SIGGRAPH: G: Enabling Reflective & Refractive Depth Representation in Computer-Generated Holography

Aaron Demolder
Hammadi Nait-Charif
Valery Adzhiev
Bournemouth University
Bournemouth, UK

Alfred Newman
Tom Durrant
Andrzej Kaczorowski
VividQ Ltd
Cambridge, UK

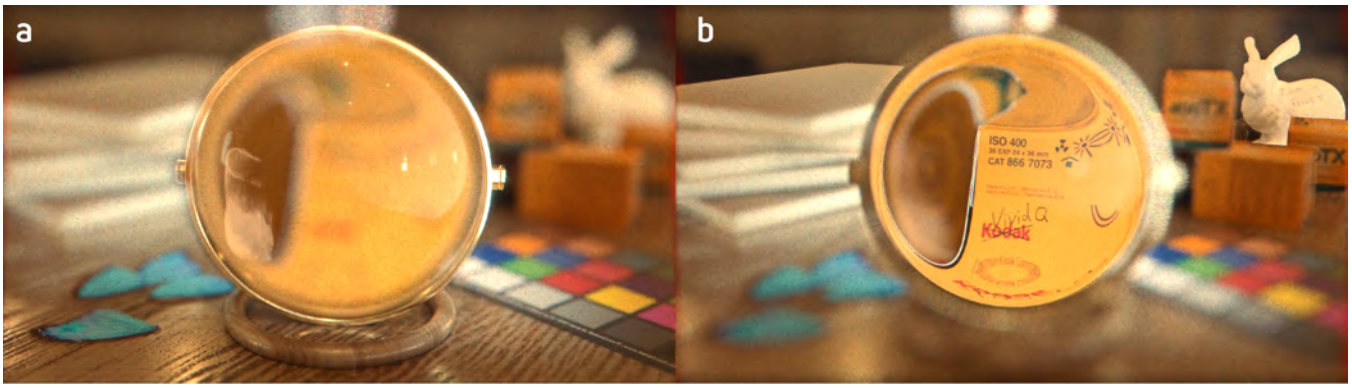

Figure 1: Holographic replays (simulated) of a scene containing a lens of f=180mm, 0.5m from camera. (a) Focus is at 0.46m on the smudge on the magnifying glass surface, the focus of the reflected lights is 0.52m. (b) Focus (1m) now on the text printed on the film box and bunny statue, the film box is physically placed 150mm behind the lens.

1 PROBLEM & MOTIVATION

Computer-Generated Holography (CGH) is a technique for the reconstruction and display of three-dimensional imagery through diffraction and interference of light. Holograms are generated digitally by calculating light propagation, and displayed on devices called Spatial Light Modulators (SLMs). When illuminated by coherent laser light, holograms then replay a full 3D scene. Recently, holography has received significant attention from industrial and academic communities, with neural holography [Peng et al. 2020] [Shi et al. 2021] and numerous other techniques [Sahin et al. 2020]. Applications of holography include Augmented/Virtual Reality [Widjanarko et al. 2020], Head-Up Displays, and larger display devices [An et al. 2020]. CGH offers "the most complete and visually satisfying 2D record of a 3D scene we know how to make"[Benton and Bove 2008], and is of particular interest as it allows a realisation of the progress towards realism in computer graphics to be translated into realistic 3D viewing.

In the real world, objects can be viewed in reflections (such as in mirrors) and through refractions (such as through glass or water); the depth at which these objects come into focus is defined by the focal power of the material in question. When rendering convincing imagery in computer graphics, realistic representation of these materials is essential to the realism of the final image. These material representations do not pose an issue for rendering for 2D display, as only the pixel intensities are required - and focus is predetermined by a depth of field routine in the renderer, or by post processing. Doing this post processing requires the same depth data of the scene as in image/layer-based CGH, where z-depth values are used that correspond to the focus distance of objects. While this technique works well for the first object hit by the raytracer, it is unable to accurately encode the depth at which virtual objects appear after reflection or refraction. It is hence necessary, in order to create perceptually correct scenes, for the renderer to pass multiple depth values into the pipeline to allow the hologram generation engine to ensure realistic representation of these reflections and refractions.

2 BACKGROUND & RELATED WORK

In computer graphics, a three-dimensional scene is rendered into a 2D image, just as a camera captures an image in the real world. This occurs with realistic lighting effects by utilising a raytracer, or via a faster approximation using a rasteriser. The traversal of the ray-tree [Whitted 1980] allows this ray-tracing simulation, where a light ray will be simulated travelling through the scene interacting with various objects and material properties. This produces a final intensity for each pixel in an image. In this realistic light simulation, objects become visible to the viewer through reflections on the surfaces of other objects, which may not directly be within the direct perspective view of the camera, for example. Conventional rasterising approaches to fast scene rendering are typically only aware of objects directly within this camera perspective [Fernando 2004]. To improve the realism of rasterised scenes, techniques such as light baking can be used to simulate realistic lighting in advance.

2.1 Depth of Field

Simulating a depth of field effect is possible in either of these cases - in a ray tracer by utilising a thin-lens approximation [Heidrich et al. 1997], and in rasterisers by post processing - making various approximations of blurring with a point spread function [Kosloff and Barsky 2010] that utilises the Z depth. Here, "Z" refers to the Z axis of the camera coordinate space. This Z value can be normalised to fit within an 8-bit variable (where 0 represents a “z-near” value, and 255 represents “z-far”), or represented in floating point in real units, typically meters. In computer-generated holography, where a field of light is projected in 3D space, any depth of field effects visible are by result of the identical process that occurs when a viewer focuses on objects in real life, and are caused by the viewer focusing at various depths within the projected scene.

2.2 Layer-Based Computer-Generated Holography

The layer-based method of computer-generated holography is the rasterising equivalent of ray-traced holography, as it displays an approximation of the 3D scene which has already been baked with any realistic lighting. Imagery from a ray-tracer can be input into a layer-based hologram generation process by utilising this same Z depth value, but instead uses the Z values to reproject pixels into 3D space. There is only one depth value per pixel, and this is usually the depth of the first geometry hit in the scene. This is sufficient for diffuse material representation, as all of the depth information is available within the perspective view, but cannot reproduce the correct depth information for any lighting effects achieved via a ray-tree, such as reflections or refractions. This is the case with any existing method of layer or point based holography [Kaczorowski 2018; Peng et al. 2020; Shi et al. 2021; Zhang et al. 2017].

2.3 Holographic Raytracing

In a holographic ray-tracer, this issue is resolved by directly forming a hologram from the scene representation rather than going via an image or pointcloud [Ichikawa et al. 2013; Nishi and Matsushima 2017; Wang et al. 2018; Zhang et al. 2016]. There are also hybrid approaches [Gilles et al. 2016]. These produce realistic effects, such as correct depths for objects in reflections - but with trade-offs. This includes image quality, as shader development in such renderers is not mature and generally not easily compatible with existing production methods, and computation time is greatly increased, as the simulation is extremely complex.

3 APPROACH & UNIQUENESS

We propose a method of utilising additional depth information required for reflective and refractive material representations on holographic displays via multiple render passes, and a first-of-its-kind method of multi-pass holographic compositing whereby render passes are combined in the hologram (frequency) domain. The resulting hologram successfully represents the input render, but now with the correct depth information for all render passes supplied. The resulting hologram can consequently be displayed on an SLM and viewed by eye with the correct perceptual cues, resulting in correct focus properties (see Figure 1) – where viewers can focus as expected on reflections or on objects through refractive materials.

Table 1: Render Passes for Composite

RGB Channels	Depth Source
Total Lighting ((Raw Lighting + Raw Global Illumination) * Diffuse Filter)	Z-depth (primary ray depth)
Reflection (Raw Reflection * Reflection Filter)	Reflection secondary ray depth + primary ray depth
Specular Highlights	Specular secondary ray depth + primary ray depth
Refraction (Raw Refraction * Refraction Filter)	Refraction Secondary ray depth + primary ray depth
Others (such as SSS, Emission, Caustics if required)	Z-depth (primary ray depth)

3.1 Depth of Primary & Secondary Rays

The ‘depth’ of a point refers to the distance from which multiple rays from that point appear to diverge. In the case of an object being viewed in the reflection of a flat mirror, the total depth is the length from the object to the mirror and then to the viewer. Figure 2(a) shows the primary ray from viewer to mirror, and a secondary ray [Whitted 1980] from mirror to object. A z-depth pass will provide the length of only the primary ray, so must be added to the length of the reflection secondary ray for the correct depth of the object in camera space.

A refraction depth pass is performed in a similar fashion, as seen in Figure 2(b). In the case of looking through glass, the primary ray is formed by the depth from viewer to the surface of the glass, but the length of the secondary ray is determined by the shape of the surface as well as the refractive index. The resulting image may be real or virtual, and have magnification. The focal length of the lens is given by the lensmaker’s equation:

$$\frac{1}{f} = (n-1)\left(\frac{1}{R_1} - \frac{1}{R_2}\right) \tag{1}$$

Where f is focal length, n is Index of Refraction (IOR) and R_1, R_2 are the radii of curvature of the surface. The new secondary ray depth of the object image visible through the lens is given by:

$$\frac{1}{S_2} = \frac{1}{f} - \frac{1}{S_1} \tag{2}$$

Where S_1 and S_2 are the depths to the lens and the object respectively. When added, S_1 and S_2 provide the total correct depth value. Most renderers will already store a value for focus depth per pixel per render-pass per trace-depth as part of the ray tree.

3.2 Holographic Compositing

Table 1 shows the combination of render passes commonly used in the compositing process, and the ray depths that are combined into the associated depth sources. Each colour pass can still be adjusted as required in compositing software. The resulting passes can then be loaded for hologram generation - here we employ VividQ’s Software Development Kit. The resulting outputs are images in Fourier space (where the image is decomposed into its constituent frequencies) for each render pass which can be displayed on an SLM.

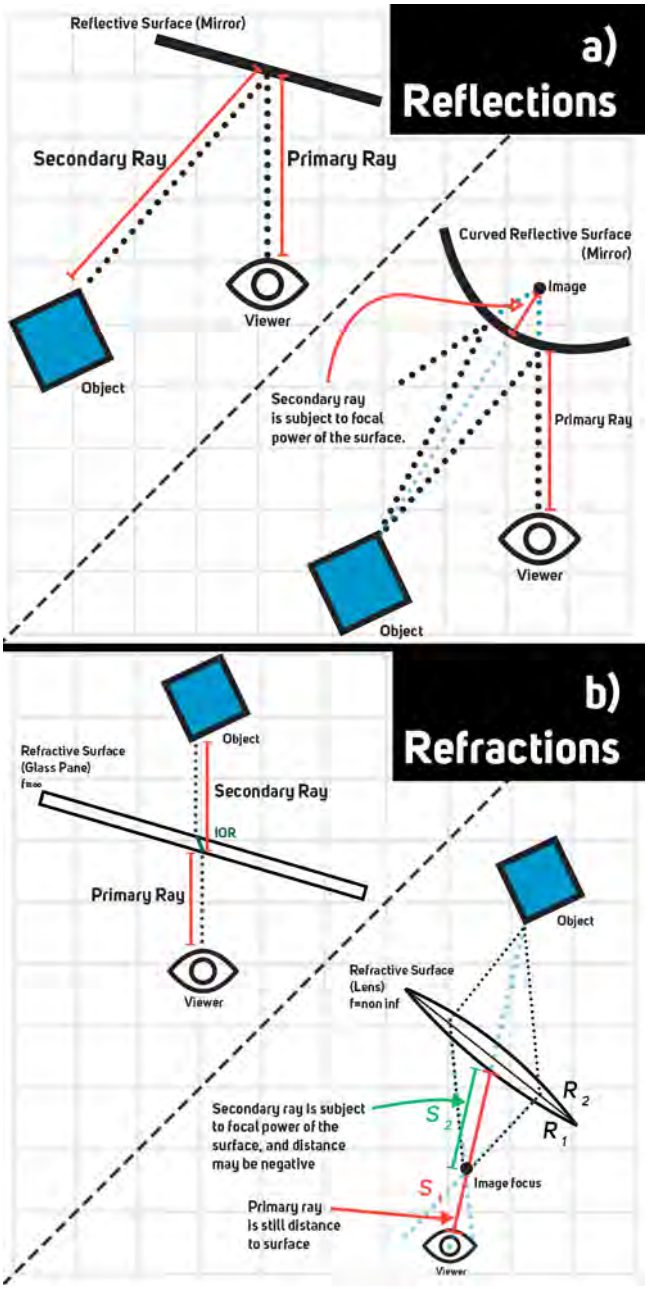

Figure 2: (a) The total depth required for reflections is the summation of primary and secondary rays, following the law of reflection and any focal power. (b) The focal length is calculated with the Index of Refraction (IOR) given by Snell's Law, and the thickness of the material. Distances of secondary rays may be negative. A material may exhibit both reflective and refractive properties simultaneously.

As holograms are additive in the same way light is - each of the holograms can consequently be composited using a familiar multi-pass process, in the same fashion as a renderer will composite the light passes, and in the same fashion as performed by a compositing artist when handling render files. However, this composite occurs whilst the hologram is still in frequency space rather than image space, this is shown in Figure 3. When the full composite is displayed on an SLM, this will 'replay' the full intended 3D scene to the viewer, with the complete correct material depth information.

This is unique in its approach as it successfully integrates an existing image renderer more closely in the hologram generation process. By taking the render passes provided by the renderer, they can be combined in frequency space to allow correct depth representations. Using this method allows existing hologram generation engines that provide only diffuse representation of scenes to also provide these reflective and refractive representations.

4 RESULTS & CONTRIBUTIONS

The results shown in both Figure 1 and Figure 4 demonstrate that the methods introduced represent a significant step towards full realism of 3D scenes in holographic display, with physically correct reflections and refractions - whilst being readily compatible with the existing computer graphics pipeline. While a fully featured holographic raytracer could in the future create such imagery, full integration of such a renderer would be significant work, as well as requiring significantly more computational power. By utilising image/layer-based holography for fast hologram processing, combined with the proposed technique, the range of materials available to be accurately viewed is expanded. This innovation enables the delivery of some of the most realistic, high-quality imagery to any holographic display to date - using image data which is already available in the renderer that would usually be discarded before display.

4.1 Additional Applications

Image-based hologram generation methods result in black artefacts in a scene (such as around the edges of the magnifying glass in Figure 1) as there is no colour information behind any of the depth layers, as it is a 2.5D perspective input. This could be minimised by ensuring some overscan is present in render passes per-object per-render-pass for any occluded objects. Additional data passes can be sourced directly from the renderer as presented in this paper, and assisted via passes such as cryptomatte [Friedman and Jones 2015] or object ID manifests [Hillman 2018]). These artefacts are most noticeable in projections that have significant defocus, or where the viewer is off-axis from the render perspective of the scene camera. Methods such as holographic stereograms [Kang et al. 2008; Padmanaban et al. 2019; Yamaguchi et al. 1993], which provide new perspectives to the hologram resolve this. Here reflections and refractions can be composited into the hologram to demonstrate view-dependant effects.

Whilst the presented method is demonstrated using offline renderers aimed at productions such as cinema or episodical media, with techniques in real-time raytracing this method can be utilised for delivering correct depth information to holograms with reflections and refractions from sources such as game-engines [Nguyen

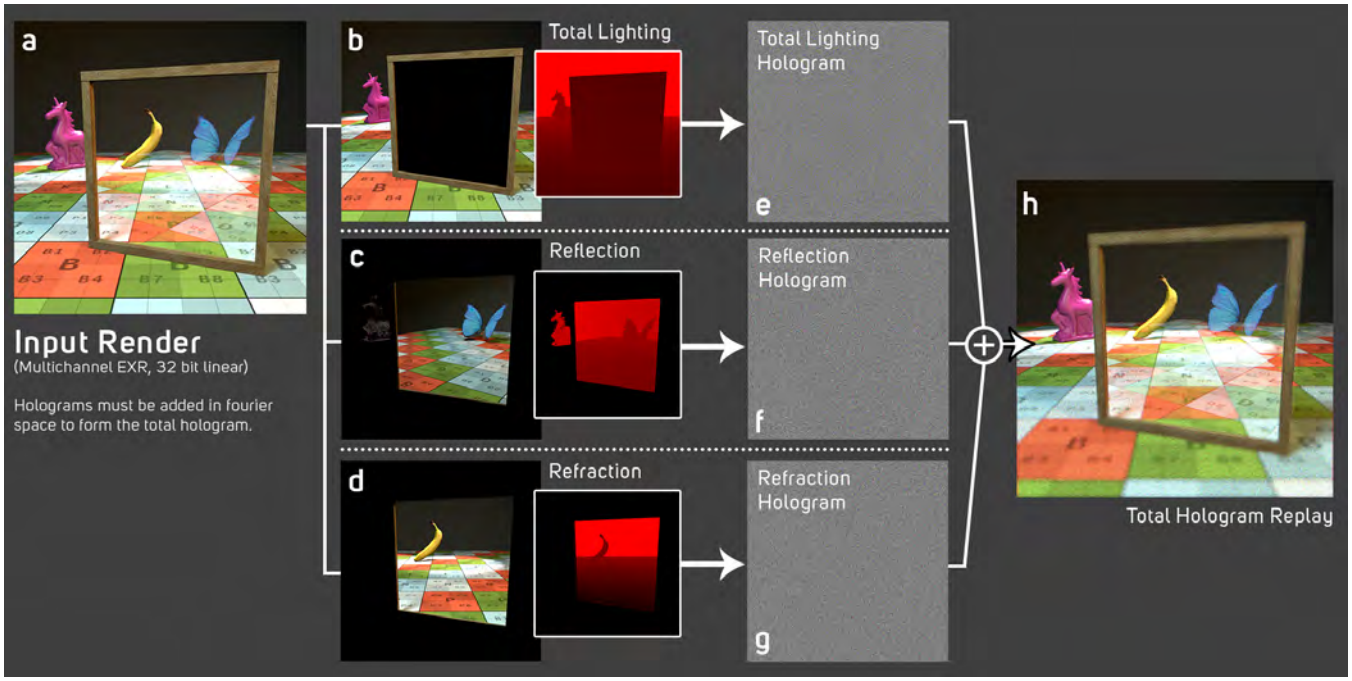

Figure 3: (a) The input multichannel file. (b) Total Lighting render, z-depth. (c) Reflection render, associated depth. (d) Refraction render, associated depth. (e) Hologram of Total Lighting only. (f) Hologram of Reflection only. (g) Hologram of Refraction Only. (h) The total composite hologram replay, via a plus operation of the holograms for each render pass. This is the final image that is viewable on the holographic display

2007]. Depths can be used in both layer-based and full 3D pointwise holograms. Layer-based CGH is currently the only viable method for delivering real-time performance, this is particularly applicable for AR/VR applications [Maimone et al. 2017] for both consumers and enterprise alike - where accurate material representations provide realistic visualisations and immersive experiences.

4.2 Research Impact

The results presented are part of a collection of work tied to my Engineering Doctorate Digital Media, a collaboration between the Centre for Digital Entertainment, Bournemouth, UK, and holographic display company VividQ, Cambridge, UK - which aims to deliver the most comprehensive research and practical methods in media production for CGH. These efforts have resulted in 3 award-winning first-author publications. This includes: artistic control of holographic bokeh [Demolder et al. 2020], which was awarded Best Short Paper at ACM CVMP 2020; standardisation of formats, pipeline integration, HDR, Colour Management and archive in CGH [Demolder 2022], which was awarded Best Student Paper 2021 by the Society of Motion Picture and Television Engineers (SMPTE); and this work, which received the Gold Medal at ACM SIGGRAPH 2021 in the graduate research competition. Each of the developments presented in the listed previous works have been used to deliver the imagery in this publication using existing production renderers. Additional papers are in preparation, alongside a patent, both of which continue this line of work and will continue to provide some of the best visual results in CGH.

ACKNOWLEDGMENTS

This work has been supported by grant funding of the Centre for Digital Entertainment by the Engineering and Physical Sciences Research Council (EPSRC) EP/L016540/1 and VividQ Ltd.

REFERENCES

Jungkwuen An, Kanghee Won, Young Kim, Jong-Young Hong, Hojung Kim, Yongkyu Kim, Hoon Song, Chilsung Choi, Yunhee Kim, Juwon Seo, Alexander Morozov, Hyunsik Park, Sunghoon Hong, Sungwoo Hwang, Kichul Kim, and Hong-Seok Lee. 2020. Slim-panel holographic video display. *Nature Communications* 11, 1 (2020).

Stephen A. Benton and V. Michael Bove. 2008. *Holographic Imaging*. John Wiley & Sons, Inc. https://doi.org/10.1002/9780470224137

Aaron Demolder. 2022. Toward the Standardization of High-Quality Computer-Generated Holography Media Production Workflow. *SMPTE Motion Imaging Journal* 131, 1 (2022), 48–58.

Aaron Demolder, Andrzej Kaczorowski, Alfred Newman, and Hammadi Nait-Charif. 2020. Artistic Control of Defocus in Computer-Generated Holography. ACM CVMP.

Randima Fernando. 2004. *GPU Gems: Programming Techniques, Tips and Tricks for Real-Time Graphics*. Pearson Higher Education.

Jonah Friedman and Andrew C. Jones. 2015. Fully Automatic ID Mattes with Support for Motion Blur and Transparency. In *ACM SIGGRAPH 2015 Posters (SIGGRAPH '15)*. Article 47.

Antonin Gilles, Patrick Gioia, Rémi Cozot, and Luce Morin. 2016. Computer generated hologram from multiview-plus-depth data

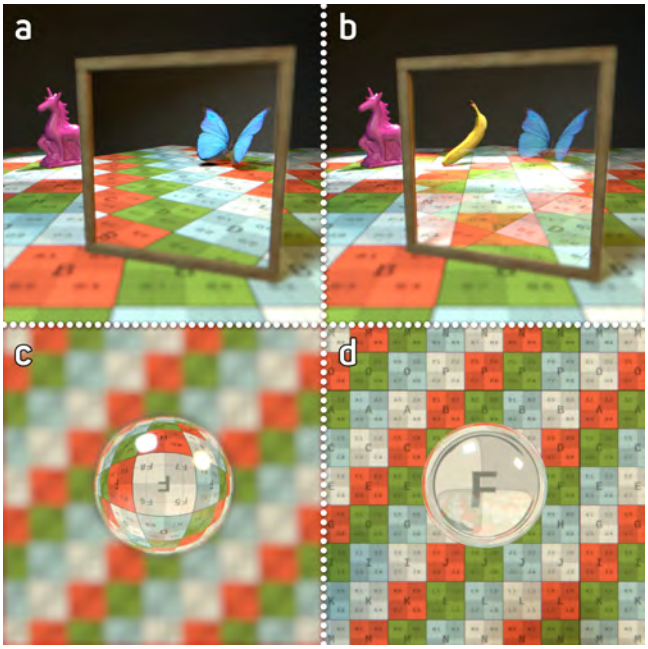

Figure 4: (a) Replay of butterfly reflection and unicorn (both 9m) (b) Changing the mirror to glass reveals a banana (9m). (c) Grid at 1m from camera focused by a lens (f=100mm) placed 850mm from the camera, the grid focuses at 550mm (real). Reflection focus is 950mm. (d) Grid, 1m from camera focused by lens (f=100mm) placed 950mm from camera, the grid focuses at 1050mm (virtual). Reflection focus is 1050mm.

considering specular reflections. In *2016 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*. IEEE, 1–6.

Wolfgang Heidrich, Philipp Slusallek, and Hans-Peter Seidel. 1997. An Image-Based Model for Realistic Lens Systems in Interactive Computer Graphics. In *Proceedings of the Conference on Graphics Interface '97* (Kelowna, British Columbia, Canada). Canadian Information Processing Society, CAN, 68–75.

Peter Hillman. 2018. A Scheme for Storing Object ID Manifests in OpenEXR Images. In *Proceedings of the 8th Annual Digital Production Symposium (DigiPro '18)*. Article 9.

Tsubasa Ichikawa, Kazuhiro Yamaguchi, and Yuji Sakamoto. 2013. Realistic expression for full-parallax computer-generated holograms with the ray-tracing method. *Appl. Opt.* 52, 1 (Jan 2013), A201–A209.

Andrzej Kaczorowski. 2018. *Adaptive aberration correction for holographic projectors.* Ph.D. Dissertation. University of Cambridge.

Hoonjong Kang, Takeshi Yamaguchi, and Hiroshi Yoshikawa. 2008. Accurate phase-added stereogram to improve the coherent stereogram. *Appl. Opt.* 47, 19 (Jul 2008), D44–D54.

Todd J Kosloff and Brian A Barsky. 2010. Three techniques for rendering generalized depth of field effects. In *Proceedings of the 2009 SIAM Conference on "Mathematics for Industry" The Art of "Mathematics for Industry"*. SIAM, 42–48.

Andrew Maimone, Andreas Georgiou, and Joel S. Kollin. 2017. Holographic Near-Eye Displays for Virtual and Augmented Reality. *ACM Trans. Graph.* 36, 4, Article 85 (jul 2017), 16 pages.

Hubert Nguyen. 2007. *Gpu Gems 3* (first ed.). Addison-Wesley Professional.

Hirohito Nishi and Kyoji Matsushima. 2017. Rendering of specular curved objects in polygon-based computer holography. *Appl. Opt.* 56, 13 (May 2017), F37–F44.

Nitish Padmanaban, Yifan Peng, and Gordon Wetzstein. 2019. Holographic Near-Eye Displays Based on Overlap-Add Stereograms. *ACM Trans. Graph.* 38, 6, Article 214 (nov 2019), 13 pages.

Yifan Peng, Suyeon Choi, Nitish Padmanaban, and Gordon Wetzstein. 2020. Neural Holography with Camera-in-the-Loop Training. *ACM Trans. Graph.* 39, 6, Article 185 (2020).

Erdem Sahin, Elena Stoykova, Jani Mäkinen, and Atanas Gotchev. 2020. Computer-Generated Holograms for 3D Imaging: A Survey. *ACM Comput. Surv.* 53, 2, Article 32 (2020).

Liang Shi, Beichen Li, Changil Kim, Petr Kellnhofer, and Wojciech Matusik. 2021. Towards real-time photorealistic 3D holography with deep neural networks. *Nature* 591, 7849 (2021).

Yuan Wang, Zhidong Chen, Xinzhu Sang, Hui Li, and Linmin Zhao. 2018. High-efficiency photorealistic computer-generated holograms based on the backward ray-tracing technique. *Optics Communications* 410 (2018), 768–773.

Turner Whitted. 1980. An Improved Illumination Model for Shaded Display. *Commun. ACM* 23, 6 (1980).

T. Widjanarko, M. El Guendy, A. O. Spiess, D. M. Sullivan, T. J. Durrant, O. A. Tastemur, A. J. Newman, D. F. Milne, and A. Kaczorowski. 2020. Clearing key barriers to mass adoption of augmented reality with computer-generated holography. In *Optical Architectures for Displays and Sensing in Augmented, Virtual, and Mixed Reality (AR, VR, MR)*, Vol. 11310. SPIE.

Masahiro Yamaguchi, Hideshi Hoshino, Toshio Honda, and Nagaaki Ohyama. 1993. Phase-added stereogram: calculation of hologram using computer graphics technique. In *Practical Holography VII: Imaging and Materials*, Stephen A. Benton (Ed.), Vol. 1914. International Society for Optics and Photonics, SPIE, 25 – 31.

Hao Zhang, Liangcai Cao, and Guofan Jin. 2016. Lighting effects rendering in three-dimensional computer-generated holographic display. *Optics Communications* 370 (2016), 192–197.

Hao Zhang, Liangcai Cao, and Guofan Jin. 2017. Computer-generated hologram with occlusion effect using layer-based processing. *Appl. Opt.* 56, 13 (May 2017), F138–F143.