhofer region

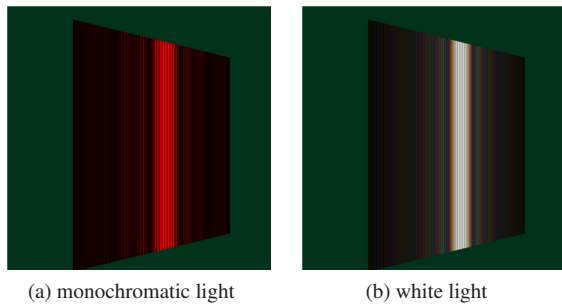(a) monochromatic light   (b) white light

Figure 9: The ALF–based rendering provides correct models of Young's double slit experiments. Two slits are illuminated by monochromatic light (Left) or white light (Right). As we change the slit size as well as the slit separation, the interference patterns also vary and we validated with Fourier optics model. This particular snapshot is taken when the slit size is 0.035 mm, the slit separation is 0.2 mm, the slit–screen distance is 1 meter, and the screen size is 200 mm × 200 mm. In the case of white light, note color smearing at the boundary of the center peak, where red is observed. The first diffraction order starts with blue since it diffracts less than red.

and distinguished by Fresnel number which is a dimensionless parameter related to the size of the aperture, the distance from the aperture to the observation plane, and the wavelength. If the propagation distance is sufficiently larger than the area of the aperture divided by the wavelength, then the observation plane is said to be in the far–zone. (Here near–zone is not near–field zone in optics.) The diffraction in the near–zone and far–zone have different characteristics with respect to the propagation distance. In the near–zone, the overall shape of the pattern is similar to the aperture shape but it has many high frequency signals due to diffraction, and the pattern changes significantly as light propagates. In contrast, the far–zone pattern is proportional to the Fourier transform of the aperture and the pattern changes only its scale as the distance varies. Figure 11 shows the intensity at the center pixel of the near–zone diffraction pattern. The intensity varies depending on the distance and this is an evidence of the wave–property of light. This phenomenon or its complimentary pattern (diffraction from an obstacle) is often referred to Poisson's blinking spot or Arago's spot in the optics community [Hec02].

5.3. Demonstration of wave effects on arbitrary and dynamic surfaces

Our rendering is compatible with traditional rendering. In Fig. 12 we show the effect of light coming from two different sources by placing a box of width 20 mm in a cubic room of width 40 mm. Diffuse light is present in the room
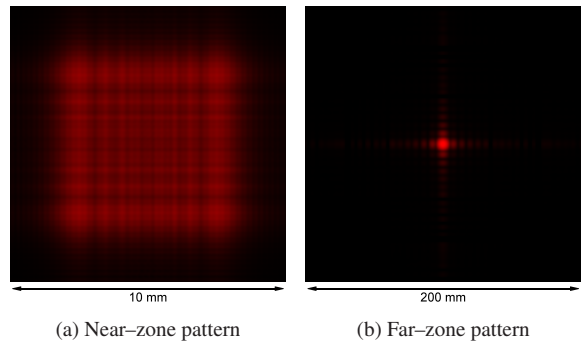

(a) Near–zone pattern   (b) Far–zone pattern

Figure 10: The ALF provides seamless transition between near and far field rendering. A square aperture whose size is 0.5 mm × 0.5 mm is probed by monochromatic light of 700 nm wavelength. Depending on the distance to the screen, significantly different patterns are observed. The near–zone pattern is captured when the distance is 10 mm and the far–zone pattern is taken at 3 meters.
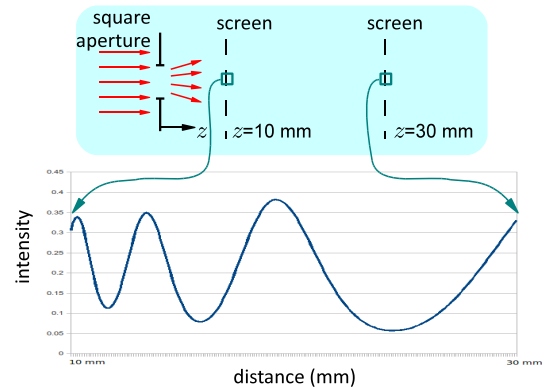

Figure 11: Validation of our rendering by the Fresnel diffraction formula using Fourier optics tools. ALF models the near–zone diffraction pattern successfully, where the intensity at the center pixel varies over 10 to 30 mm distance. The mean variance compared to Fourier optics simulation is less than 1%.

which lights up the box. Another source of light, at a distance of about 50 meters pours through a square aperture of size 1 mm × 1 mm. Figure 13 shows the interference pattern on a dynamically changing surface. The geometry is 200 mm×200 mm and the aperture is 1000 mm away from the surface.

5.4. Airy spot in a camera lens

In many computer vision and graphics applications, cameras are often modeled as pin–hole cameras and the point spread function (PSF) is a infinitesimally small point. However, in
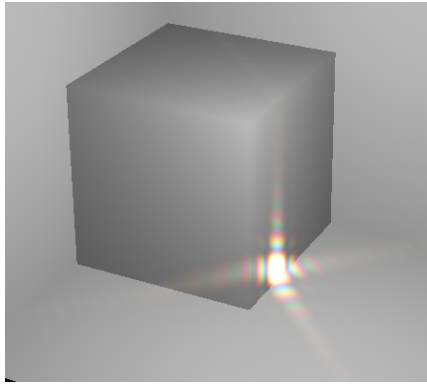
Figure 12: Rendering interference on arbitrarily oriented surfaces and changing light source direction (see video). Our implementation can mix and match traditional diffuse shading effects with wave effects.
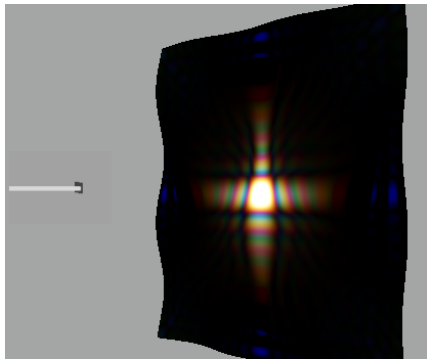

Figure 13: Rendering interference on a dynamic surface. The shader implementation can conveniently display diffraction and interference effects on dynamic geometry (see wavy mesh animation in supplementary video).

practice, cameras have finite apertures, thus even diffracted–limited lens produces finite sized–PSF. Here, we compute the PSF for two different camera lenses: 1) $f = 18$ mm and F/5.6 and 2) $f = 18$ mm and F/16. Although these two lenses have the same focal length, the size of the PSFs are different. To demonstrate the different size of PSF, we simulate images blurred by the PSF. The test image has a sharp edge between black and white regions, and the two lenses generate different images.

5.5. Fabric

Our final rendering example is a more realistic situation: diffraction from a fabric. Here, a fabric curtain is located between the scene and the camera, as shown in Fig. 1. We modeled the fabric as a binary amplitude grating and decomposed the fabric into multiple tiles. This is not only because computing the LF transformer of the entire fabric requires
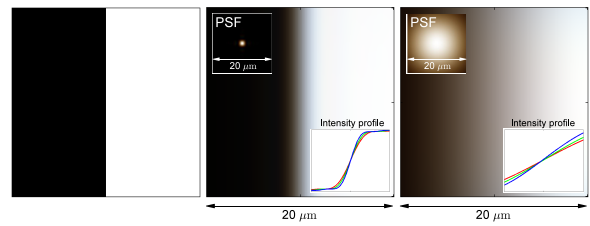
(a) original object   (b) F/5.6, f=18 mm   (c) F/32, f=18 mm

Figure 14: Reducing aperture shows diffraction and color smearing at a sharp black–white step edge. Ray–based models do not predict the decrease in image quality when a smaller aperture is used. The ALF correctly shows that with lower aperture sharp edges exhibit blur because the PSF of the lower F/# lens is narrower. A smaller aperture also leads to color dispersion, smaller wavelengths (blue) create smaller spots.

intensive computation but also it is easy to include irregularities of the real fabric patterns. In addition, the periodicity of real fabric is relatively larger than the pitch of optics–graded gratings so that the diffracted rays do not diverge over wide angles and the camera sees only small area of the fabric along the direct path from a point in the scene to the camera. We rendered a diffraction pattern from a single point and simulated an overall scene observed though the fabric, which is shown in Fig. 1.

All renderings were performed on nVidia 8800 GTS 320 MB card and shown in log–intensity domain. To achieve color effects, we render the scene at multiple wavelengths (typically 7) and use a simple color response linear weights for R, G, and B channels. The execution times were as follows. With 250,000 samples on the grating, the Young's slit experiment is shown at 512×512 resolution with 49 seconds/frame. The box scene is rendered at 1024×1024 resolution with 8 minutes/frame. The wavy mesh with 40,000 samples on grating and 512×512 resolution requires 2 minues/frame.

6. Conclusion

Ray–based representations can render complex scenes but fail to describe light transport associated with diffraction and interference. We have introduced the ALF framework that can handle diffraction in a purely ray–based representation. This challenges the earlier notion that diffraction is a purely wave-phenomenon and requires Fourier optics or optical path length based formulations. The only modification in the ALF is allowing virtual projectors that may have negative radiance. To locate virtual projectors for diffractive occluders, we introduced the light field transformers. To demonstrate the compatibility of the ALF with ray–based representation, we implemented an OpenGL shader and ren-

dered various diffraction phenomena. We validated the rendering results with respect to Fourier optics simulations.

Our current method is limited to single layer interface. Our gratings are separable patterns and the rendering method does not support arbitrary 2D patterns or non-planar gratings. In addition, our OpenGL shader based implementation limits total number of pairwise patches. In the future, rendering for multiple layers, global illumination effects and volumetric scattering will be useful. Future work also includes extending the ALF to polarized light [Alo04] to render light transport in birefringent materials. In computer graphics applications, we have shown lens aperture effects. Additional effects such as lens geometric and chromatic aberrations will allow more realistic insertion of synthetic elements in live action footage.

We speculate that the ALF framework would lead us to a variety of new applications. RayÐbased and beam-based representations have been used for sound rendering [TFNC01]. Audio wavelengths are subject to significant diffraction. Similarly fluids animation and haptics require solving wave propagation equations. Via source-based modeling, we hope that the ALF framework will simplify many such forward and inverse problems.

Acknowledgement

References

[Alo01] ALONSO M. A.: Radiometry and wide-angle wave fields. iii. partial coherence. *J. Opt. Soc. Am. A 18*, 10 (2001), 2502–2511. 2

[Alo04] ALONSO M. A.: Wigner functions for nonparaxial, arbitrarily polarized electromagnetic wave fields in free space. *J. Opt. Soc. Am. A 21*, 11 (2004), 2233–2243. 2, 10

[Bas09] BASTIAANS M. J.: Wigner distribution in optics. In *Phase-Space Optics: Fundamenals and Applications*, Testorf M., Hennelly B., Ojeda-Castañeda J., (Eds.). McGraw-Hill, New York, 2009, pp. 1–44. 2, 3, 5

[BW99] BORN M., WOLF E.: *Principles of optics*, 7th (expanded) ed. Cambridge University Press, Cambridge [England], New York, 1999. 2

[Cas08] CASTAÑEDA R.: Phase space representation of spatially partially coherent imaging. *Applied Optics 47*, 22 (August 2008), E53–E62. 5

[CCST00] CHAI J.-X., CHAN S.-C., SHUM H.-Y., TONG X.: Plenoptic sampling. In *SIGGRAPH* (2000), pp. 307–318. 3

[DHS*05] DURAND F., HOLZSCHUCH N., SOLER C., CHAN E., SILLION F. X.: A frequency analysis of light transport. *ACM Trans. Graph. 24*, 3 (2005), 1115–1126. 3

[FMCJS06] FURLAN W. D., MARTÍNEZ-CORRAL M., JAVIDI B., SAAVEDRA G.: Analysis of 3-D Integral Imaging Displays Using the Wigner Distribution. *Journal of Display Technology 2*, 2 (2006), 180–185. 2

[GGSC96] GORTLER S., GRZESZCZUK R., SZELISKI R., COHEN M.: The lumigraph. In *SIGGRAPH* (1996), pp. 43–54. 3

[Goo05] GOODMAN J. W.: *Introduction to Fourier optics*, 3rd ed. Roberts & Co., Englewood, Colo., 2005. 3

[Hec90] HECKBERT P.: Adaptive radiosity textures for bidirectional ray tracing. In *Siggraph* (1990), pp. 145–154. 7

[Hec02] HECHT E.: *Optics*, 4th ed. Addison-Wesley, Reading, Mass., 2002. 8

[IMG00] ISAKSEN A., MCMILLAN L., GORTLER S.: Dynamically reparameterized light fields. In *SIGGRAPH* (2000), pp. 297–306. 3

[LH96] LEVOY M., HANRAHAN P.: Light field rendering. In *SIGGRAPH 96* (1996), pp. 31–42. 3

[LW93] LITTLEJOHN R. G., WINSTON R.: Corrections to classical radiometry. *J. Opt. Soc. Am. A 10*, 9 (1993), 2024–2037. 2

[MMH*06] MAYCOCK J., MCELHINNEY C. P., HENNELLY B. M., NAUGHTON T. J., MCDONALD J. B., JAVIDI B.: Reconstruction of partially occluded objects encoded in three-dimensional scenes by using digital holograms. *Applied Optics 45*, 13 (2006), 2975–2985. 3

[Mor81] MORAVEC H. P.: 3d graphics and the wave theory. In *SIGGRAPH* (1981), vol. 15, pp. 289–296. 1, 3

[OBR08] OH S. B., BARBASTATHIS G., RASKAR R.: Augmenting light field to model wave optics effects. *tech report* (Nov 2008). 1

[PA07] PETRUCCELLI J. C., ALONSO M. A.: Propagation of partially coherent fields through planar dielectirc boudnaries using angle-impact wigner functions i. two dimensions. *J. Opt. Soc. Am. A 24*, 9 (2007), 2590–2603. 2

[ROAZ09] RASKAR R., OH S. B., ACCARDI A., ZHANG Z.: Light fields: Present and future. In *IEEE CVPR Tutorial* (Miami, FL, 2009), http://scripts.mit.edu/~raskar/lightfields/. 1

[SJ04] STERN A., JAVIDI B.: Sampling in the light of Wigner distribution. *J. Opt. Soc. Am. A 21*, 3 (2004), 360–366. 3

[Sta99] STAM J.: Diffraction shaders. In *SIGGRAPH* (1999), pp. 101–110. 3

[TFNC01] TSINGOS N., FUNKHOUSER T., NGAN A., CARLBOM I.: Modeling acoustics in virtual environments using the uniform theory of diffraction. In *SIGGRAPH* (2001). 10

[VRA*07] VEERARAGHAVAN A., RASKAR R., AGRAWAL A., MOHAN A., TUMBLIN J.: Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocusing. *ACM Trans. Graph. 26*, 3 (July 2007), 69:1–69:12. 3

[WAF99] WOLF K. B., ALONSO M. A., FORBES G. W.: Wigner functions for helmholtz wave fields. *J. Opt. Soc. Am. A 16*, 10 (1999), 2476–2487. 2

[Wal73] WALTHER A.: Radiometry and Coherence. *J. Opt. Soc. Am. 63*, 12 (1973), 1622–1623. 2

[WMZ98] WOLF K. B., MENDLOVIC D., ZALEVSKY Z.: Generalized Wigner function for the analysis of superresolution systems. *Applied Optics 37*, 20 (1998), 4374–4379. 3

[ZCG08] ZIEGLER R., CROCI S., GROSS M.: Lighting and occlusion in a wave-based framework. *Computer Graphics Forum 27*, 2 (2008), 211–228. 3

[ZL09] ZHANG Z., LEVOY M.: Wigner distributins and how they relate to the light field. In *IEEE International Conference on Computational Photography* (2009). 2, 3