

A Framework for Holographic Scene Representation and Image Synthesis

Remo Ziegler, *Member, IEEE*, Peter Kaufmann, and Markus Gross, *Senior Member, IEEE*

Abstract—We present a framework for the holographic representation and display of graphics objects. As opposed to traditional graphics representations, our approach reconstructs the light wave reflected or emitted by the original object directly from the underlying digital hologram. Our novel holographic graphics pipeline consists of several stages including the digital recording of a full-parallax hologram, the reconstruction and propagation of its wavefront, and rendering of the final image onto conventional, framebuffer-based displays. The required view-dependent depth image is computed from the phase information inherently represented in the complex-valued wavefront. Our model also comprises a correct physical modeling of the camera taking into account optical elements, such as lens and aperture. It thus allows for a variety of effects including depth of field, diffraction, interference, and features built-in anti-aliasing. A central feature of our framework is its seamless integration into conventional rendering and display technology which enables us to elegantly combine traditional 3D object or scene representations with holograms. The presented work includes the theoretical foundations and allows for high quality rendering of objects consisting of large numbers of elementary waves while keeping the hologram at a reasonable size.

Index Terms—Holography, light propagation, wave, diffraction, aliasing, image synthesis, graphics representation.

1 INTRODUCTION

TRADITIONALLY, graphics objects or scenes have been represented through geometry and appearance. Most often, the former is described by geometric primitives including triangles, points, basis functions, and others, while the latter is encoded in texture maps. Due to the difficulty of representing some real-world objects, such as fur, hair, or trees, with traditional techniques, research has also focused on image-based rendering using light-fields [1], [2], [3], lumigraphs [4], [5], reflectance fields [6], [7], [8], sprites, and other models. All of these graphics representations share the property that their treatment of light is motivated through ray optics and that rendering involves projection and rasterization, or ray tracing.

In recent years, significant progress has been made in the field of display technology with the goal to replace the 2D flat screen by more advanced and immersive three-dimensional displays. Besides various autostereoscopic displays [9], [10], there has been an increased interest in holographic displays [11] which reconstruct the 3D object or scene wavefront directly from the underlying holographic representation. In spite of some current limitations, such as computational costs and sampling rates, the rapid development of compute power and data storage makes holographic representation, image generation, and display a technology of greatest potential for the future of computer graphics and interactive visual simulation.

In this paper, we present a novel framework for graphics representation, processing and display which is entirely based on digital holography and wave optics, while leaving the possibility of combining it with traditional rasterized rendering. Holograms are elegant structures capturing the phase and amplitude of an object or scene wavefront as seen from all possible views through a window of given aperture. The major difference to a lightfield, however, is its intrinsic wave optics handling anti-aliasing implicitly and its ability to reproduce object depth through phase. Thus, holograms overcome some of the inherent limitations of image-based methods including defocus and compositing with conventional graphics scenes.

A central feature of this framework is its seamless integration into conventional, framebuffer-oriented 2D display technology as well as its compliance with future 3D holographic displays. The core component of our novel representation is a digital hologram which can be recorded both from real-world objects and from conventional, triangle or point-based graphics objects. In order to reconstruct 2D images at a given camera position, the original wavefront has to be reconstructed from the hologram and propagated through space. To this end, we utilize fast, discrete approximation methods based on Fourier theory and angular spectrum, which have been described in a continuous form in [12] and in a refined way in [13]. In this paper, we specifically describe required resampling steps for this discrete angular spectrum propagation. In addition, we will show that this wave-based approach allows for an elegant modeling of the camera optics. The introduction of two different camera models, namely, a pinhole camera as well as a physically based camera model including a thin lens allow optical effects, such as depth of field, varying focal length, chromatic aberration or diffraction, which would not easily

- R. Ziegler is with ETH Zentrum, IFW D27.1, Haldeneggsteig 4, 8092 Zurich, Switzerland. E-mail: zieglerr@inf.ethz.ch.
- P. Kaufmann is with Kindergartenweg 4, CH-5243 Muelligen, Switzerland. E-mail: kaufmann@pop.agri.ch.
- M. Gross is with ETH Zentrum, IFW D28.1, Haldeneggsteig 4, 8092 Zurich, Switzerland. E-mail: zieglerr@inf.ethz.ch.

Manuscript received 12 May 2006; revised 28 Sept. 2006; accepted 4 Oct. 2006; published online 10 Jan. 2007.

For information on obtaining reprints of this article, please send e-mail to: tcvg@computer.org, and reference IEEECS Log Number TVCG-0067-0506.

be possible using a theoretical camera model as it is used in traditional rasterization systems. A crucial step in the rendering of digital holograms on a conventional 2D screen is the reconstruction of the depth buffer. We will present a novel algorithm to reliably reconstruct object depth from the digital hologram and to render it along with conventional graphics representations.

Despite some earlier work on holographic representations in graphics [14], [15], [16] and computer generated holography (CGH) [17], [11], this paper, for the first time, presents a full framework for a holography-inspired graphics pipeline, which allows to generate, process and render holograms from synthetic and real 3D objects. Rendering holograms from synthesized objects allows using small holograms, since the wavelength can be chosen outside of the visible spectrum. Our main contributions stated above are both fundamental and practical. They include a summary of the relevant theory as well as a variety of algorithms and methods for the digital recording, wave propagation and 2D rendering of holograms, while elaborating the inherently aliasing free image generation by wave-based rendering. Some rendering and compositing examples demonstrate the versatility and quality of the presented methods.

2 RELATED WORK

A lot of work has been done in holography since Gabor invented “wavefront reconstruction” in 1948 [18], which today is known as holography. More recently, the progress in computer technology provided the possibility to propagate light waves in space efficiently and therefore to record as well as to reconstruct a hologram digitally. Goodman [19] gives a good introduction to fast paraxial light wave propagation of planar wavefronts based on the Fourier spectrum. Restrictions concerning tilted planes are lifted in [20], [21], [22] while still applying paraxial propagation of waves. Delen and Hooker [12] extend Tommasi’s system such that propagation to parallel planes becomes feasible while having the precision of Rayleigh-Sommerfeld. Fourier analysis has also successfully been applied to light transport based on ray optics by Durand et al. [23] where the influence of shading, occlusion, and transport on the frequency content of radiance is studied.

Fourier spectrum propagation has recently been applied by Matsushima [13], [24] for digital hologram recording of textured objects. Similarly to the method presented in [25], the 3D Fourier spectrum has been used for fast 2D slice projection for computer generated holography in [26]. Further efficient propagation of objects consisting of points, lines, or triangles were presented in [27], [28], [29] showing a clear quality loss in the hardware-based approaches. Lucente and Galyean [14], [15] and Halle and Kropp [16] take advantage of graphics hardware to render higher quality holographic stereograms to a holographic screen achieving interactive rates while loosening the constraint of full-parallax. An excellent overview of holographic screens is given in [11].

Digital reconstructions of captured holograms allow mathematical filtering [30], [31] improving the final image quality. Filtered holograms can adopt negative or complex

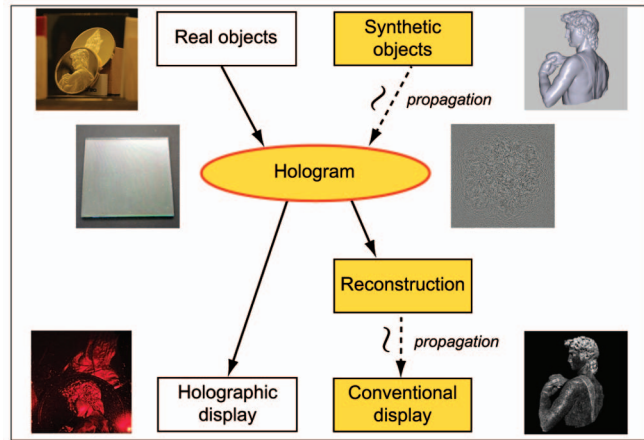


Fig. 1. Overview of the wave framework with the components of our new wave pipeline highlighted in orange. (The David data set is courtesy of Stanford University, Digital Michelangelo Project.)

values and, therefore, not be usable for holographic screens anymore. Still, the digital reconstruction of the filtered hologram can be achieved as shown in [32]. However, Schnars and Jüptner did not simulate a camera model with a lens showing limited depth of field. Recent interest in realistic camera models with limited depth of field has been shown in computer graphics in [3], [33], [34]. Ng [35] even built a prototype of a hand-held plenoptic camera capturing a 4D light field of which a refocused 2D image can be calculated.

3 OVERVIEW

In this section, a conceptual overview of the complete wave framework, including our novel holographic pipeline (highlighted in orange in Fig. 1) will be given.

The first step of this framework involves recording of holograms. We distinguish between recording and acquisition of holograms. Acquisition stands for the physical capturing of a hologram of *real objects* onto a high resolution photographic film or a CCD as in digital holography, whereas recording (see Section 5.2) will be used for the computation of a hologram from *synthesized objects* as in computer generated holography (CGH). In order to generate a hologram from a wavefront given on the surface of a *synthetic object*, we have to be able to propagate this wavefront to any position and time in space. We explain different techniques of propagation in Section 4. Furthermore, the hologram can be used as input to a holographic screen as well as an input for further processing (see Section 6) leading to an output on a conventional display. This output requires a discrete reconstruction of the wavefront at the hologram followed by a propagation of this wavefront to an arbitrarily placed camera as described in Section 7, which provides the input to a conventional display.

We will give a thorough introduction into theory of scalar waves in Section 4 since it builds the fundamental description of how waves are traveling through space in our pipeline and can therefore be used for recording and reconstruction of holograms as well as image generation. In

the remainder of the paper, we concentrate on the theoretical aspects and problems of the wave pipeline and refer to [19] for issues only being relevant to practical holography.

4 PROPAGATION

Light propagation is affected by the phenomena of refraction, reflection, diffraction, and interference. In our current model, we do not account for refraction or reflection, but focus on the influence of diffraction and interference. We describe the evaluation of spherical waves and plane waves at arbitrary points in space in Section 4.1 and elaborate the efficient propagation of a wavefront given on a surface S to an arbitrary point P' or to an arbitrarily placed surface S_A in the remainder of the section. Since most of our synthetic objects consist of planar primitives like triangles and rectangles, we will focus on their propagation efficiency and refer to [27] for a fast evaluation of point sampled objects based on recurrence formula. Issues arising when evaluating the continuous equations in discrete space are discussed in Section 4.4.

4.1 Scalar Wave Representation

Throughout the paper, the wave representation is based on scalar diffraction theory and is therefore an approximation of the complete Maxwell equations. A general time harmonic scalar wave $u(P, t)$ at position P and time t can be represented as

$$u(P, t) = \Re\{U(P)e^{-i\omega t}\}, \quad (1)$$

with $\omega = 2\pi\nu$ being the angular frequency while ν describes the frequency or oscillation of the wave. The complex function $U(P)$ with the real-valued amplitude $A(P)$ and phase $\varphi(P)$ at position P is defined as

$$U(P) = A(P)e^{i\varphi(P)}. \quad (2)$$

The time-dependent term $e^{-i\omega t}$ of (1) can be neglected since the simulated light of our pipeline is of monochromatic nature, as is most often the case when working with holograms and, therefore, $U(P)$ is sufficient to adequately describe $u(P, t)$ and is used as the wave representation in our paper.

A spherical wave originating at point P_0 is being represented as

$$U(P) = A_0 \frac{e^{ikr + \varphi_0}}{r}, \quad (3)$$

where A_0 is the real-valued amplitude at P_0 , k is the wave number, φ_0 the initial phase at P_0 , P is the position of evaluation and $r = \|P - P_0\|_2$. The wave number k is defined as $k = \frac{2\pi}{\lambda}$ with λ being the wavelength. A monochromatic plane wave can be represented as

$$U(P) = A_0 e^{i\mathbf{k} \cdot \mathbf{r} + \varphi_0}, \quad (4)$$

where A_0 and φ_0 are the real-valued amplitude and phase at the origin P_0 and \mathbf{r} is defined as $\mathbf{r} = P - P_0$. The vector \mathbf{k} is defined as $\mathbf{k} = k * (\alpha, \beta, \gamma)$ with k being the wave number and (α, β, γ) being the unit vector pointing in the direction




Fig. 2. Propagation of a wavefront based on Huygen's, Kirchhoff's, and Rayleigh-Sommerfeld's formulations.

of the wave propagation. The components of the vector are called *directional cosines*.

4.2 Wavefront Propagation

In 1678, Christian Huygens [19] presented a way of describing the propagation of a wavefront as the "envelope" of "secondary" spherical waves being emitted by every point of the primary wavefront at time $t = 0$ (see Fig. 2). Later, Kirchhoff put the theory on firmer mathematical ground by finding a formulation to evaluate the wave at an arbitrary point P' in space knowing the evaluation and the derivative of the homogeneous wave equation on a closed surface S , enclosing volume V containing P' (see Fig. 2):

$$U(P') = \frac{1}{4\pi} \iint_S \frac{e^{ikr}}{r} \frac{\partial U(P)}{\partial \mathbf{n}} - U(P) \frac{\partial}{\partial \mathbf{n}} \left(\frac{e^{ikr}}{r} \right) ds, \quad (5)$$

$$U(P') = \begin{cases} U(P') & \text{if } P' \in V \\ 0 & \text{if } P' \notin V. \end{cases}$$

\mathbf{n} is the surface normal, $U(P')$ is the evaluation at point P' , whereas $U(P)$ is the evaluation of the wave at $P \in S$. The so-called *Rayleigh-Sommerfeld* theory can be derived from the Kirchhoff formula with the difference of only imposing boundary conditions on the evaluation of the wave function or on its derivative. The limitation of only being valid for planar geometries is not a real restriction, since the hologram as well as different primitives of synthetic objects like triangles and rectangles are planar as well. Assuming point P' away from the immediate vicinity ($k \gg 1/r$) of the aperture S_A , a propagation in the positive z direction perpendicular to S_A with S_A at $z = 0$ and $P \in S_A$ leads to the Rayleigh-Sommerfeld formula

$$U(P') = \frac{1}{2\pi} \iint_{S_A} U(P) \frac{\partial}{\partial \mathbf{n}} \left(\frac{e^{ikr}}{r} \right) ds, \quad (6)$$

$$U(P') = \frac{i}{\lambda} \iint_{S_A} U(P) \frac{e^{ikr}}{r} (\mathbf{r} \cdot \mathbf{n}) ds, \text{ with } \frac{k}{2\pi} = \frac{1}{\lambda}. \quad (7)$$

The Rayleigh-Sommerfeld equation (7) can be interpreted from a physical point of view as a superposition of spherical waves $\frac{e^{ikr}}{r}$ located at the aperture S_A with amplitude $\frac{U(P)}{\lambda}$ multiplied by a phase shift of 90° caused by the multiplication of i . Additionally the spherical waves are multiplied by a directional factor $(\mathbf{r} \cdot \mathbf{n})$. For further details and derivations, please see [19] and included references.




Fig. 3. The angular spectrum splits a wave field in multiple plane waves and can be used to propagate a wave from an aperture S_A to an aperture $S_{A'}$. (a) For an arbitrary placement of $S_{A'}$ “straight” propagation, (b) rotation of the spectrum, and (c) propagation on parallel offset planes is needed. The figure shows the 2D case, where the third dimension can be imagined to be perpendicular to x and z .

Evaluating these general forms of scalar diffraction theory can be very time consuming. There exist two approximations one for the *near field* called the *Fresnel approximation* and one for the *far field* called the *Fraunhofer approximation*. As shown in [19], the Fresnel approximation can be efficiently calculated by a convolution, while the Fraunhofer approximation can be interpreted as a Fourier transform.

However, direct integration as well as the approximations have the deficiency of either being inefficient or inaccurate. Therefore, we apply a method first presented by Delen and Hooker [12] and refined by Matsushima [13] where they present an accurate and fast propagation based on angular spectrum propagation presented in [36], [19].

4.3 Angular Spectrum

The complex monochromatic planar field $U(x, y, 0)$ given at an aperture S_A can be split into multiple uniform infinite plane waves traveling in different directions (see Fig. 3a) by applying a Fourier transform on $U(x, y, 0)$ leading to the angular spectrum $\mathcal{A}(\nu_x, \nu_y, 0)$. Therefore, the angular spectrum $\mathcal{A}(\nu_x, \nu_y, 0)$ is an assembly of all plane waves defined as $e^{-2\pi i(\nu_x x + \nu_y y)}$ in (8). The frequencies ν_x and ν_y can be substituted by the directional cosines α and β of the plane wave (4) as $\nu_x = \frac{\alpha}{\lambda}$ and $\nu_y = \frac{\beta}{\lambda}$ leading to (9).

$$\mathcal{A}(\nu_x, \nu_y, 0) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} U(x, y, 0) e^{-2\pi i(\nu_x x + \nu_y y)} dx dy \quad (8)$$

$$= \mathcal{F}\{U(x, y, 0)\},$$

$$\mathcal{A}\left(\frac{\alpha}{\lambda}, \frac{\beta}{\lambda}, 0\right) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} U(x, y, 0) e^{-2\pi i\left(\frac{\alpha}{\lambda}x + \frac{\beta}{\lambda}y\right)} dx dy. \quad (9)$$

To propagate the wave field along z to a parallel plane at distance z , the angular spectrum $\mathcal{A}(\nu_x, \nu_y, 0)$ has to be multiplied by a propagation phase term:

$$\mathcal{A}\left(\frac{\alpha}{\lambda}, \frac{\beta}{\lambda}, z\right) = \mathcal{A}\left(\frac{\alpha}{\lambda}, \frac{\beta}{\lambda}, 0\right) e^{\frac{2\pi i}{\lambda} z \sqrt{1 - \alpha^2 - \beta^2}} \quad \text{for } \alpha^2 + \beta^2 < 1. \quad (10)$$

In case of $\alpha^2 + \beta^2 > 1$, the wave would evaluate to a real-valued so-called evanescent wave and, therefore, not carry any energy away from the plane at $z = 0$. The propagation defined in [19] is restricted to propagation between parallel planes sharing a common center.

4.3.1 Angular Spectrum Rotation

Using the rotation of the angular spectrum from plane S_A with coordinate system (x, y, z) to $S_{A'}$ with coordinate system (ξ, η, ζ) (see Fig. 3b) as proposed by [20] lifts the restriction of propagation onto parallel planes. The rotational matrix M transforming (ξ, η, ζ) into (x, y, z) also relates the corresponding spatial frequencies $(\nu_\xi, \nu_\eta, \nu_\zeta)$ to (ν_x, ν_y, ν_z) by

$$(\nu_\xi, \nu_\eta, \nu_\zeta)^T = M^T (\nu_x, \nu_y, \nu_z)^T. \quad (11)$$

The coordinate transformation causes a shift of the sampling points requiring interpolation of the values for the new coordinate system. Additionally, backward propagating waves are set to 0. According to [12], linear interpolation is appropriate to find amplitudes at intermediate points.

Propagation to Parallel Plane with Offset: If the parallel plane $S_{A'}$ is offset to the propagation direction of S_A by $(x_0, y_0, 0)$ as depicted in Fig. 3c, the propagation can be calculated by means of the Fourier shift theorem [12] as

$$U(x, y, z) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \hat{\mathcal{A}}(\nu_x, \nu_y, 0) e^{i2\pi z \sqrt{\frac{1}{\lambda^2} - \nu_x^2 - \nu_y^2}} \times e^{i2\pi(\nu_x x + \nu_y y)} d\nu_x d\nu_y, \quad (12)$$

where

$$\hat{\mathcal{A}}(\nu_x, \nu_y, 0) = \mathcal{A}(\nu_x, \nu_y, 0) e^{i2\pi(\nu_x x_0 + \nu_y y_0)} \quad (13)$$

expresses the Fourier shift theorem. Using the propagation of the angular spectrum in continuous space, yields to the identical result as the Rayleigh-Sommerfeld (RS) formula (6). A practical implementation of RS using convolution affords a discretization leading to different problems which are treated in the following sections.

4.4 Discrete Propagation

The Fourier transform (8) and the continuous propagation (12) are defined as an integral over an infinite plane, whereas in the discrete case the Discrete Fourier Transform (DFT) and the summation is restricted to a finite size leading to artifacts. Furthermore, the energy conservation does not hold, since the wave is not being evaluated over an infinite plane after propagation and, therefore, decreases with an increasing distance z [22]. By zero-padding the plane before propagation those artifacts can be limited.

DPPO: For the sake of simplicity, the derivation for discrete propagation to parallel planes with offset (DPPO) will be based on a propagation between two squares, but the extension to rectangles is straightforward.

The aperture S_A is sampled using a uniform grid of $N \times N$. The discretized wave field on S_A is given as U_{xy} and Fourier transformed using DFT leading to the angular




Fig. 4. The top row depicts the setting of S_A and $S_{A'}$ before modification. The bottom row shows the solution for the matched size after zero-padding in the spatial domain or cutting in the spatial domain and the adjusted sampling by zero-padding in the frequency domain and cutting in the frequency domain.

spectrum \mathcal{A}_{ml} . The propagated wave field U'_{xy} can be evaluated at distance z based on (12) in the discrete case as

$$U'_{xy} = \frac{1}{N^2} \sum_{m=0}^{N-1} \sum_{l=0}^{N-1} \hat{\mathcal{A}}_{ml} e^{i2\pi z \sqrt{\frac{1}{\lambda^2} - \left(\frac{m}{N}\right)^2 - \left(\frac{l}{N}\right)^2}} e^{i2\pi\left(\frac{m}{N}x + \frac{l}{N}y\right)} \quad (14)$$

for $x = 0..N-1$ and $y = 0..N-1$ with

$$\hat{\mathcal{A}}_{ml} = \mathcal{A}_{ml} e^{i2\pi\left(\frac{m}{N}x_0 + \frac{l}{N}y_0\right)}. \quad (15)$$

The asymptotic complexity can be reduced from $O(N^4)$ for the direct integration of RS to $O(N^2 \log N)$ for DPPO whereas the complexity of the rotation described in Section 4.3.1 is negligible. The complexity of DPPO differs solely by a constant factor from the Fresnel and Fraunhofer approximation while still evaluating the propagation with the precision of RS.

Discrete Propagation Pipeline: So far, we summarized one method for rotation or tilting in Section 4.3.1 and one for propagation in Section 4.4. These two methods can now be combined to evaluate a wave originating on S_A on any arbitrarily placed aperture $S_{A'}$. The accuracy of the result relies on the sampling of S_A and $S_{A'}$ and on the order of the tilt and the propagation step. In case propagation and tilt would be needed, the tilt operation has to be evaluated first. For minimal sampling requirements, refer to Section 6.1.

Since the source aperture S_A does not necessarily have the same resolution and size as $S_{A'}$, resampling or resizing might be necessary. The four possible cases and their solution are depicted in Fig. 4. The sizes of the apertures S_A and $S_{A'}$ are preferably chosen as a power of 2 in order to guarantee an efficient FFT.

5 HOLOGRAM

5.1 Basics

A full-parallax hologram can be imagined as a window revealing limited viewing possibilities from a 3D scene (see Fig. 1). There exist many different holograms, being able to capture different properties of complex waves, such as phase, amplitude, or intensity. Different setups and different shapes increase the variety even further [19]. We are simulating a transmission hologram as in Fig. 5, a setup where the laser and the object are placed on the same side of




Fig. 5. Possible setup of a real transmission hologram using a beam-splitter to warrant monochromatic waves.

the finite planar hologram during recording. The reconstruction of this type of hologram requires a laser placed at the same relative position to the hologram and having the same wavelength as for recording. For the remainder of the paper, we will refer to a transmission hologram simply as *hologram*. The hologram is a thin real-valued and densely sampled plane measuring the static wavefield intensity created by interference of multiple monochromatic waves. In a physical setup of a hologram, the laser light source is used to illuminate the object as well as the hologram in order to guarantee monochromatic nature of the waves. The laser can be split up in these two parts using a beam-splitter as depicted in Fig. 5. The light from the object will be called the *object wave*, whereas the light from the laser during recording will be called *reference wave* and during reconstruction *reconstruction wave*. The following two sections will conceptually describe the recording of a hologram in Section 5.2 as well as the reconstruction of the object wave in Section 6. Capturing the interference pattern affords a very high sampling rate of the hologram. Sampling rate requirements will be discussed in Section 6.1.

5.2 Recording

For simplicity, we will base the evaluation of a wave in space in this section and Section 6 on equations described in Section 4.1. The same concept of recording and reconstruction holds when evaluating waves as described in Section 4. Furthermore, having a purely theoretical approach allows assumptions of monochromatic laser properties implying spatial and temporal coherence and no motion artifacts. Two incident light waves and a holographic plate are needed to generate a hologram. One of the light waves being of type (4) stems from the laser source and will be called reference wave R . Another wave stems from the object and will be referred to as object wave O . In a theoretical recording, we do not have to simulate the reflection of the laser light at the object as depicted in Fig. 5, but sample our object into many point sources or planar segments emitting waves of type (2) with the same wavelength λ as the laser light R . Therefore, O is the superposition of all waves emitted by the subsegments of the object. Fig. 6a shows the applied setup for hologram recording. The two light waves O and R are mutually coherent, leading to a time independent interference pattern in space. The intensity of this interference pattern is being recorded by a holographic plate as function H :

$$H = |O + R|^2. \quad (16)$$

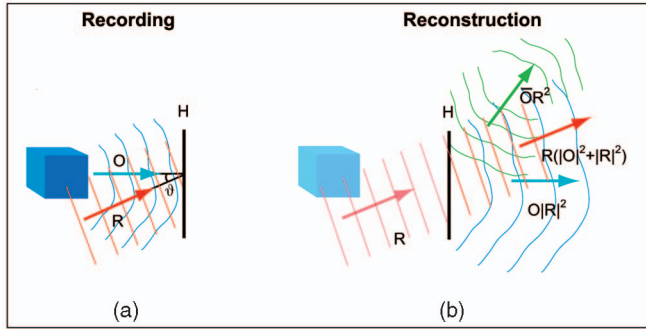


Fig. 6. (a) Hologram recording by interference of object wave O and reference wave R . (b) Hologram reconstruction using the reconstruction wave R resulting in the reference term $R(|O|^2 + |R|^2)$, the real image $\bar{O}R^2$ and the virtual image $O|R|^2$.

Methods as presented in Section 4 can be used to evaluate the object wave at the hologram in an efficient manner.

5.2.1 Recording of Specific Objects

This pipeline supports two different primitives. On the one hand, points are represented by point sources implemented as spherical waves with the center at the position of the point. The waves can efficiently be evaluated using shaders on graphics hardware. Traditional point-based objects can, therefore, be represented by a number of point sources. Backfacing or occluded points are rejected during a splatting step. The remaining points are represented as point sources in our pipeline. Textured triangular meshes can also be sampled into point sources and, therefore, be evaluated similar to point-based objects. In both cases, a view dependent density sampling or a view dependent scaling of the point sources is needed in order to avoid areas of inconsistent intensity. Our current pipeline handles this effect only for point-based objects.

On the other hand, objects like the cube or planes are represented by aperture primitives. The textured apertures can be directly transformed into a rectangular discrete wavefield by using the texture value to evaluate the amplitude and by assigning a random phase distribution to propagate the light in all directions. This rectangular discrete wavefield can then be propagated using the method described in Section 4.4.

6 RECONSTRUCTION

The holographic material is assumed to provide a linear mapping between the recorded intensity and the amplitude transmitted or reflected during the reconstruction process. A reconstruction wave R' being the same as the former reference wave R illuminates the holographic plate and is modulated by the transmittance T

$$T = t_b + \beta' H = H, \text{ with } t_b = 0 \text{ and } \beta' = 1. \quad (17)$$

Thereby, t_b is a "bias" transmittance which can be set to 0 and β' is a scaling factor which will be set to 1 in our theoretical setup leading to $T = H$. The incident reconstruction wave R' is therefore modulated by H leading to a reconstructed wave U as:

$$\begin{aligned} U &= R' \cdot H = R'(|O|^2 + |R|^2) + \bar{O}R'R' + \bar{O}R'R' \\ &= \underbrace{R(|O|^2 + |R|^2)}_{\text{reference term}} + \underbrace{O|R|^2}_{\text{virtual image}} + \underbrace{\bar{O}R^2}_{\text{real image}}. \end{aligned} \quad (18)$$

Reference Term. The reference term $R(|O|^2 + |R|^2)$, can be approximated as a scaled version of the reference wave. This assumption is true if $|O|^2$ and $|R|^2$ are constant over the whole hologram. For R being a planar wave, this is trivial. However, for O the assumption is only true if the object is not close to the hologram or if $|O|^2$ is small compared to $|R|^2$ and, therefore, variations would be negligible. The influence of O is further discussed in Section 6.2.

Virtual Image. The term $O|R|^2$ can be considered proportional to O , since $|R|^2$ is constant over the whole hologram. This is also called the virtual image and is the part we are interested in, since it describes the same wave as has been emitted by the object during recording.

Real Image. The last term $\bar{O}R^2$ is the conjugate of the object wave and produces the real image which corresponds to an actual focusing of light in space. The conjugate wave \bar{O} is propagating in the opposite direction of O and, therefore, focuses in an image in front of the hologram. If R is perpendicularly incident to the hologram plane, R^2 is a constant value and \bar{O} will be mirrored at the hologram.

6.1 Sampling

Recording the interference pattern affords a *very high* sampling rate. It depends on the wavelength λ and the angle ϑ between the direction of propagation of O and the propagation direction of R . The sampling distance d can be computed by

$$d = \frac{\lambda}{2\sin(\vartheta)}. \quad (19)$$

The chosen sampling distance during recording also has an influence on the maximum viewing angle of the hologram which has to be considered when positioning a camera in Section 7. To achieve a maximum viewing angle $\vartheta = \frac{\pi}{2}$ the sampling distance d has to be at least as small as $\frac{\lambda}{2}$ in order to guarantee that the Nyquist frequency of the hologram is equal to the carrier frequency of the interference pattern. Note that, by simulating synthetic objects, only three orders of magnitude bigger than the wavelength we are able to keep the hologram size reasonable and still get nice images, since the resolution of the final image and not the physical size of it is important when rendering to a screen (see Table 1).

6.2 Filtering

The reconstruction leads to three different terms (see Fig. 6b), whereas we are only interested in the virtual image $O|R|^2$. Therefore, the hologram is filtered in a way to suppress the unwanted terms during reconstruction while maintaining the quality of the virtual image. To suppress unwanted terms, we have to rewrite (16) as

$$H = |O|^2 + |R|^2 + \bar{O}R + \bar{O}R. \quad (20)$$

Considering (2), O and R can be substituted such that $O = A_O e^{i\varphi_O}$ and $R = A_R e^{i\varphi_R}$, where A_O is the amplitude of O , φ_O

TABLE 1
Statistics

Figure	# Of Prim.	Holo. Size	Pict. Size	Time per It.	# Of It.
Fig. 13	100K	1024 ²	1024 ²	185s	70
Fig. 14a)-d)	50K	512 ²	512 ²	43s	24
Fig. 14e)	50K	1024 ²	1024 ²	155s	74
Fig. 14f)	300K	1024 ²	1024 ²	325s	70
Fig. 15a)/d)	137K	1024 ²	1024 ²	197s	90
Fig. 15b)	173K	1024 ²	1024 ²	241s	90

is the phase of O , A_R is the amplitude of R , and φ_R is the phase of R leading to

$$H = A_O^2 + A_R^2 + \underbrace{A_O A_R e^{i\varphi_O} e^{-i\varphi_R} + A_O A_R e^{i\varphi_O} e^{-i\varphi_R}}_{2A_O A_R \cos(\varphi_O - \varphi_R)} \quad (21)$$

$$H = A_O^2 + A_R^2 + 2A_O A_R \cos(\varphi_O - \varphi_R).$$

For better readability the x and y dependency of H , A_O , A_R , φ_O , and φ_R have been omitted. The first two terms of (21) are close to be constant over the whole hologram, while the third term is statistically varying from $2A_O A_R$ to $-2A_O A_R$. Thus, the average intensity \hat{H} over the whole hologram H can be approximated by $\hat{H} \approx A_O^2 + A_R^2$. The first two terms can therefore be suppressed by subtracting the average intensity \hat{H} from the hologram H leading to $H' = H - \hat{H}$. If H' is used for reconstruction, the reference term will be suppressed. This method is known as DC-term suppression in [31], since the almost constant term $A_O^2 + A_R^2$ leads to the DC-term $\mathcal{H}(0, 0)$, where \mathcal{H} is the Fourier transform of the hologram H . Setting the reference wave sufficiently off-axis will deviate \bar{O} such that it cannot be seen in front of the hologram.

A different method consisting of a high-pass filtering or asymmetrical masking of the hologram as presented in [30] would only improve our results if the spectra of R , O , and \bar{O} would not overlap.

7 IMAGE GENERATION

Based on the propagation described in Section 4, any wave field given on a synthetic object or on a reconstructed hologram can be evaluated on an aperture of an arbitrarily placed camera in space. The image generation consists of




Fig. 7. Each row of pictures corresponds to a certain angle ϑ . The DC-Term of the hologram H gets suppressed leading to hologram H' depicted in the third column. In the reconstruction based on H , we can clearly see the reference term, while a reconstruction based on H' leads to much better results. Best results are obtained with a big angle ϑ leading, however, to a high sampling rate of the hologram.




Fig. 8. Image generation simulating a pinhole camera as in (a) can be achieved by DPPO (b) or by making use of the directional cosines (c). A perspective distortion leads to the final image (d).

four steps, namely, the simulation of an optical camera model Section 7.1, a multiwavelength rendering for colored holograms Section 7.2, a depth evaluation for compositing of multiple objects Section 7.3, and a speckle noise reduction Section 7.5.

7.1 Camera Model

In a first step, the simplest architecture of a camera, the pinhole camera, is studied for image generation. Artifacts occurring because of limited aperture size are described in Section 7.1 while a more complex camera model including a lens is introduced in Section 7.1.

Pinhole Camera: As shown in Fig. 8a, a wave is propagating through a pinhole and producing an image on a parallel plane at distance z . Modeling the slit by an aperture S_A and applying DPPO to propagate the wave to the image plane $S_{A'}$ would be possible, but would require zero padding of S_A (see Fig. 8b). Actually, assuming an indefinitely small slit and knowing the directional components (α, β, γ) of the wave front (see Section 4.3) is sufficient to generate the image at a distance z assuming geometrical propagation from the slit (see Fig. 8c). The directional components can be calculated from the frequency components ν_x and ν_y of the angular spectrum as

$$(\alpha, \beta, \gamma) = (\lambda\nu_x, \lambda\nu_y, \sqrt{1 - (\lambda\nu_x)^2 - (\lambda\nu_y)^2}).$$

Finally, a mapping between the directional cosines and the image pixels given by coordinates u and v as in Fig. 8d) can be found by a projection matrix P as $(\alpha, \beta, \gamma, 1)P \Rightarrow (u, v)$ leading to

$$\left. \begin{aligned} \alpha &= -\frac{u}{s_x} \sqrt{\left(\frac{u}{s_x}\right)^2 + \left(\frac{v}{s_y}\right)^2 + 1}^{-1} \\ \beta &= -\frac{v}{s_y} \sqrt{\left(\frac{u}{s_x}\right)^2 + \left(\frac{v}{s_y}\right)^2 + 1}^{-1} \end{aligned} \right\} s_y = \frac{1}{\tan(\frac{\vartheta}{2})}, s_x = r_a s_y. \quad (22)$$

θ is the field of view in y direction with the aspect ratio $r_a = \frac{w}{h}$, where w is the image width and h the image height. The pixels of the image are computed by a weighted sum (e.g., bilinear interpolation) of the directional cosines.

Resolution. The maximal resolution of the image is determined by the size of the pinhole W . This can be shown by the relation of the spacing δ_f in the frequency and the spacing δ in the spacial domain given as $\delta\delta_f = \frac{1}{N}$, where N is the number of samples. The smallest nonzero frequency is given as $\delta_f = \frac{1}{W}$, since $W = N\delta$. This corresponds to the direction cosine $\alpha_{min} = \frac{1}{W}$ and gets close to the Rayleigh resolution criterion $\alpha = 1.22 \frac{\lambda}{D}$ describing the critical angle

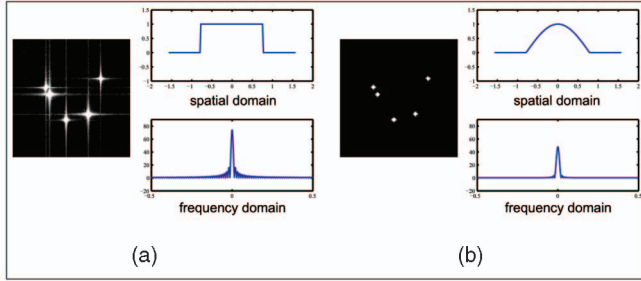


Fig. 9. Applying a rect-function for the aperture leads to (a) ringing artifacts. The Fourier transform of the aperture is shown below. Multiplying an apodization function to the aperture in the spatial domain results in an (b) image free of ringing.

at which two point sources can still be distinguished when diffracting at a circular aperture of diameter D (see [37]).

Apodization: The effect of a rectangular aperture as used in our pipeline can be simulated as a multiplication by a rect-function in the spatial domain and as a convolution with a sinc-function in the frequency domain (see Fig. 9a). The convolution with the sinc-function can lead to visible artifacts called ringing appearing as blur in the axial directions. This blur can be minimized by applying a tapering function as proposed in [38] and [39]. In our implementation, we use a cosine window function being 1 in the middle of the aperture and decreasing to 0 at the border of the aperture resulting in an apodization as shown in Fig. 9b).

Thin Lens: Placing the camera at arbitrary positions while choosing the object appearing in focus requires a lens. We use a thin lens model:

$$\frac{1}{d_o} + \frac{1}{d_i} = \frac{1}{f} \quad (23)$$

with d_o being the distance from the lens to the object, d_i being the distance from the lens to the image, and f being the focal length.

The effect of a thin lens can be simulated by multiplying a complex-valued function $L(x, y)$

$$L(x, y) = e^{-ikr}, \text{ with } r = \sqrt{x^2 + y^2 + f^2} \quad (24)$$

inducing a phase shift to the wavefront [19]. The phase shift transforms the spherical wave with origin at focal distance f into an almost planar wave:

$$\left(A_0 \frac{e^{ikr}}{r} \right) L(x, y) = \frac{A_0}{r}. \quad (25)$$

Even though it does not lead to a perfect planar wave, it results in a good approximation for a point source being far away compared to the size of the lens. Fig. 10 shows a scene consisting of five points being placed at different distances from the lens taken with three different focal lengths. Off-axis points at the edge of the image can produce a comet-shaped image, an effect known as *coma aberration*. For more general applications, one would need to implement a lens simulation, which distributes aberration and artifacts over the whole image.




Fig. 10. Scene with five points placed at different depth and rendered with three different focal lengths f .

Due to a physically-based model of the holography, pipeline effects like *depth of field* or *refocusing* are automatically generated (see Section 8).

7.2 Color Holograms

A nonmonochromatic wave could be split into all spectral frequencies, which could be propagated separately in the wave framework. However, instead of simulating all frequencies separately, we chose three primary colors ($\lambda_r, \lambda_g, \lambda_b$, having approximately the wavelength of red, green, and blue) from which other visible colors can be combined linearly. The amplitudes A_0 at the origin of each wave are set as $A_0 = \sqrt{I}$, where I is the intensity of one of the primary colors. The final image can thus be rendered by calculating the scene for each monochromatic primary color once.

Simulating nonmonochromatic colors introduces another artifact called *chromatic aberration* when using lenses in the optical path. This means that when choosing a focal length f_r for a lens, the effective focal length for the image generated with wavelength λ_g is $f_r \left(\frac{\lambda_g}{\lambda_r} \right)$. Applying a Fourier based image generation and perspective distortion does not introduce further chromatic artifacts.

When choosing the wavelength for the primary colors we are not restricted to take realistic values since the intensity I of the color channel only influences the amplitude. The wavelength λ_r, λ_g , and λ_b could be exactly the same avoiding all aberrations and artifacts related to wavelength differences. In physical lens systems, multiple lenses and lens materials are minimizing the effect of chromatic aberration.

7.3 Evaluate Depth

So far, we have shown the generation of an image of a single object. However, if we have scenes composed of several objects occlusions have to be taken care of. We present a depth reconstruction from a hologram in order to compose objects using the depth buffer.

Depth From Phase: Calculating depth from phase is a straightforward approach also used in digital holographic interferometry as shown in [17]. Knowing the phase of the source and the image, the relative depth can be calculated out of this phase difference modulo λ . In order to unambiguously reconstruct the surface, the depth difference of two neighboring pixels has to be smaller than $\frac{\lambda}{2}$. However, a relative surface is not sufficient for depth composition.

Depth from Phase Difference: A more advanced approach consists in generating several images with slightly different wavelengths λ_i . In our case, we generate two images with two wavelengths λ_1 and λ_2 being related by $\lambda_1 = \epsilon \lambda_2$ with $\epsilon < 1$. The phase φ_i can be calculated as $\varphi_i = \left(\frac{2\pi}{\lambda_i} * r + \varphi_0 \right) \bmod \pm\pi$,

where $\varphi_i \bmod \pm\pi = ((\varphi_i + \pi) \bmod 2\pi) - \pi$. This leads to a phase difference $\Delta\varphi$ in $[0, 2\pi)$ as $\Delta\varphi = 2\pi r \left(\frac{1}{\lambda_1} - \frac{1}{\lambda_2}\right)$. Considering the maximal phase difference 2π , we get the maximal radius r_{max} as

$$r_{max} = \frac{2\pi}{2\pi \left(\frac{1}{\epsilon\lambda_1} - \frac{1}{\lambda_1}\right)}, \text{ with} \quad (26)$$

$$\epsilon = \frac{1}{1 + \frac{\lambda_1}{r_{max}}}. \quad (27)$$

Therefore, the bigger r_{max} is, the closer ϵ has to be to 1 leading to a small difference between λ_1 and λ_2 . Applying this depth reconstruction results in an absolute depth needed for depth composition. Before having a closer look at the influence of interference problems in Section 7.3, we will analyze the impact of a lens on depth reconstruction.

Depth Reconstruction Using a Lens: Applying a lens for image generation imposes a consideration of influence on depth reconstruction as well. Considering the mathematical representation of a lens with focal length f as described in (24), we can find images $U_{\lambda_1}(P)$, $U_{\lambda_2}(P)$ at distance R from the origin as follows:

$$U_{\lambda_j} = A_0 \frac{e^{ik_j R}}{R} e^{-ik_j f}, \text{ where } k_j = \frac{2\pi}{\lambda_j}, j = 1, 2. \quad (28)$$

The phase difference $\Delta\varphi$ evaluates to

$$\Delta\varphi = 2\pi \left(\frac{1}{\lambda_2} - \frac{1}{\lambda_1}\right) (r_f - f) \quad (29)$$

for the phases given as $\varphi_1 = k_1(r_f - f)$ and $\varphi_2 = k_2(r_f - f)$. The distance of the source can thus be found by solving (29) for r_f . Comparing the reconstruction based on a camera without a lens leading to depth r_f with the reconstruction with a lens leading to depth r we find the simple relation to be $r_f = r + f$. In order to find the desired distance values d_{max} around the focal point $[f - d_{max}, f + d_{max}]$, we have to evaluate a new factor ϵ_l such that $\lambda_2 = \epsilon_l \lambda_1$ to:

$$\epsilon_l = \frac{1}{1 + \frac{\lambda_1}{2d_{max}}}. \quad (30)$$

Influence of Interference: Despite of having better results using a depth reconstruction approach based on phase difference, we still have the problem of interference. We explain the case of two interfering point sources PS_1 and PS_2 with origins P_1 and P_2 at distance r_1 and r_2 from the point of evaluation P' (see Fig. 11a) and, finally, infer the general case of possible multiple point sources in the end. Since r_1 and r_2 cannot be reconstructed separately, we would like to find a depth r' which lies in between the two distance values such that $r_1 \leq r' \leq r_2$. This is equal to guaranteeing $\Delta\varphi_1 \leq \Delta\varphi' \leq \Delta\varphi_2$ for $\Delta\varphi_1$ and $\Delta\varphi_2$ being the individual phase differences of PS_1 and PS_2 , respectively, and $\Delta\varphi'$ being the phase difference of the superposition of PS_1 and PS_2 at P' .

The complex valued image U_{λ_j} describes the interference pattern of PS_1 with amplitude A_{1_0} and PS_2 with amplitude A_{2_0} with wavelength λ_j as follows:

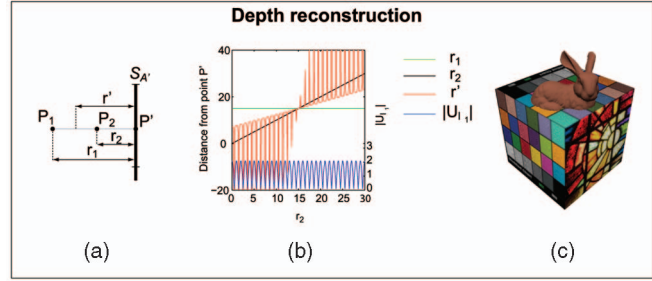


Fig. 11. Moving P_2 in (a) from P' to $r_2 = 30$ is leading to a reconstruction of r' as in (b). Improving results in (b) by multiple iterations for depth refinement leads to results as in (c).

$$U_{\lambda_j} = A_{1_0} \frac{e^{ik_j r_1}}{r_1} + A_{2_0} \frac{e^{ik_j r_2}}{r_2}, \text{ with } k_j = \frac{2\pi}{\lambda_j}. \quad (31)$$

Ignoring interference would lead to a depth reconstruction as seen in Fig. 11b. The desirable values of r' lying in between r_1 and r_2 always coincide with the maxima of $|U_{\lambda_1}|$. This can be shown since $|U_{\lambda_1}| = \max(|U_{\lambda_1}|)$ implies $\min(\Delta\varphi_1, \Delta\varphi_2) \leq \Delta\varphi' \leq \max(\Delta\varphi_1, \Delta\varphi_2)$. The associative property of vector addition allows the extension of two point sources to multiple point sources, so the proof holds for the addition of a new point source with the sum of the remaining point sources as well. In order to improve depth quality, in case of $|U_{\lambda_1}| \neq \max(|U_{\lambda_1}|)$, the image is rendered several times while randomizing the phase of the point sources for every iteration. Depending on the number of iterations the depth reconstruction can be arbitrarily close to an optimal r' .

Finally, the depth buffer is multiplied by a mask created by a thresholding of an introduced opacity channel, verifying that a pixel belongs to a certain object. Note, that once the wave based object has been evaluated the compositing with rasterized objects as depicted in Fig. 11c can be performed in real time.

7.4 Aliasing Free Images

A rasterized image is essentially a discrete function defined on a 2D regular grid. To generate an aliasing free image, the sampling process has to satisfy the Nyquist criterion. Since the Nyquist criterion in traditional computer graphics depends on the frequency of the projected textures, appropriate filtering has to be applied to obtain antialiased images as has been shown by Heckbert [40], [41]. However, in our pipeline, generated images are inherently free of aliasing.

The scenes used in our pipeline can consist of point primitives or textured planar primitives emitting waves of equal wavelength λ . The color of the point sources as well as the planar primitives is encoded in the amplitude of the propagating waves. The interference generated by the superposition of all these primary waves create a wavefront containing the complete information of the whole scene. Since the wavefront does not have a higher frequency than the interference pattern, the sampling rate will not have to be higher either. In order to avoid any aliasing the sampling rate of the interference pattern has to guarantee the Nyquist frequency and therefore a sampling distance of $\frac{\lambda}{2}$, which has been shown to be the case in Section 6.1. By simulating

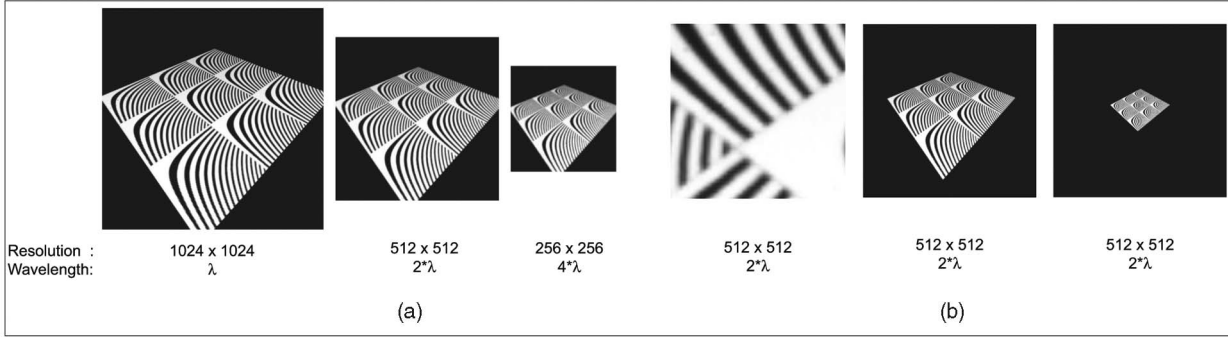


Fig. 12. (a) The three images show the same scene evaluated at different image resolutions and, therefore, different wavelengths. (b) shows different magnifications of the same textured plane without having to apply any filtering. (Please look at the pdf for correct anti-aliasing appearance.)

monochromatic waves, we implicitly guarantee a distinct lower bound for the wavelength λ and, therefore, a distinct upper bound for the required sampling rate $\frac{2}{\lambda}$.

Finally, by Fourier transforming the sampled interference pattern we get the aliasing free directional components of the wavefront. After an according mapping described in Section 7.1, we get the aliasing free image.

Changing image resolution implies a modified frequency in the sampling process of the image function. In traditional computer graphics, this requires a pre or postfiltering of the texture in order to comply with the Nyquist frequency. In our pipeline, however, a lower image resolution does not require any filtering, but simply an increased wavelength of the monochromatic source. In order to avoid a scaling of the object wave, the hologram is regenerated using the new wavelength. Fig. 12a shows three renderings with different resolutions and, therefore, three different wavelengths. We did not apply any filtering on the original texture and still obtain aliasing free images. Furthermore, magnification or minification does not require any filtering or wavelength adjustments as can be seen in Fig. 12b.

7.5 Image Improvement

Throughout our pipeline, coherence is a fundamental property to generate a standing wave which again is necessary to record and reconstruct a hologram. However, generating images based on coherent light can cause visible artifacts due to the static interference pattern called speckle noise, which can even be visible by the naked eye in physical setups as shown in [42].

Speckle Noise Reduction through Higher Resolution: The speckle size δ_S depends on the distance x of two interfering light waves with wavelength λ being at distance z of the interference measurement as follows (see [17]):

$$\delta_S = \frac{\lambda z}{x}. \quad (32)$$

Therefore, the minimal speckle size $\delta_{S_{min}}$ can be calculated dependent on the aperture size W as $\delta_{S_{min}} = \frac{\lambda z}{W}$.

These properties imply a higher frequency of speckle noise with an increasing image resolution meaning a bigger W (see Section 7.1), whereas the frequency of the actual image content stays the same. Hence, filtering out high frequency speckle noise does not affect the actual image content of the rendered image. Additional reduction can be

achieved by down sampling the image using a low pass filter.

Speckle Noise Reduction by Temporal Averaging: Instead of filtering in space as in Section 7.5, it might be faster to filter in time by rendering multiple images from a diffuse object with randomized phases. Computing the final image is achieved by either blending all the images or choosing the maximal value per pixel. The SNR is better for the maximal composition of N images if only two sources with amplitude A_1 and A_2 are contributing to an image point P of image k with intensity $I_k(P)$, since the expectation $E[\lim_{N \rightarrow \infty} \max_k(I_k(P))] = (A_1 + A_2)^2$ is bigger than the expectation for averaging $E[\frac{1}{N} \sum_{k=1}^N I_k(P)] = A_1^2 + A_2^2$. However, the more interfering sources are involved the smaller the probability that the maximal value is found after N iterations. The necessary and sufficient condition for a maximal composition of interfering sources consists in all sources having equal phases as shown in Section 7.3. The signal to noise ratio (SNR) is improving very quickly up to roughly 20 iterations. As stated in Table 1, the SNR for most objects lead to good results around 70 iterations.

8 RESULTS

Using our framework, we are capable of rendering high quality images, while taking into account depth of field and refocusing of the scene by simply adjusting the aperture size or focal length as if a real camera would have been used (see Fig. 13). Even big objects with primitives count up to 300K (see Fig. 14f) can be evaluated in a reasonable amount of time (see Table 1), due to usage of fragment shaders and angular spectrum propagation. By using propagation in frequency space, time can be reduced considerably. A $1,024^2$ hologram as used in Fig. 14e can be propagated in 50s instead of 5min 30s. Creating a hologram of $1,024^2$ with double precision per primary color channel leads to a total size of 24MB. Undersampling of a high-frequency texture during the rasterization step in ray-based image generation can lead to severe aliasing artifacts whereas in a wave-based framework anti-aliasing comes for free Fig. 15c. All the images were generated on a Pentium 4 3.2GHz containing a NVidia GeForce 7800GT. Currently, we make use of fragment shaders for point evaluations.



Fig. 13. The top scene has been calculated with a small aperture leading to a large depth of field. The two images below were computed using a bigger aperture and two different focal lengths.

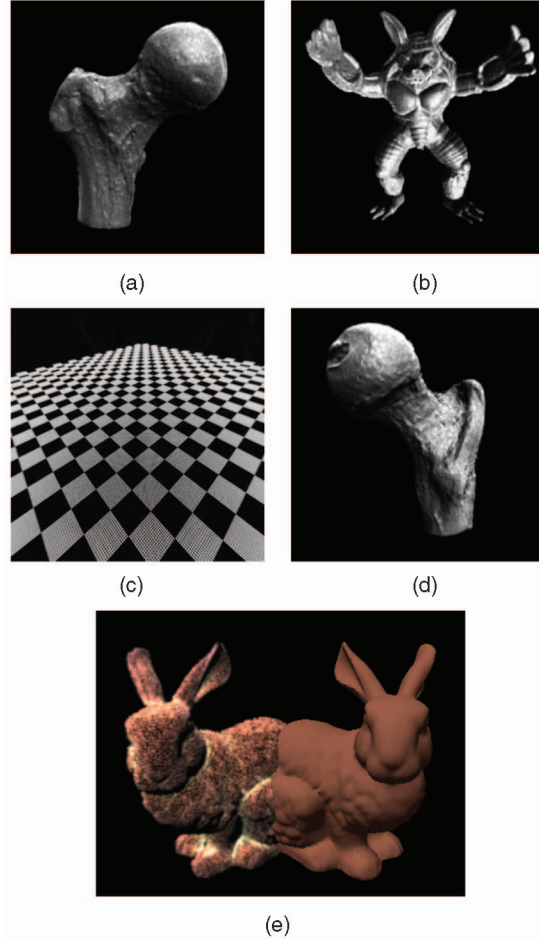


Fig. 15. Wave-based evaluation of point-based objects can represent surfaces of high geometric detail, as long as the sampling of the point-based object is dense enough in the first place. On the lower left (c), an aliasing-free checkerboard generated using our pipeline is shown. (e) is showing a scene composed of a holographic and a mesh based representation of the same object.

9 HOLOGRAPHY VERSUS LIGHT FIELDS

Introducing holograms as a data representation also suggests a comparison to light fields. Independent of the scene complexity, light fields can be acquired quite easily and can be rendered in an efficient way, leading to high quality images. Simulation of depth of field and refocusing is possible, although inadequately sampled camera planes can lead to aliasing. The lack of geometry makes compositing and occlusion difficult and may lead to defocus artifacts. Furthermore, seamless full parallax light fields require a considerable amount of data and proper reconstruction filters for an antialiased image generation.

Holograms on the other hand are inherently free of aliasing. Providing information of phase and amplitude allows to reconstruct depth, which enables scene compositing, occlusion, future reshading or morphing. Further effects, such as depth of field, refocusing or diffraction can be intuitively adjusted by changes of the aperture, focal length, wave length or shape of the aperture, without requiring knowledge of the scene. Although holograms of reasonable size can lead to high quality images, capturing of real objects requires a big amount of data. Additionally,




Fig. 14. (a), (b), (c), and (d) are pairs of images of the same scene rendered with different foci. (b) and (d) are images focused on the holographic plate revealing its rectangular shape. Using a bigger hologram in (e) avoids evanescence of the object at the border of the hologram. The chameleon in (f) is rendered with a bigger depth of field than the wasp in (e).

rendering time is significantly more extensive than rasterized rendering, while the image quality can still suffer from speckle noise.

This short comparison does not outweigh holograms over light fields or vice versa, but rather shows their properties and possibilities of complementing their strength.

10 DISCUSSION AND FUTURE WORK

We presented a novel framework for holographic scene representation and rendering that integrates smoothly into conventional graphics architectures. The wave nature of holographic image generation provides a variety of interesting features including depth of field and high quality anti-aliasing. The phase information inherently embedded into the representation allows us to reliably reconstruct view-dependent depth and to compose digital holograms with standard 2D framebuffer graphics. To the best of our knowledge, this is the first pipeline which allows the rendering of holographic representations as well as traditional geometric objects in the same scene. While holograms require a very high sampling rate in theory, our practical implementation computes high quality renderings from reasonably sized digital holograms.

In spite of hardware acceleration, holographic image generation is still far from real-time, which sets one of the limitations of the presented method. In future work, we plan to explore the parallel nature of wave evaluation to further accelerate the rendering. In addition, relighting, reshading, and deformation operators in wave domain are of highest interest to us. Finally, efficient compression of digital holograms could be achieved by exploiting their spectral composition.

ACKNOWLEDGMENTS

The authors would like to thank Tim Weyrich for valuable discussions as well as Miguel A. Otaduy for the received advice while writing this paper.

REFERENCES

- [1] M. Levoy and P. Hanrahan, "Light Field Rendering," *Computer Graphics*, vol. 30, pp. 31-42, 1996, citeseer.ist.psu.edu/levoy96light.html.
- [2] D.N. Wood, D.I. Azuma, K. Aldinger, B. Curless, T. Duchamp, D.H. Salesin, and W. Stuetzle, "Surface Light Fields for 3D Photography," *Proc. 27th Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '00)*, pp. 287-296, 2000.
- [3] A. Isaksen, L. McMillan, and S.J. Gortler, "Dynamically Reparameterized Light Fields," *Proc. 27th Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '00)*, pp. 297-306, 2000.
- [4] S.J. Gortler, R. Grzeszczuk, R. Szeliski, and M.F. Cohen, "The Lumigraph," *Computer Graphics*, vol. 30, pp. 43-54, 1996, citeseer.ist.psu.edu/gortler96lumigraph.html.
- [5] C. Buehler, M. Bosse, L. McMillan, S. Gortler, and M. Cohen, "Unstructured Lumigraph Rendering," *Proc. 28th Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '01)*, pp. 425-432, 2001.
- [6] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar, "Acquiring the Reflectance Field of a Human Face," *Proc. 27th Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '00)*, pp. 145-156, 2000.
- [7] W. Matusik, H. Pfister, A. Ngan, P.A. Beardsley, R. Ziegler, and L. McMillan, "Image-Based 3D Photography Using Opacity Hulls," *ACM Trans. Graphics*, vol. 21, no. 3, pp. 427-437, 2002.
- [8] T. Weyrich, "Rendering Deformable Surface Reflectance Fields," *IEEE Trans. Visualization and Computer Graphics*, vol. 11, no. 1, pp. 48-58, Jan./Feb. 2005.
- [9] T. Okoshi, *Three-Dimensional Imaging Techniques*. Academic Press, 1976.
- [10] S.A. Benton, "Selected Papers on Three-Dimensional Displays," *SPIE—Int'l Soc. for Optical Eng.*, 2001.
- [11] C. Slinger, C. Cameron, and M. Stanley, "Computer-Generated Holography as a Generic Display Technology," *Computer*, vol. 38, no. 8, pp. 46-53, Aug. 2005.
- [12] N. Delen and B. Hooker, "Free-Space Beam Propagation between Arbitrarily Oriented Planes Based on Full Diffraction Theory: A Fast Fourier Transform Approach," *J. Optical Soc. Am. A*, vol. 15, pp. 857-867, Apr. 1998.
- [13] K. Matsushima, "Computer-Generated Holograms for Three-Dimensional Surface Objects with Shade and Texture," *Applied Optics*, vol. 44, pp. 4607-4614, Aug. 2005.
- [14] M. Lucente and T.A. Galyean, "Rendering Interactive Holographic Images," *Proc. 22nd Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH '95)*, pp. 387-394, 1995.
- [15] M. Lucente, "Interactive Three-Dimensional Holographic Displays: Seeing the Future in Depth," *SIGGRAPH Computer Graphics*, vol. 31, no. 2, pp. 63-67, 1997.
- [16] M. Halle and A. Kropp, *Fast Computer Graphics Rendering for Full Parallax Spatial Displays*. 1997, citeseer.csail.mit.edu/halle97fast.html.
- [17] U. Schnars and W. Jüptner, *Digital Holography: Digital Hologram Recording, Numerical Reconstruction, and Related Techniques*. Springer, 2005.
- [18] D. Gabor, "A New Microscope Principle," *Nature*, vol. 161, pp. 777-779, 1948.
- [19] J.W. Goodman, *Introduction to Fourier Optics*. McGraw-Hill Book Company, 1968.
- [20] T. Tommasi and B. Bianco, "Frequency Analysis of Light Diffraction between Rotated Planes," *Optics Letters*, vol. 17, pp. 556-558, Apr. 1992.
- [21] T. Tommasi and B. Bianco, "Computer-Generated Holograms of Tilted Planes by a Spatial Frequency Approach," *J. Optical Soc. Am. A*, vol. 10, pp. 299-305, Feb. 1993.
- [22] K. Matsushima, H. Schimmel, and F. Wyrowski, "Fast Calculation Method for Optical Diffraction on Tilted Planes by Use of the Angular Spectrum of Plane Waves," *Optical Soc. Am. J. A*, vol. 20, pp. 1755-1762, Sept. 2003.
- [23] F. Durand, N. Holzschuch, C. Soler, E. Chan, and F. Sillion, "A Frequency Analysis of Light Transport," *ACM Trans. Graphics (Proc. SIGGRAPH Conf.)*, vol. 24, no. 3, Aug. 2005, <http://artis.imag.fr/Publications/2005/DHSC505>.
- [24] K. Matsushima, "Exact Hidden-Surface Removal in Digitally Synthetic Full-Parallax Holograms," *Proc. SPIE*, vol. 5742, pp. 53-60, Practical Holography XIX: Materials and Applications, pp. 25-32, June 2005.
- [25] T. Totsuka and M. Levoy, "Frequency Domain Volume Rendering," *Computer Graphics*, vol. 27, pp. 271-278, 1993, citeseer.ist.psu.edu/totsuka93frequency.html.
- [26] S. Yusuke, I. Masahide, and Y. Toyohiko, "Color Computer-Generated Holograms from Projection Images," *Optics Express*, vol. 12, pp. 2487-2493, May 2004.
- [27] K. Matsushima and M. Takai, "Recurrence Formulas for Fast Creation of Synthetic Three-Dimensional Holograms," *Applied Optics*, vol. 39, pp. 6587-6594, Dec. 2000.
- [28] M. Koenig, O. Deussen, and T. Strothotte, "Texture-Based Hologram Generation Using Triangles," *Proc. SPIE*, vol. 4296, pp. 1-8, Practical Holography XV and Holographic Materials VII, S.A. Benton, S.H. Stevenson, and T.J. Trout, eds., pp. 1-8, June 2001.
- [29] A. Ritter, J. Böttger, O. Deussen, M. König, and T. Strothotte, "Hardware-Based Rendering of Full-parallax Synthetic Holograms," *Applied Optics*, vol. 38, pp. 1364-1369, Feb. 1999.
- [30] E. Cuhe, P. Marquet, and C. Depeursinge, "Spatial Filtering for Zeroorder and Twin-Image Elimination in Digital Off-Axis Holography," *Applied Optics*, vol. 39, pp. 4070-4075, Aug. 2000.
- [31] T. Kreis and W. Jüptner, "Suppression of the DC Term in Digital Holography," *Optical Eng.*, vol. 36, pp. 2357-2360, Aug. 1997.

- [32] U. Schnars and W. Jüptner, "Digital Recording and Numerical Reconstruction of Holograms," *Measurement Science and Technology*, vol. 13, pp. 85-101, 2002.
- [33] V. Vaish, B. Wilburn, N. Joshi, and M. Levoy, "Using Plane + Parallax for Calibrating Dense Camera Arrays," *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, vol. 1, pp. 2-9, 2004.
- [34] M. Levoy, B. Chen, V. Vaish, M. Horowitz, I. McDowall, and M. Bolas, "Synthetic Aperture Confocal Imaging," *ACM Trans. Graphics*, vol. 23, no. 3, pp. 825-834, 2004.
- [35] R. Ng, "Fourier Slice Photography," *ACM Trans. Graphics*, vol. 24, no. 3, pp. 735-744, 2005.
- [36] J.A. Ratcliffe, "Some Aspects of Diffraction Theory and Their Application to the Ionosphere," *Reports on Progress in Physics*, vol. 19, no. 1, pp. 188-267, 1956, <http://stacks.iop.org/0034-4885/19/188>.
- [37] P.A. Tipler, *Physics for Scientists and Engineers*, third ed., extended version. Worth Publishers, Inc., 1991.
- [38] G. Gbur, *OPTI 6271/8271 Lecture Notes Spring 2005*, UNCC, Jan. 2005.
- [39] G. Gbur and P. Carney, "Convergence Criteria and Optimization Techniques for Beam Moments," *Pure and Applied Optics*, vol. 7, pp. 1221-1230, Sept. 1998.
- [40] P.S. Heckbert, "Fundamentals of Texture Mapping and Image Warping," Master's thesis, 1989, citeseer.ist.psu.edu/150876.html.
- [41] N. Greene and P.S. Heckbert, "Creating Raster Omnimax Images from Multiple Perspective Views Using the Elliptical Weighted Average Filter," *IEEE Computer Graphics and Applications*, vol. 6, no. 6, pp. 21-27, 1986.
- [42] J. Dainty, A. Ennos, M. Françon, J. Goodman, T. McKechnie, and G. Parry, *Laser Speckle*. Springer-Verlag, 1975.



projective technology. He is member of the IEEE and the ACM.



Peter Kaufmann studied computer science at the ETH Zurich, Switzerland, specializing in computational science, computer graphics, and computer vision. He received the MS degree at ETH Zurich in computer science in 2005, and is currently working in the fields of GPU-based real-time video processing, multiprojector setups, and distributed systems in London, England.



Markus Gross received the MS degree in electrical and computer engineering and the PhD degree in computer graphics and image analysis, both from the University of Saarbrücken, Germany. He is a professor of computer science and director of the Computer Graphics Laboratory of the ETH Zurich, Switzerland. From 1990 to 1994, Dr. Gross worked for the Computer Graphics Center in Darmstadt, where he established and directed the Visual Computing Group. His research interests include point-based graphics, physics-based modelling, multiresolution analysis, and virtual reality. He has been widely publishing and lecturing on computer graphics and scientific visualization, and he authored the book *Visual Computing* (Springer, 1994). Dr. Gross has taught courses at major graphics conferences, including ACM SIGGRAPH, IEEE Visualization, and Eurographics. He is the associate editor of the *IEEE Computer Graphics and Applications* and has served as a member of international program committees of many graphics conferences. Dr. Gross has been a papers cochair of the IEEE Visualization 1999, the Eurographics 2000, and the IEEE Visualization 2002 conferences. He chaired the papers committee of ACM SIGGRAPH 2005. He is a senior member of the IEEE.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.