

MIT Open Access Articles

A compressive light field projection system

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Hirsch, Matthew, Gordon Wetzstein, and Ramesh Raskar. "A Compressive Light Field Projection System." *ACM Transactions on Graphics* 33, no. 4 (July 27, 2014): 1–12.

As Published: <http://dx.doi.org/10.1145/2601097.2601144>

Publisher: Association for Computing Machinery

Persistent URL: <http://hdl.handle.net/1721.1/92775>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



A Compressive Light Field Projection System

Matthew Hirsch*

Gordon Wetzstein*
MIT Media Lab

Ramesh Raskar




Figure 1: Compressive light field projection for glasses-free 3D display. The system comprises a single light field projector and a completely passive screen. The angular range of the light field emitted from the projector is limited to the size of the projection lens aperture, hence very small. Keplerian telescopes inspire our screen design—the angular range of incident light is expanded for an observer on the other side, creating a field of view that is suitable for glasses-free 3D display. A prototype projector was implemented from scratch using two high-speed spatial light modulators (SLMs); a prototype screen was fabricated from two lenticular sheets with different focal lengths, mounted back-to-back. With the implemented system, we achieve high-rank light field synthesis (center) for human observers with a critical flicker fusion threshold that is smaller than the product of the SLM refresh rates and the rank of the synthesized light field. Note that color results above are composited from multiple images captured from our grayscale prototype.

Abstract

For about a century, researchers and experimentalists have strived to bring glasses-free 3D experiences to the big screen. Much progress has been made and light field projection systems are now commercially available. Unfortunately, available display systems usually employ dozens of devices making such setups costly, energy inefficient, and bulky. We present a compressive approach to light field synthesis with projection devices. For this purpose, we propose a novel, passive screen design that is inspired by angle-expanding Keplerian telescopes. Combined with high-speed light field projection and nonnegative light field factorization, we demonstrate that compressive light field projection is possible with a single device. We build a prototype light field projector and angle-expanding screen from scratch, evaluate the system in simulation, present a variety of results, and demonstrate that the projector can alternatively achieve super-resolved and high dynamic range 2D image display when used with a conventional screen.

CR Categories: B.4.2 [Hardware]: Input/Output and Data Communications—Image display; I.3.1 [Mathematics of Computing]: Numerical Linear Algebra; I.3.1 [Computer Graphics]: Hard-

ware Architecture—Three-dimensional displays; I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms;

Keywords: compressive displays, light fields

Links: [DL](#) [PDF](#) [WEB](#) [VIDEO](#) [DATA](#) [CODE](#)

1 Introduction

Within the last few years, 3D movie theaters have become so popular and wide-spread that most new movies are released in 3D; even classics are often re-rendered to fit the increasing demand for 3D content. For many people, the experience of watching a 3D movie on a large screen is significantly more immersive than conventional 2D screenings or watching smaller-scale 3D content on TV. Commercially available 3D projection technology is based on stereoscopic principles. An image pair depicting a scene from slightly different perspectives is displayed such that observers perceive each of these images with a different eye. Most often, this is achieved by optically encoding the image pair in different polarization states or color spectra and then decoding it with special glasses worn by the observer. This approach can create viewer discomfort; furthermore, the correct perspective is only observed from a single sweet-spot in center of the theater.

As opposed to stereoscopic image generation, light field displays provide physically correct views for a wide range of perspectives and do not require an observer to wear special glasses. Interestingly, inventors worldwide have investigated large-scale light field projection systems throughout the last century [Funk 2012]. Several light field movie theaters were open to the public in Russia and France in the 1940s. Most of these and subsequent installations employ large parallax barrier-type screens, resulting in severe loss of image resolution and light throughput. Today, larger-scale light field projection systems are commercially available but require dozens of devices [Balogh 2006], making these systems expensive, power hungry, bulky, and difficult to calibrate.

*The indicated authors acknowledge equal contributions by sharing first authorship.




Figure 2: Overview of prototype light field projection system. The projector (right) emits a 4D light field with a narrow field of view that only varies over the projection lens (b, Nikkor 35mm $f/1.4$ AI-s). This angular range is expanded by the screen (left) for an observer on the other side. The screen (a) is composed of passive pixels that each expand the angles of all incident light, just like a Keplerian telescope. No special calibration w.r.t. the projector is necessary beyond focusing the latter on the screen. The projector emits a 4D light field, which is synthesized by two reflective spatial light modulators (SLMs, Silicon Micro Display ST1080). Their contribution is optically combined by a 1:1 relay lens (h, $2\times$ Canon EF 50 mm $f/1.8$ mounted face-to-face). The light source (10W LED) is synchronized to the refresh rate (240 Hz) of the SLMs by a custom board (e). The SLMs use liquid crystal on silicon (LCoS) technology, which requires polarizing beam splitter cubes (c), and are connected to a standard graphics card via a driver board (d).

Inspired by recent advances in compressive light field display, we present the first compressive light field projection system. The proposed system combines a novel, passive screen, a single high-speed light field projector, and light field factorization algorithms. The employed factorization routines directly exploit redundancy in the target content. Just like image and video compression, this allows us to represent light fields in a compressed form that is automatically computed via low-rank factorization. However, the factorizations in our approach directly map to pixels states of the optical system, thereby not only reducing the memory footprint needed to store the light fields but also the number of projection devices. Hence, the proposed system is compressive in a computational and an optical sense. Through the co-design of display optics and computational processing, we devise a practical solution to large-scale light field display.

Our primary contribution is to introduce and characterize a compressive light field projection system. Additional technical contributions include the following:

- By combining new optical designs (passive screens and modified high-speed projectors) with compressive light field factorization, we demonstrate the first compressive glasses-free 3D projection system.
- We introduce a new screen design that is inspired by Keplerian telescopes; the screen optically expands the field of view of a light field projected on it. It is completely passive, thin, and could be fabricated in large-scale.
- We analyze the proposed setup and show that it is more light efficient than alternative time-sequential parallax barrier approaches and also relaxes requirements on display refresh rates.

- We also demonstrate that the proposed projector, in combination with conventional diffuse screens, allows for superresolution and high dynamic range 2D image projection.
- We design and fabricate a prototype projection system consisting of a custom high-speed, LCoS-based duallayer light field projector and an angle-expanding screen.

2 Related Work

Light field displays Glasses-free 3D or light field displays were invented more than a century ago. Gabriel Lippmann [1908] and Frederic Ives [1903] were the first to discover how arrays of lenslets or pinholes can be used to synthesize a light field. These technologies have evolved into head-tracked [Perlin et al. 2000] and large-scale multi-monitor [Sandin et al. 2005] systems. However, one of the main shortcomings of these approaches is the loss of image resolution—spatial resolution is traded for angular light field resolution. Hence, most commercial systems employ dynamically switchable 2D/3D technologies (e.g., [Jacobs et al. 2003]). Alternative light field display configurations include spinning mirrors [Cosairt et al. 2007; Jones et al. 2007], time-multiplexed diffusers [Sullivan 2003], multi-focal displays [Akeley et al. 2004], adaptive optics [Tompkin et al. 2013] and time-multiplexed shutters [Travis 1990]. A state-of-the-art review of 3D displays can be found in Urey et al. [2011] and in Masia et al. [2013]. With the proposed projection system, we adopt a compressive display approach to overcome limitations in achievable resolution and required cost, power, and form factor over existing technologies.

Compressive light field displays Through the co-design of display optics and computational processing, compressive displays

strive to transcend limits set by purely optical designs. It was recently shown that tomographic light field decompositions displayed on stacked films of light-attenuating or polarizing materials can achieve higher resolutions than previously possible [Wetzstein et al. 2011; Lanman et al. 2011]. Compression is achieved in the number of layer pixels, which is significantly smaller than the number of emitted light rays. Low-rank light field synthesis was also demonstrated for duallayer [Lanman et al. 2010] and multi-layer displays with directional backlighting [Wetzstein et al. 2012; Maimone et al. 2013]. In these display designs, an observer perceptually averages over a number of patterns that are displayed at refresh rates beyond the critical flicker frequency of the human visual system (HVS). The limited temporal resolution of the HVS is directly exploited by decomposing a target light field into a set of patterns, by means of nonnegative matrix or tensor factorization, and presenting them on high-speed spatial light modulators; this creates a perceived low-rank approximation of the target light field. Conceptually, these compressive displays extend earlier work on time-multiplexed 3D display (e.g., [Travis 1990]) using modern mathematical approaches to low-rank factorization; display brightness is improved and requirements on display refresh rates relaxed. Inspired by compressive displays, we design and fabricate the first compressive light field projection system comprised of modified, active projectors and a novel passive screen.

Light field projection Large-scale autostereoscopic and multi-scope projection systems have been actively investigated throughout the last century. Most of the proposed systems are variants of integral imaging or parallax barriers; by combining active projection and large barrier screens, theater-sized installations have been built in France and Russia starting in the 1940s [Funk 2012]. Today, barrier-type light field projection systems are still an active area of research (e.g., [Yang et al. 2008; Kim et al. 2012]). The fundamental limitation of these displays, as every other integral imaging or barrier-based method, is the loss of spatial resolution. 2D/3D switchable solutions for projectors have been proposed [Hong et al. 2010] but require multiple devices.

Resolution limits can be overcome using multiple projectors combined with front or rear-projected lenticular screens [Matusik and Pfister 2004; Hsu 2008] or unidirectional diffusers [Balogh 2006; Jones et al. 2013]. In these systems, the number of devices roughly matches the number of viewing zones. Projector arrays can also be directly observed [Jurik et al. 2011], but require one device per pixel. Dodgson et al. [2000] investigate multi-projector devices combined with time-multiplexed image synthesis for 3D display. Compared to existing multi-device solutions, we present a new optical configuration that is well-suited for compressive light field synthesis with a single device.

Single-device configurations have been explored [Nims and Lo 1972; Cossairt and Favallora 2006; Meuret et al. 2010; Bogaert et al. 2010]. These methods project individual viewing zones sequentially and at a high-speed onto special screens. Unfortunately, these screens require mechanically-moving parts that translate in unison with the high-speed projection. We present a compressive light field projection system that requires only a single device (it can be enhanced using a few additional devices), operates at the full display resolution, and does not require active components in the screen. The proposed display combines a novel screen design based on Keplerian angle expansion with high-speed light field projection and compressive factorizations. In addition, we show applications to 2D superresolution and high dynamic range projection.

Superlenses Though not widely used today, back-to-back lenticular sheets have been explored historically for their unique imag-

ing properties. Dennis Gabor [1944] demonstrated that varying the pitch and focal length of back-to-back lenticular sheets can create configurations that perform analogously to physically larger standard lens systems. Such arrangements are described as Gabor superlenses in optics literature [Hembd-Sölner et al. 1999]. In a closely related work Eichenlaub et al. [2005] demonstrate that superlenses can be used to enlarge a volumetric display, though other means of generating large volumetric displays have been shown [Kimura et al. 2006; Smoot et al. 2011]. We present a computational and optical system for light field projection. One part of this system incorporates an angle expanding screen and one possible implementation of such a screen is a superlens composed of back-to-back lenticular sheets.

Superresolution and high dynamic range display Approaches to superresolution projection have been described in previous work. Examples include optical configurations that combine the contribution of multiple overlapping devices [Damara-Venkata and Chang 2009] or single devices with either two stacked liquid crystal displays (LCDs) [Sajadi et al. 2012] or one LCD and a double-lens system [Sajadi et al. 2013]. Superresolution display with monitors, as opposed to projectors, can be achieved by fast mechanical motion of the screen [Berthouzoz and Fattal 2012] or using two stacked LCDs with a diffuser mounted on top [Heide et al. 2014]. In Section 7, we briefly outline that the proposed light field projector in combination with a conventional diffuse screen can achieve superresolution projection. Our approach is closely related to [Sajadi et al. 2012] and [Heide et al. 2014]; as opposed to the former we use formal optimization to derive optimal pixel states in the display and compared to the latter, demonstrate superresolution on a diffuse projection screen rather than a monitor.

In addition, we also demonstrate that the proposed projection system is capable of high dynamic range display, which is inspired by the methods proposed by Seetzen et al. [2004]. We note that both superresolution projection and HDR display have been described in the past. In Section 7 we simply demonstrate that our projection system is flexible enough to support these applications *in addition* to glasses-free 3D display.

3 Compressive Light Field Synthesis

In this section, we derive the optical image formation of the proposed system as well as related optimization techniques. The formulations are derived in 2D “flatland”, but extensions to the full 4D case are straightforward.

3.1 Optical Image Formation

Consider a conventional rear-projection system. The projection lens re-images and magnifies the pattern displayed on an internal spatial light modulator (SLM). A diffusing transmissive screen is placed at the conjugate plane of the SLM, such that an image can be observed over a wide range of viewing angles from the other side of the screen. The light field on the viewer side is engineered to be as *view-independent* as possible and directly corresponds to the SLM image g :

$$\tilde{l}(x, \nu) = g(x). \quad (1)$$

While this approach is effective for presenting two-dimensional images, we are interested in emitting *view-dependent* 4D light fields. For this purpose, two modifications to conventional projection systems are necessary. First, the projector has to emit a light field and not just a 2D image. Second, the screen has to preserve the incident angular variation. Unfortunately, diffusing screens in most existing projection setups optically average an incident light field in the




Figure 3: Overview of light field projection system. Two spatial light modulators, g and h , synthesize a light field inside a projector (top right). The projection screen is composed of an array of angle-expanding pixels (bottom). Inspired by Keplerian telescopes, these pixels expand the field of view of the emitted light field for an observer on the other side of the screen.

angular domain and eliminate high frequency directional variation. To overcome this limitation, we introduce the notion of an angle-preserving screen that changes the image formation to

$$\tilde{l}(x, \nu) = g(\phi(x, \nu)), \quad (2)$$

where each light ray (x, ν) on the observer side of the screen is mapped to the SLM inside the projector by the function $\phi: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$. We adopt a two-plane parameterization of the light field, where x is the spatial coordinate on the screen and $\nu = \tan(\theta)$ the point of intersection with a relative plane at unit distance (see Fig. 3).

In addition to the angle-preserving screen, the projector also needs to be modified so as to emit a light field. Such projectors have been proposed in the past; possible options include microlenses or a pinhole mask near the image SLM and coded projector apertures (e.g., [Grosse et al. 2010]). We follow Lanman et al. [2010] and use two programmable, light-attenuating SLMs inside the projector (see Fig. 3). The image formation is now given by the multiplication the the patterns g and h shown on the two SLMs:

$$\tilde{l}(x, \nu) = g(\phi(x, \nu)) h(\psi(x, \nu)). \quad (3)$$

Similar to ϕ for g , $\psi: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ maps each ray in the light field to a position on the second SLM h . Using ray transfer matrices [Hecht 2002], these mapping functions can easily be derived given the distance between screen and aperture d , the ray transfer matrix of the screen \mathbf{T}_s , the focal length of the projection lens f_p , and the distance d_g from the aperture to the SLM:

$$\begin{pmatrix} \phi(x, \nu) \\ \zeta \end{pmatrix} = \begin{pmatrix} 1 & d_g \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -1/f_p & 1 \end{pmatrix} \begin{pmatrix} 1 & d \\ 0 & 1 \end{pmatrix} \mathbf{T}_s \begin{pmatrix} x \\ \nu \end{pmatrix} \quad (4)$$

The incident ray angle ζ is disregarded in the following; ψ is similar to ϕ but replaces d_g by d_h . All system parameters are illustrated

in Figure 3. While Equation 4 models the ray transfer under the assumption of perfect optics, aberrations can be incorporated into ϕ as well.

3.2 Angle-expanding Screen Design

Independent of the specific method of light field synthesis within the projector, the resulting light field will have a narrow angular range that varies only over the aperture of the device. Unfortunately, this limited range is insufficient for an observer to freely move and enjoy glasses-free 3D display within a reasonable field of view. To address this problem, we propose a screen that not only preserves angular variation but *expands* it.

Angle expansion is a common technique in optics that is for instance used in Keplerian telescopes. These telescopes perform angle expansion with two lenses of different focal lengths mounted such that the distance between them is equal to the sum of their focal lengths. Inspired by this idea, we propose a screen that comprises an array of miniature angle-expanding telescopes—one for each screen pixel. This design is illustrated in Figure 3 (close-up). Whereas the spatial extent of a beam incident from the right is reduced, its incident angle is amplified on the observer side of screen. The ray transfer matrix \mathbf{T} of such a Keplerian angle expander can be modeled as

$$\mathbf{T} = \begin{pmatrix} 1 & 0 \\ -1/f_s^{(p)} & 1 \end{pmatrix} \begin{pmatrix} 1 & f_s^{(p)} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & f_s^{(o)} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -1/f_s^{(o)} & 1 \end{pmatrix} \quad (5)$$

where $f_s^{(o)}$ and $f_s^{(p)}$ are the focal lengths of screen lenslets facing the observer and the projector, respectively. As illustrated in Figure 3, a simple design of the proposed screen uses two lenslet or lenticular arrays with the same lens pitch but different focal lengths. Mounted back to back and with a lens size corresponding to the pixel size on the screen, \mathbf{T}_s becomes

$$\mathbf{T}_s = \begin{pmatrix} 1 & 0 \\ 0 & -f_s^{(o)}/f_s^{(p)} \end{pmatrix} \quad (6)$$

Note that the dependence on ray position in a single angle-expander vanishes for the entire screen in Equation 6, because each lenslet has the same size as a projected image pixel. The refractive effect of the proposed screen only depends on the incident ray angle (i.e. $\nu_p = -f_s^{(o)}/f_s^{(p)} \nu_o$), which is flipped and amplified by an angle-expansion factor of $M = f_s^{(p)}/f_s^{(o)}$. Although the screen is fundamentally limited by diffraction, this effect is negligible in the proposed system because pixels on projection screens are usually large (millimeters as opposed to microns).

3.3 Efficient Light Field Synthesis

The most intuitive way to generate a light field inside the duallayer projector is to display an array of pinholes on one screen and the interlaced views of the light field on the other [Ives 1928]. Unfortunately, this approach generates low-resolution images and is also extremely light-inefficient. We follow recent proposals on light field factorization that optimize both resolution and image brightness. These algorithms have previously been described for TV-sized dual and multilayer liquid crystal displays (LCDs) [Lanman et al. 2010; Wetzstein et al. 2012]; we briefly review light field factorization in this section and adapt the factorization algorithm to the proposed system by incorporating the effects of projection lens and angle-amplifying screen via mapping functions ϕ and ψ .

Specifically, the image formation (Eq. 3) is discretized as

$$\tilde{\mathbf{I}} = (\Phi \mathbf{g}) \circ (\Psi \mathbf{h}), \quad (7)$$

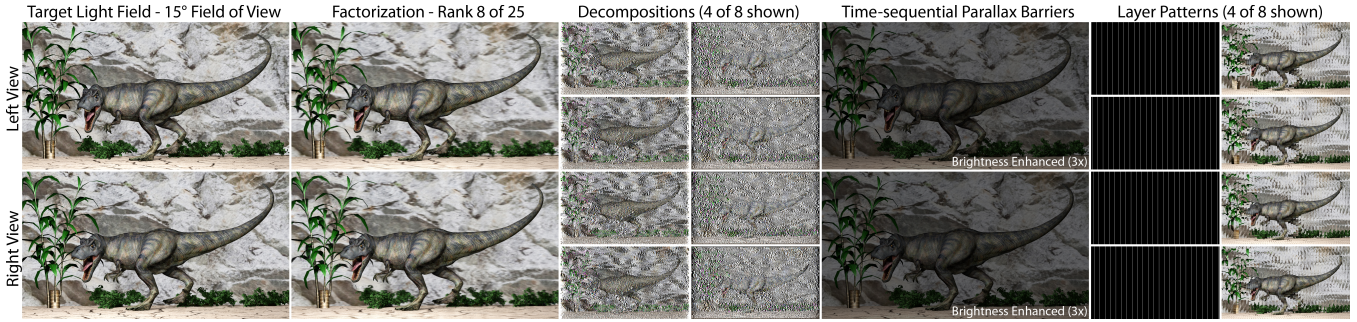


Figure 4: Light field factorization and comparison to time-sequential parallax barriers. Two views of a target light field with horizontal-only parallax and 25 views equally distributed over a field of view of 15° are shown on the left. Using the framework proposed in this paper, the light field is decomposed into a set of patterns for two spatial light modulators (SLMs) running at 480 Hz (center). When observed by a human, these decompositions create a rank-8 approximation of the light field (center left). The alternative to factorized image synthesis is display of time-sequential parallax barriers (right), which are $7.5\times$ darker than our method and require 1500 Hz SLMs to achieve the same resolution (center right).

where $\Phi \in \mathbb{R}^{L \times N}$ and $\Psi \in \mathbb{R}^{L \times M}$ are matrices that permute the rows of the discrete SLM patterns $\mathbf{g} \in \mathbb{R}^N$ and $\mathbf{h} \in \mathbb{R}^M$ according to the mapping in $\phi(x, \nu)$ and $\psi(x, \nu)$, respectively, and \circ is the Hadamard or element-wise product. In this notation, the emitted light field is represented as a discrete vector $\tilde{\mathbf{I}} \in \mathbb{R}^L$. The matrices Φ and Ψ are sparse (usually one non-zero value per row) and constructed via raytracing for simulations (Eqs. 4, 5) or using calibration that accounts for optical aberrations in practice (Sec. 5).

Equation 7 makes clear that the emitted light field is the product of two permuted vectors, hence rank-1. Following recent proposals on light field factorization, we employ high-speed SLMs that operate at refresh rates beyond the critical flicker frequency of the human visual system. Images displayed at such refresh rates are perceptually averaged. In particular, we model high-speed SLMs in the proposed setup as

$$\tilde{\mathbf{I}} = \frac{1}{T} \sum_{t=1}^T (\Phi \mathbf{g}_t) \circ (\Psi \mathbf{h}_t) \quad (8)$$

Here, T pairs of displayed patterns are averaged by the visual system and create a perceived rank- T light field $\tilde{\mathbf{I}}$. The temporally-changing patterns on the SLMs at time t are \mathbf{g}_t and \mathbf{h}_t . Given a target light field $\mathbf{l} \in \mathbb{R}^L$, an optimization problem can be formulated to find the best set—in a least-squared error sense—of time-varying patterns as

$$\begin{aligned} & \underset{\{\mathbf{g}, \mathbf{h}\}}{\text{minimize}} && \left\| \beta \mathbf{l} - \sum_{t=1}^T (\Phi \mathbf{g}_t) \circ (\Psi \mathbf{h}_t) \right\|_2^2 \\ & \text{subject to} && 0 \leq g_{ik}, h_{jk} \leq 1, \forall i, j, k \end{aligned} \quad (9)$$

Note that β absorbs the factor $1/T$ as well as a user-defined brightness boost (see Fig. 5). The nonnegativity constraints ensure that optimized patterns are physically feasible. Although this is a non-linear and nonconvex problem, it is biconvex in \mathbf{g} and \mathbf{h} ; fixing one results in a convex problem for updating the other. Such updates are usually performed in an alternating and iterative manner. We derive multiplicative matrix update rules for our problem as:

$$\begin{aligned} \mathbf{g}_t & \leftarrow \mathbf{g}_t \circ \frac{\Phi^T (\beta \mathbf{l} \circ (\Psi \mathbf{h}_t))}{\Phi^T (\tilde{\mathbf{I}} \circ (\Psi \mathbf{h}_t)) + \epsilon} \\ \mathbf{h}_t & \leftarrow \mathbf{h}_t \circ \frac{\Psi^T (\beta \mathbf{l} \circ (\Phi \mathbf{g}_t))}{\Psi^T (\tilde{\mathbf{I}} \circ (\Phi \mathbf{g}_t)) + \epsilon} \end{aligned} \quad (10)$$




Figure 5: Quantitative evaluation of convergence and brightness boosting factor β in Equation 9. In this example, the proposed update rules converge after about 200 iterations (left). The brightness of the target light field can be boosted as compared to conventional, time-sequential methods; a higher brightness, however, results in a slight decrease in reconstructed light field quality.

where \circ and $-$ denote element-wise product and division, respectively, ϵ is a small value that prevents division by zero, and $\tilde{\mathbf{I}}$ is computed via Equation 8.

Multiplicative update rules for nonnegative matrix factorization problems have become increasingly popular in the scientific computing community (e.g., [Lee and Seung 1999; Cichocki et al. 2009]). We extend these methods by including the projection matrices Φ and Ψ into the solver. The update rules in Equation 10 are *mathematically distinct* but *numerically equivalent* to the conventional multiplicative update rules, as for instance used by Lanman et al. [2010]. However, our formulation has the advantage of not only allowing for an elegant mathematical formulation of arbitrary optical setups, but also for extremely *efficient implementations*. As discussed in more detail in Section 5 and the supplement, Φ and Ψ can be implemented as a multiview rendering step whereas Φ^T and Ψ^T correspond to projective texture mapping. These operations are hardware-accelerated on the GPU and can be implemented in real-time. Source code and pseudo-code for these update rules can be found in the supplemental material.

4 Analysis

We simulate factorized light field synthesis for the proposed projection system in Figure 4. For this experiment, we decompose a light field with 25 views (two of them shown) into eight pairs of time-multiplexed patterns. This choice simulates 480 Hz spatial light modulators that create a rank-8 light field approximation for an observer with a critical flicker frequency of 60 Hz. Device di-

mensions match those of the prototype (see Sec. 5). The light field can be reproduced with a high image quality (center left, PSNR 26.4 dB) using this configuration. In comparison, a time-sequential parallax barrier display would require 25 time-multiplexed images to achieve the same resolution and 1500 Hz SLMs. In addition, the normalized brightness boost β/T for the factorized result is chosen to be 0.3, which makes the observed light field $7.5\times$ brighter than the parallax-barrier display mode.

Looking at the layer decompositions (Fig. 4) shows how light-inefficient parallax barriers are. While the patterns for one of the SLMs contain the interlaced views of the light field, the other comprises a set of vertical slits that block most of the light (images may appear black in printout). The factorized patterns are much more light efficient but less intuitive. As observed for previous approaches to dual and multilayer light field factorization [Lanman et al. 2010; Wetzstein et al. 2012], we interpret the patterns as distributing low image frequencies in the 3D scene to the closest SLM while depth discontinuities in the light field create high-frequency, temporally-varying structures. These can be interpreted as content-adaptive parallax barriers that are automatically created where needed: around edges and scene features that extrude from the physical device.

We also plot the convergence of the proposed algorithm in Figure 5 (left). After about 200 iterations, no significant improvements in image quality, measured in peak signal-to-noise (PSNR) ratio, are observed. The brightness boosting factor β (see Eqs. 9, 10) can be freely chosen to trade 3D image quality for brightness. We analyze this tradeoff in Figure 5 (right). A normalized value of β/T of 0.2 – 0.3 results in high-quality reconstructions. For the example shown in Figure 4, we chose $\beta/T = 0.3$ which results in a direct brightness boost of factor $7.5\times$ over conventional time-sequential parallax barriers.

We also show a quantitative evaluation of light field compressibility in Figure 6. Both horizontal-only and full parallax light fields are considered for decompositions with rank 6, 12, and 18. In all cases, the target light field has 25 views equally spaced over the entire 2D field of view (FOV). Intuitively, light fields containing only horizontal parallax are much more compressible, which is confirmed by higher PSNR values. As the FOV increases, compressibility of the light field decreases due to larger parallax. The small “bumps” in the left plots are discretization artifacts.




Figure 6: Light field compressibility. We simulate reconstructions of the “t-rex” scene for a varying field of view. The target light field has either 25 or 5×5 views equally distributed in a horizontal-only (left) or horizontal and vertical (right) viewing zone, respectively. Horizontal-only parallax (HOP) light fields are much more compressible; higher-rank decompositions achieve a better quality. We observe that rank-6 decompositions with HOP for fields of view up to 20° achieve high-quality reconstructions.

5 Implementation

Our prototype projection system comprises two optical hardware parts which can conceptually be implemented independently of one another: an angle-expanding screen and a light field projector. This section provides recipes for both parts. In the supplemental material, we discuss additional details and outline an alternative light field projector implementation that can be used with the same screen. Source code is also provided in the supplement.

Light Field Projector The projector places two spatial light modulators (SLMs) at different distances behind a projection lens to create angular variation across the lens aperture. As is apparent in Figure 3, the field of view of the system will be maximized by choosing a projection lens with a large aperture relative to its focal length or, similarly, a small f-number. For a fixed screen distance, the image size will be maximized with a shorter focal length lens. We choose a Nikon Nikkor 35mm f/1.4 AI-s lens for our prototype (Fig. 2, b).

The SLMs are reflection mode Liquid Crystal on Silicon (LCoS) modulators (Silicon Micro Display ST1080, Fig. 2 g). To achieve an optical path equivalent to that of Figure 3 with reflective modulators, we employ two polarizing beamsplitter cubes (Fig. 2, c). The physical extent of the beamsplitter cubes requires an additional 1:1 relay lens to optically place both SLMs close to each other. The f-number of the relay lens should match that of the projection lens. We use two Canon EF 50mm f/1.8 II lenses mounted face to face (Fig. 2, h). Although this compound relay lens limits the f-number of the system, it provides high image quality and minimizes optical aberrations. The ST1080 modulator operates at 240Hz and is driven by a driver board (Fig. 2, d) that is intended to run the LCoS for a head mounted display. Assuming a critical flicker fusion rate of about 40Hz for the human visual system, which is reasonable for low-light conditions, the available refresh rates allow a rank-6 monochrome light field decomposition.

The illumination unit in the projector has to match the f-number of the system. It should also be uniform over its spatio-angular extent and be synchronized with the frame updates of the SLMs, meaning the illumination source must be switchable at 240Hz. These constraints can be met with high-power LEDs; we place a 10W LED (similar models can be purchased from Cree, Inc.) mounted on a heat sink behind a mirrored light pipe (Fig. 2, f). The light pipe is taken out of a conventional projector and acts as a “kaleidoscope”, virtually cloning the LED to a larger illumination range. Care is taken to place the LED image out-of-focus with any of the SLM planes, screen, or viewer location. Additional off-the-shelf lenses are used to form a converging beam on the rear SLM. A custom circuit board (Fig 2, e) employs a microcontroller and a power field-effect transistor to switch the LED in sync with the frame updates of the SLM.

Angle-expanding Screen In principle, a horizontal-only expander can be implemented by placing two lenticular sheets of different focal lengths back-to-back (Fig. 2, a). However, the design tolerances of off-the-shelf lenticulars make it difficult to fabricate angle-expanding screens with suitable characteristics in practice. We were able to have a horizontal-only angle-expanding screen with an expansion power of $M = 3$ custom manufactured by Microsharp Innovation. To support a range of vertical viewpoints, the screen requires an additional vertical-only diffuser. We use a horizontally oriented 100 lpi 31° Lenstar Plus 3D lenticular from Pacur. Alternatively, holographic uni-directional diffusers, for instance from Fusion Optix or Luminit, can be used.

Unlike typical lenticular displays, the proposed screen does not de-




Figure 7: Top: parts of the prototype angle-expanding screen: a) Fresnel lens, back-to-back lenticular sheets on b) projector-side, and c) viewer-side, and d) overlaid vertical diffuser. Bottom: test images captured from the extreme viewing angles demonstrating parallax. The closeups show vertical stripes caused by the lenticular of the angle expanding screen. These are not apparent when observed by eye.

crease the resolution of the projected image. Ideally, each lenslet or lenticular has the size of a projected pixel on the screen. To maximize image resolution, these properties should be optically matched. In our current implementation, the projection lens is located 50cm away from the screen and produces an image with a size of 21.3×11.7 cm. The lenticular pitch of the screen is 0.5mm, which currently limits the achieved image resolution in the prototype to 426×720 pixels. A larger image size or smaller lenticulars could increase this resolution.

The screen lenticulars of our prototype have the same pitch. However, to achieve a viewing zone at a distance from the screen greater than that of the projector, the pitch should be adjusted such that the screen acts as an angle-expander and simultaneously as a lens focusing illumination into the viewing zone. For the prototype setup, we can achieve the same effect with an additional Fresnel lens mounted close to the screen on the projector side. The entire optical stack can be seen in Figure 7 and the supplemental document illustrates optical rays paths for various screen configurations, with and without the Fresnel lens. With complete design freedom the entire screen optics could equivalently be fabricated as a single, large-scale sheet in a roll-to-roll process.

System Calibration Creating a projection system from scratch relies on careful calibration of each component. An optical rail system constrains many unnecessary degrees of freedom in the prototype projector. Approximate alignment of each component is achieved by probing with a laser. Once both SLM images can be observed through the projection lens, a checkerboard pattern is projected to focus both SLMs independently and overlay them precisely. As a final verification step, a bar target (Fig. 7) is displayed on the prototype and photographed. The top bars, displayed on SLM 1 (Fig. 2) are in sharp focus on the screen, while the bottom bars form a virtual image in front of the screen, and demonstrate motion parallax as the camera is moved.

Note that the front-focused image cannot be focused as sharply as the rear-focused image due to optical aberrations in the angle-expanding screen. We characterize the point spread function (PSF) of the angle-expanding screen by displaying a point on SLM 2 and taking a RAW photograph with subtracted blacklevel. The recorded

PSF is approximated as a 2D Gaussian and incorporated into the light field factorization as described in more detail in the supplement. We also characterize the intensity transfer function of the SLMs, which are not well approximated by a standard Gamma curve. For this purpose, RAW photos of the screen are taken from the center of the viewing zone while the prototype displays a sequence of different intensities over the input range of the SLM driver. The inverses of the resulting curves are applied to the factorizations computed by the solver.

Software Implementation Target light fields are rendered using POV-Ray, but any graphics engine could be used alternatively. We implement the nonnegative light field factorization routines (Eq. 10) on the GPU using OpenGL and Cg. Pseudo-code is listed in the supplemental material along with source code for a Matlab-based implementation. Decomposing a light field with eight horizontal views and an image resolution of 1280×720 pixels takes about one minute on an Intel Core i7-2600 PC with an Nvidia GeForce GTX 690 GPU. Including convolution operations with the point spread function modeling screen aberrations increases processing times by a factor of $10 - 20\times$, depending on the PSF size. The finite blacklevel of each SLM is taken into consideration by clamping the values of g and h to the feasible range after each iteration (see Eq. 10). More implementations details are discussed in the supplement.

6 Assessment

The prototype was configured to achieve 5° field of view for horizontal-only parallax. Light fields with eight views are decomposed into rank-6 patterns with a brightness scale $\beta/T = 0.2$. Color channels are captured separately and digitally combined. This setting achieves a brightness benefit of $1.6\times$ over time-sequential parallax barriers and also provides compression in the number of required patterns (six vs eight).

Figure 8 shows an overview of four different scenes we captured. We show the central light field views for each scene as predicted in simulation and compare them to photographs of the prototype. Other than dark vertical stripes, the simulations predict observed image quality well. The stripes are due to the borders between screen lenticulars, which would completely disappear for lenticulars with a higher pitch or, similarly, for larger images projected onto the same screen.

We also show two extreme views for one of the scenes in Figure 9. These images are arranged so that they can be viewed using the “cross-eye” 3D viewing method¹. Parallax over the target field of view is well-reproduced. Animations of all results illustrating the observed 3D effect are included in the supplemental video. As described in the previous section, we incorporate an approximation of the optical aberrations of the screen in the solver. This results in decompositions that are blurrier for one layer than the other (Fig. 9, bottom).

7 Additional Applications

In this section, we briefly outline two additional application of the proposed *light field projector* when used with *conventional diffusing screens* as opposed to angle-expanding screens. In particular, we show that superresolution and high dynamic range projection can be achieved with such system configurations. Note that these are equally applicable to rear and front-projection systems.

¹Try crossing your eyes until the stereo image pair fuses. It does not work for everyone!




Figure 8: Overview of experimental results. Each column shows the central view of a light field that comprises eight views with horizontal-only parallax. Simulated results (top row) are compared with photographs of the prototype light field projector (bottom row). Color results are composited from three photos of our grayscale prototype.

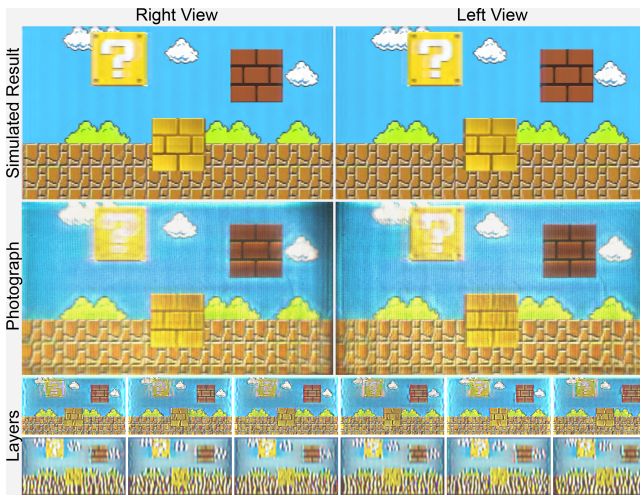


Figure 9: Two extreme views of an example scene with eight light field views for simulation (top row) and photographed prototype (center row). Try crossing your eyes to see it in 3D! The six-frame layer decompositions show how the solver accounts for optical aberrations of the screen prototype – by only considering blurry patterns on one of the SLMs (bottom row). Color results are composited from three photos of our grayscale prototype.

7.1 Superresolution Projection

Increasing display resolution has been one of the main driving factor of the industry for the last years. This is a challenging problem for projectors in particular, because not only do the electronics in the spatial light modulators have to decrease in size but also the projection optics must be designed to support the desired image resolutions. These requirements make projection optics for high-resolution display system, such as 4K and 8K projectors, expensive, complex, and large. In the following, we present an approach to super-resolved compressive image synthesis that uses two low-resolution SLMs in combination with a projection lens that only needs to support the low SLM resolution.

Following Equation 3, we can model the image formation of the proposed light field projector illuminating a conventional, diffuse screen as

$$\tilde{i}(x) = g(\phi_{2D}(x))h(\psi_{2D}(x)), \quad (11)$$

where $\tilde{i}(x) = \int \tilde{l}(x, \nu) d\nu$ is the observed image and $\phi_{2D}, \psi_{2D} : \mathbb{R} \rightarrow \mathbb{R}$ map the image coordinates of \tilde{i} to 2D coordinates on the SLMs. A single pixel of g or h is likely to affect multiple pixels in \tilde{i} , which is due to defocus blur and the fact that we wish to synthesize a projected image with a higher pixel count than that of either SLM. Similar to Equation 8, we can discretize this formulation, account for time-multiplexing over T frames, and incorporate the mapping functions into projection matrices as

$$\tilde{\mathbf{i}} = \frac{1}{T} \sum_{t=1}^T (\Phi_{2D} \mathbf{g}_t) \circ (\Psi_{2D} \mathbf{h}_t). \quad (12)$$

In practice, the projection matrices $\Phi_{2D} \in \mathbb{R}^{S \times N}$ $\Psi_{2D} \in \mathbb{R}^{S \times M}$ encode a convolution of the SLM images $\mathbf{g}_t \in \mathbb{R}^N$ and $\mathbf{h}_t \in \mathbb{R}^M$ at time t with Gaussian-type kernels modeling defocus blur and super-resolved image formation. The optical setup and image formation in both primal and frequency domains are illustrated in Figure 10.

A target image $\mathbf{i} \in \mathbb{R}^S, S > M, N$ can be approximated by the proposed setup by solving an objective function that closely follows the formulations in Equation 10 but replaces the 4D light field projection matrices with the 2D convolution matrices Φ_{2D}, Ψ_{2D} . In practice, all forward and transpose matrix multiplications are implemented as convolution operations in Matlab with point spread functions measured on the physical display. We include additional details and source code in the supplemental material.

Figure 11 shows results that we captured with our prototype. Two target images (top row) with a resolution exceeding that of either SLM are decomposed into three time-multiplexed frames that are displayed on the SLMs at a high speed. For a human viewer, or a camera capturing photographs with exposure times similar to those of the human visual system, the observed image (third row) has a higher resolution than that of a conventional projection with a single SLM (second row).

As illustrated in Figure 10, the improvement in resolution cannot exceed a factor of $2\times$. The highest frequencies displayed on either SLM optically multiply, which is a convolution of Dirac peaks in the Fourier transform. For two SLMs, this convolution shifts the maximum achievable frequency out to twice its original distance. Let us assume the maximum image frequencies of the panels, as observed on the screen, are f_g and f_h , respectively. Displaying sinusoidal patterns at the maximum frequencies results in the fol-




Figure 10: Illustration of superresolution display with light field projector and conventional, diffuse screen. The maximum image frequencies of each SLM (center) in the projector (left) are optically multiplied, which corresponds to a convolution in the frequency domain (right). Although the image of one of the SLMs is out of focus on the screen due to defocus blur, the effective resolution of the display is increased by a factor of $2\times$ compared to a conventional projection system (vertical dotted lines on the right).

lowing image:

$$\begin{aligned} \tilde{i}(x) &= \cos(f_g x) \cos(f_h x) \\ &= 1/2 (\cos((f_g - f_h)x) + \cos((f_g + f_h)x)). \end{aligned} \quad (13)$$

Although we ignore positivity constraints of the displayed image in this formulation, it does not change the fact that the maximum observable frequency in the composed image is $f_{max} = f_g + f_h$, allowing for super-resolution by a factor of $2\times$.

7.2 High Dynamic Range Projection

Inspired by high dynamic range (HDR) display systems [Seetzen et al. 2004], we also show application of the proposed projector to increased image contrast. For this purpose, we apply the same update rules as for superresolution to a target image \tilde{i} that has a higher dynamic range than either SLM. In our implementation, we apply the outlined update rules and clamp the resulting SLM patterns to the physically-achievable contrast after each iteration. Figure 11 (third row) shows that the dynamic range of a projected image—in addition to image resolution—can be significantly increased compared to a conventional projection system (second row).

A variety of additional results, animations, and layer decompositions for superresolution and HDR projection can be found in the supplemental material. For extensive numerical evaluations of all system parameters of dual and multilayer display devices, we refer the interested reader to [Heide et al. 2014].

8 Discussion

In this paper, we introduce a compressive light field projection system. Through the co-design of active, factored light field projection and a novel passive screen design, we present the first single device approach to glasses-free 3D projection that does not require mechanical movement of screen elements. We believe that the proposed system has promise to scale to large sizes, such as movie theaters, although additional engineering efforts are necessary to achieve the required image quality and dimensions.

What makes this a compressive light field projection system?

Most existing approaches to light field display use multiplexing. Each ray in the displayed light field maps to exactly one display element. Multiplexing in space, angle, or time allows each ray to be




Figure 11: Superresolution and high dynamic range projection. The proposed light field projector can be used with a conventional, diffuse screen to increase image resolution and contrast (third row) as compared to conventional projection with a single spatial light modulator (second row). Color results are composited from three photos of our grayscale prototype. The above source images are used with permission from the Peach Open Movie Project ©copyright 2008, Blender Foundation / www.bigbuckbunny.org.

addressed by some optical element in isolation. This has different implications for different solutions to 3D display: spatial multiplexing trades image resolution for angular resolution, time multiplexing trades SLM refresh rate for angular resolution, and multi-device systems use multiple projectors to create angular resolution. Unfortunately, low image resolutions significantly decrease the user experience, super-high-speed SLMs providing color and grayscale imagery are currently not available, and multiple devices make display systems expensive, bulky, and power hungry.

The proposed system addresses these challenges with compression. First, this is optically achieved by the screen’s expanding the angular range of the light field from the projector to an observer. (Or, equivalently, the screen compresses the observed light field into

the projector lens.) Second, it is well-known that natural images and videos are highly compressible. This notion directly extends to light fields—why build optics and electronics that address each light ray in isolation? Compressive display architectures present a compressed form of the target light field to an observer—limitations in visual integration speed allow the visual system to act as a de-compressor. The factorization algorithms discussed in this paper provide the means to compute the optimal compressed light field representation for a given optical system. The proposed projection system is the first to provide a compression w.r.t. the number of devices in an application space that is normally addressed (for similar fields of views) with multiple projectors. This is a direct benefit of the proposed screen.

Another benefit of our screen, for instance compared with using a simple Fresnel lens as a screen [Dodgson et al. 2000], is its optical flexibility. Although we only demonstrate angle amplification by a factor of 3, this number can be increased by altering the ratio of focal lengths within the screen. A Fresnel lens would fix the distance and size of the viewing zone for a fixed distance between projector and screen: a ratio governed by the lens-makers equation. The optical design space of the proposed screen w.r.t these parameters is more flexible in that the viewing zone size and location can be scaled independently of the projector-to-screen distance. An angle expanding screen fabricated from lenses with identical pitch and having a power $M = f_s^{(p)}/f_s^{(o)}$, will optically image an object at distance d behind the screen to a distance d/M in front of the screen for small values of d . We use a Fresnel lens in the prototype to focus the projector illumination to a viewing zone at an arbitrary distance to the screen. In future setups, the pitch between screen lenticulars can be modified such that the screen directly focuses light to the desired viewing zone while amplifying the angular range.

Alternative light field projectors Although we implement a light field projector using a duallayer design, conceptually our system can employ any light field projector. We built and tested both the described setup and also a coded aperture system [Grosse et al. 2010]. For the latter, we disassembled a Canon EF 50mm f/1.8 lens and mounted a polarizing beam splitter cube and an LCoS display as close to the aperture plane as possible. By extending the optical path lengths between these lens components, aberrations are introduced that severely limit observed image quality and that also introduce cross-talk between light field views. The duallayer setup described in this paper was significantly easier to build, align, and calibrate than the coded aperture device. Nevertheless, both configurations, and possible others as well, are well-suited for light field projection with the proposed screen. We report additional simulations and results for a home-built coded aperture systems in the supplemental document.

Possible extensions to multi-device systems Although we explore single device configurations throughout this paper, the proposed light field projection techniques can be extended to multi-device setups. For instance, multiple light field projectors with abutting projection lenses could be mounted in a horizontal array to extend the horizontal field of view of the system. Similar multi-device configurations were described by Dodgson et al. [2000] for a CRT-based, time-sequential projection system. A 2D array of projectors could increase the field of view both horizontally and vertically for a full parallax projection system. A challenge for such setups would be to arrange the devices close enough. However, periodic viewing zones of the lenslet-based screen could be exploited to place the projection devices at a larger physical distance but precisely located within different angular periods of the screen.

One could also consider the combination of a lenticular-based

horizontal-only parallax screen with vertically stacked projectors. The required vertical diffuser in the screen would optically combine (i.e. sum) their contribution, effectively increasing image brightness and also the achievable rank of the synthesized light field. The vertical summation caused by the diffuser is mathematically equivalent to increasing the speed of the SLMs in a single device. Multiple vertical projectors in such a setup would therefore be able to increase the rank of the system by a factor corresponding to the number of devices. In addition to increased rank, color and luminance projection could be decoupled, allowing for high-rank luminance and lower-rank chrominance light field display. Note that the human visual system is much more sensitive to depth cues in luminance than in chrominance. A projection system could be devised that presents image content optimized for the human visual system.

Although we have not yet built multi-device compressive light field projection systems, we simulate the benefits of a few configurations in the supplemental material and derive image formation and corresponding update rules for the factorization problem there as well.

Scaling the system A system of the size depicted in the concept sketch to the left of Figure 1 (2m screen size) can reasonably be obtained by scaling the optical properties of the current projector components. The field of view (FOV) of a projection system following the schematic shown in Figure 3 is given by

$$\text{FOV} = 2M \arctan \left(\frac{h_i}{2N(h_s - h_i)} \right), \quad (14)$$

where h_i and h_s are the SLM and screen heights, respectively, N is the effective f-number of the projection optics, and M is the power of the angle expanding screen. Note that the dependence of FOV on f-number alone, rather than the focal length or exit pupil size of the projection lens, remains so long as the projector-to-screen distance is not fixed.

Plausible real-world values for the variables in Equation 14 can be obtained from commercial catalogs and published academic work. The catalogs of Pacur and Micro Lens Technology Inc. contain commodity microlenses that range in focal length from 6.35mm to 0.26mm (although they differ in pitch), suggesting a plausible value for M approaching 25. A conservative estimate derived from analysis of similar screen optics in Eichenlaub et al. [2005] is $M = 10$. Commodity 35mm camera lenses with f-number as small as f/1.1 exist, such as the Voigtlander Nokton 50mm f/1.1. Therefore, from Equation 14 it follows that with an $N = 1.1$ lens, $M = 10$ power screen, $h_s = 2000\text{mm}$ and $h_i = 36\text{mm}$, equivalent to a typical working area for a 35mm lens, a projector placed 2.7m from the screen would produce a 2m wide light field image with a 10° FOV.

As previously discussed, wider fields of view can be obtained by employing multiple devices, such that a 20° FOV can be obtained with two devices, a 30° with three, and so on. Real-time color can be achieved using three devices of the type described, or one device with an SLM capable of switching at $3 \times 240\text{Hz} = 720\text{Hz}$. Achieving wider field of view with a single device will require smaller f-number projection lenses, or alternative screen optics, which we leave to future work.

Limitations The major limitation of the proposed system is the image quality achieved with the prototype setup. A maximum refresh rate of 240 Hz limits us to show rank-4–6 grayscale light fields for a human observing the prototype. Higher-speed SLMs with field sequential color or multi-device setups could address this limitation. The image quality of the prototype projector is limited by vignetting, optical field curvature by the beam splitter cubes, scatter-

ing in the screen, as well as color aberrations from the Fresnel lens. Further, the f-number of the projection system is currently limited to $f/1.8$ by the relay lens. The contrast of the SLMs is reduced by low f-number illumination, but this is inherently addressed and partially corrected for by the solver. The prototype screen provides an angular amplification factor of approx. $3\times$, resulting in a total field of view of approx. 5° achieved with the prototype. Future screen implementations should significantly increase this factor. Finally, the factorization adds additional computational cost to the system, but we are confident that real-time implementations are possible with optimized software on modern GPUs.

Future work Exploring multilayer or custom-designed coded aperture projection devices is an interesting avenue of future work. Most importantly, however, we would like to experiment with alternative spatial light modulators and evaluate alternative screen implementations. Higher-speed spatial light modulators, such as digital micromirror devices (DMDs) or fast transmissive microshutters (e.g., pixtronix.com) provide refresh rates in the KHz range as opposed to 240 Hz. Increasing the SLM speed would directly increase the observed 2D and 3D quality. In addition, we would like to scale the screen to a significantly larger size and also increase the angle-amplification factor. Fabrication techniques like embossing and interference lithography could facilitate the manufacture of low-cost and large-scale classical or holographic optical elements (HOEs) that achieve the same effect as our lenticular-based prototype screen. HOEs have precedent in projection screens; the transparent 2D projection screen designed by HoloPro employs HOEs. Our application may require volume holograms, which are widely used in large-scale color holograph (e.g. Zebra Imaging). Alternative screen designs, such as reflective screens, are also possible [Gabor 1944] and could be useful for devising front projection systems. The exploration of this space is outside the scope of this paper and proposed as future research. Improved image formation techniques, such as Adaptive Image Synthesis [Heide et al. 2013] can improve computational complexity for a range of compressive display problems, including light field projection.

9 Conclusion

Compressive light field displays are an emerging technology that offer significant improvements in resolution, brightness, contrast, and device form factors as compared to conventional 3D displays. With this work, we present the first compressive display approach to glasses-free 3D projection, which is facilitated via light field projection with a single device and a passive screen. In future implementations, we envision our technology to be seamlessly integrated into consumer products. Such projectors could be operated in different modes: as 2D projectors providing increased resolution and dynamic range as compared to conventional projectors or, in combination with the proposed screen, as 3D displays. With the demonstrated system flexibility and improvements over alternative projection technologies, the proposed system takes the first step toward affordable glasses-free 3D projection.

Acknowledgements

The authors thank Vincent Lee, who contributed to an earlier version of this work, and members of the Camera Culture group and Matthew O'Toole for helpful discussion. The authors express their gratitude to Phil Surman and Xin Li of Microsharp Innovation for fabricating the screen optics used in our prototype, and to Mike Jin and Paul Jin of Silicon Micro Display Inc., for providing access to and support for their SLM devices. This work was supported by Media Lab consortia funding. Gordon Wetzstein was supported by

an NSERC Postdoctoral Fellowship, and NSF Grant IIS-1116718 (AdaCID).

References

- AKELEY, K., WATT, S. J., GIRSHICK, A. R., AND BANKS, M. S. 2004. A Stereo Display Prototype with Multiple Focal Distances. *ACM Trans. Graph. (SIGGRAPH)* 23, 804–813.
- BALOGH, T. 2006. The HoloVizio System. In *Proc. SPIE 6055*, vol. 60550U.
- BERTHOUSOZ, F., AND FATTAL, R. 2012. Resolution Enhancement by Vibrating Displays. *ACM Trans. Graph.* 31, 2, 15:1–14.
- BOGAERT, L., MEURET, Y., ROELANDT, S., AVCI, A., SMET, H. D., AND THIENPONT, H. 2010. Single Projector Multiview Displays: Directional Illumination Compared to Beam Steering. In *Proc. SPIE 7524*, vol. 75241R.
- CICHOCKI, A., ZDUNEK, R., PHAN, A. H., AND ICHI AMARI, S. 2009. *Nonnegative Matrix and Tensor Factorizations*. Wiley.
- COSSAIRT, O., AND FAVALORA, G., 2006. Minimized-Thickness Angular Scanner of Electromagnetic Radiation, Apr. 26. US Patent App. 11/380,296.
- COSSAIRT, O. S., NAPOLI, J., HILL, S. L., DORVAL, R. K., AND FAVALORA, G. E. 2007. Occlusion-Capable Multiview Volumetric Three-Dimensional Display. *Applied Optics* 46, 8, 1244–1250.
- DAMERA-VENKATA, N., AND CHANG, N. L. 2009. Display Supersampling. *ACM Trans. Graph.* 28, 1, 9:1–9:19.
- DODGSON, N. A., MOORE, J. R., LANG, S. R., MARTIN, G., AND CANEPA, P. 2000. A time-sequential multi-projector autostereoscopic display. *Journal of the SID* 8, 2, 169–176.
- EICHENLAUB, J. B. 2005. Optical System Which Projects Small Volumetric Images to Very Large Size. In *Electronic Imaging 2005*, International Society for Optics and Photonics, 313–322.
- FUNK, W. 2012. History of Autostereoscopic Cinema. In *Proc. SPIE 8288*, vol. 82880R.
- GABOR, D., 1944. Optical System Composed of Lenticules, June 13. US Patent 2,351,034.
- GROSSE, M., WETZSTEIN, G., GRUNDHÖFER, A., AND BIMBER, O. 2010. Coded Aperture Projection. *ACM Trans. Graph.* 29, 22:1–22:12.
- HECHT, E. 2002. *Optics, fourth edition*. Addison Wesley.
- HEIDE, F., WETZSTEIN, G., RASKAR, R., AND HEIDRICH, W. 2013. Adaptive Image Synthesis for Compressive Displays. *ACM Trans. Graph. (Proc. SIGGRAPH)* 32, 4, 132:1–132:12.
- HEIDE, F., GREGSON, J., WETZSTEIN, G., RASKAR, R., AND HEIDRICH, W. 2014. A Compressive Multi-Mode Superresolution Display. *ArXiv e-prints* (Apr.).
- HEMBD-SÖLNER, C., STEVENS, R. F., AND HUTLEY, M. C. 1999. Imaging Properties of the Gabor Superlens. *Journal of Optics A: Pure and Applied Optics* 1, 1, 94.
- HONG, J., KIM, Y., PARK, S.-G., HONG, J.-H., MIN, S.-W., LEE, S.-D., AND LEE, B. 2010. 3D/2D Convertible Projection-type Integral Imaging using Concave Half Mirror Array. *Optics Express* 18.

- HSU, F.-H., 2008. Three-Dimensional (3D) Image Projection. US patent 7425070 B2.
- IVES, H., 1903. Parallax Stereogram and Process of Making Same. US patent 725,567.
- IVES, H. 1928. Camera for Making Parallax Panoramagrams. *J. Opt. Soc. Amer.* 17, 435–439.
- JACOBS, A., MATHER, J., WINLOW, R., MONTGOMERY, D., JONES, G., WILLIS, M., TILLIN, M., HILL, L., KHAZOVA, M., STEVENSON, H., AND BOURHILL, G. 2003. 2D/3D switchable displays. *Sharp Technical Journal*, 4, 15–18.
- JONES, A., MCDOWALL, I., YAMADA, H., BOLAS, M., AND DEBEVEC, P. 2007. Rendering for an Interactive 360° Light Field Display. *ACM Trans. Graph. (SIGGRAPH)* 26, 40:1–40:10.
- JONES, A., LIU, J., BUSCH, J., DEBEVEC, P., BOLAS, M., AND YU, X., 2013. An Autostereoscopic Projector Array Optimized for 3D Facial Display. *SIGGRAPH Emerging Technologies*.
- JURIK, J., JONES, A., BOLAS, M., AND DEBEVEC, P. 2011. Prototyping a Light Field Display Involving Direct Observation of a Video Projector Array. In *Proc. ProCams*, IEEE.
- KIM, Y., HONG, K., YEOM, J., HONG, J., JUNG, J.-H., LEE, Y. W., PARK, J.-H., AND LEE, B. 2012. A Frontal Projection-type Three-dimensional Display. *Optics Express* 20.
- KIMURA, H., UCHIYAMA, T., AND YOSHIKAWA, H. 2006. Laser Produced 3D Display in the Air. In *SIGGRAPH Emerging Technologies*, ACM, 20.
- LANMAN, D., HIRSCH, M., KIM, Y., AND RASKAR, R. 2010. Content-Adaptive Parallax Barriers: Optimizing Dual-Layer 3D Displays using Low-Rank Light Field Factorization. *ACM Trans. Graph. (SIGGRAPH Asia)* 28, 5, 1–10.
- LANMAN, D., WETZSTEIN, G., HIRSCH, M., HEIDRICH, W., AND RASKAR, R. 2011. Polarization Fields: Dynamic Light Field Display Using Multi-Layer LCDs. *ACM Trans. Graph. (SIGGRAPH Asia)* 30, 6, 186.
- LEE, D. D., AND SEUNG, S. 1999. Learning the Parts of Objects by Non-negative Matrix Factorization. *Nature* 401, 788–791.
- LIPPMANN, G. 1908. La Photographie Intégrale. *Academie des Sciences* 146, 446–451.
- MAIMONE, A., WETZSTEIN, G., LANMAN, D., HIRSCH, M., RASKAR, R., AND FUCHS, H. 2013. Focus 3D: Compressive Accommodation Display. *ACM Trans. Graph. (TOG)* 32, 5, 153:1–153:13.
- MASIA, B., WETZSTEIN, G., DIDYK, P., AND GUTIERREZ, D. 2013. A survey on computational displays: Pushing the boundaries of optics, computation, and perception. *Computers & Graphics* 37, 8, 1012–1038.
- MATUSIK, W., AND PFISTER, H. 2004. 3D TV: a Scalable System for Real-time Acquisition, Transmission, and Autostereoscopic Display of Dynamic Scenes. *ACM Trans. on Graph. (SIGGRAPH)* 23, 814–824.
- MEURET, Y., BOGAERT, L., ROELANDT, S., VANDERHEIJDEN, J., AVCI, A., SMET, H. D., AND THIENPONT, H. 2010. LED Projection Architectures for Stereoscopic and Multiview 3D Displays. In *Proc. SPIE 7690*, vol. 769007.
- NIMS, J., AND LO, A., 1972. 3-D Screen and System. US patent 3,814,513.
- PERLIN, K., PAXIA, S., AND KOLLIN, J. S. 2000. An Autostereoscopic Display. In *ACM SIGGRAPH*, ACM, 319–326.
- SAJADI, B., GOPI, M., AND MAJUMDER, A. 2012. Edge-Guided Resolution Enhancement in Projectors via Optical Pixel Sharing. *ACM Trans. Graph.* 31, 4, 79.
- SAJADI, B., LAI, D.-Q., IHER, A., GOPI, M., AND MAJUMDER, A. 2013. Image Enhancement in Projectors Via Optical Pixel Shift and Overlay. In *Proc. IEEE ICCP*, 1–8.
- SANDIN, D. J., MARGOLIS, T., GE, J., GIRADO, J., PETERKA, T., AND DEFANTI, T. A. 2005. The Varrier Autostereoscopic Virtual Reality Display. *ACM Trans. Graph. (SIGGRAPH)* 24, 3, 894–903.
- SEETZEN, H., HEIDRICH, W., STUERZLINGER, W., WARD, G., WHITEHEAD, L., TRENTACOSTE, M., GHOSH, A., AND VOROZCOVS, A. 2004. High Dynamic Range Display Systems. *ACM Trans. Graph. (SIGGRAPH)* 23, 3, 760–768.
- SMOOT, L. S., SMITHWICK, Q., AND REETZ, D. 2011. A Volumetric Display Based On A Rim-Driven Varifocal Beamsplitter And LED Backlit LCD. In *SIGGRAPH Emerging Technologies*, ACM, 22.
- SULLIVAN, A. 2003. A Solid-State Multi-Planar Volumetric Display. In *SID Digest*, vol. 32, 207–211.
- TOMPKIN, J., HEINZLE, S., KAUTZ, J., AND MATUSIK, W. 2013. Content-Adaptive Lenticular Prints. *ACM Trans. Graph. (TOG)* 32, 4, 133.
- TRAVIS, A. R. L. 1990. Autostereoscopic 3-D Display. *OSA Appl. Opt.* 29, 29, 4341–4342.
- UREY, H., CHELLAPPAN, K. V., ERDEN, E., AND SURMAN, P. 2011. State of the Art in Stereoscopic and Autostereoscopic Displays. *Proc. IEEE* 99, 4, 540–555.
- WETZSTEIN, G., LANMAN, D., HEIDRICH, W., AND RASKAR, R. 2011. Layered 3D: Tomographic Image Synthesis for Attenuation-based Light Field and High Dynamic Range Displays. *ACM Trans. Graph. (SIGGRAPH)*.
- WETZSTEIN, G., LANMAN, D., HIRSCH, M., AND RASKAR, R. 2012. Tensor Displays: Compressive Light Field Synthesis using Multilayer Displays with Directional Backlighting. *ACM Trans. Graph. (SIGGRAPH)* 31, 1–11.
- YANG, R., HUANG, X., LI, S., AND JAYNES, C. 2008. Toward the Light Field Display: Autostereoscopic Rendering via a Cluster of Projectors. *IEEE TVCG* 14, 1, 84–96.