Pinlight Displays: Wide Field of View Augmented Reality Eyeglasses using Defocused Point Light Sources

Andrew Maimone[1] Douglas Lanman[2] Kishore Rathinavel[1] Kurtis Keller[1] David Luebke[2] Henry Fuchs[1]
[1]University of North Carolina at Chapel Hill [2]NVIDIA Research

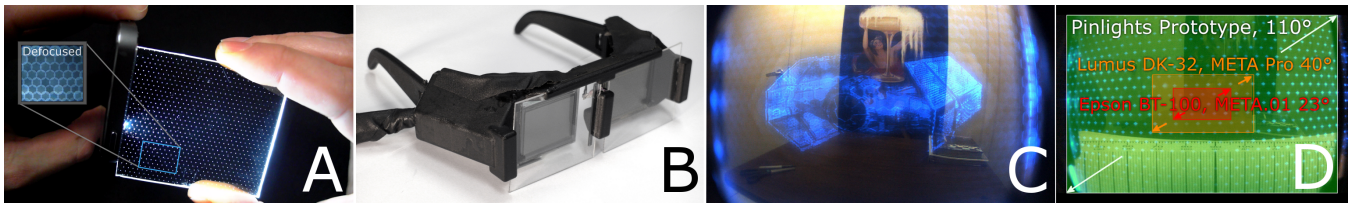

Figure 1: *A) A sparse array of point light sources, called* pinlights*, fill the eye's image plane when held* defocused *near the eye. B) Prototype optical see-through display consisting of pinlight arrays and spatial light modulators (LCDs). The spatial light modulators code the defocused pinlights to form an image on the retina. C) A photograph taken through our prototype display using a camera that approximates the human eye. D) A comparison of the field of view of our prototype display* (110°) *to state-of-the art commercial optical see-through glasses. Ship model by Staffan Norling.*

Abstract

We present a novel design for an optical see-through augmented reality display that offers a wide field of view and supports a compact form factor approaching ordinary eyeglasses. Instead of conventional optics, our design uses only two simple hardware components: an LCD panel and an array of point light sources (implemented as an edge-lit, etched acrylic sheet) placed directly in front of the eye, out of focus. We code the point light sources through the LCD to form miniature see-through projectors. A virtual aperture encoded on the LCD allows the projectors to be tiled, creating an arbitrarily wide field of view. Software rearranges the target augmented image into tiled sub-images sent to the display, which appear as the correct image when observed out of the viewer's accommodation range. We evaluate the design space of tiled point light projectors with an emphasis on increasing spatial resolution through the use of eye tracking. We demonstrate feasibility through software simulations and a real-time prototype display that offers a 110° diagonal field of view in the form factor of large glasses and discuss remaining challenges to constructing a practical display.

CR Categories: B.4.2 [Input/Output and Data Communications]: Input/Output Devices—[Image display] H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—[Artificial, augmented, and virtual realities]

Keywords: head-mounted displays, see-through displays, computational displays, augmented reality

Links: DL PDF VIDEO

ACM Reference Format
Maimone, A., Lanman, D., Rathinavel, K., Keller, K., Luebke, D., Fuchs, H. 2014. Pinlight Displays: Wide Field of View Augmented Reality Eyeglasses using Defocused Point Light Sources. ACM Trans. Graph. 33, 4, Article 89 (July 2014), 11 pages. DOI = 10.1145/2601097.2601141 http://doi.acm.org/10.1145/2601097.2601141.

0730-0301/14/07-ART89 $15.00.
DOI: http://dx.doi.org/10.1145/2601097.2601141

1 Introduction

Augmented reality (AR) offers a tantalizing vision for the future. Imagine leaving home to proceed along directions placed neatly on the sidewalk; along the way, a glance across the street yields the menu for a cafe, prompting us to stop and share coffee with a remote friend apparently seated across the table. In this example, we imagine *casually* harnessing graphics with meaningful *spatial connections* to the world, at a moment's notice and at many moments throughout the day. We imagine computer graphics transitioning from a distinctly external entity into a part of human vision.

Realizing this dream requires advances in many disciplines – low-latency rendering, tracking, application development, mobile computing, localization, networking – but perhaps the most fundamental problem is obtaining a suitable display. A display that satisfies the long-term potential envisioned for AR must satisfy two key requirements:

Wide Field of View: Field of view (FOV) is a critical attribute of a *spatially registered* AR display. A synthetic object or information overlay registered to the world, however small, may over time appear anywhere in a viewer's field of view as the viewer moves. Most envisioned AR applications also expect that a synthetic overlay could, at any given moment, fill an arbitrarily large portion of the viewer's FOV. Therefore, if the field of view of an AR display is less than the viewer's total field of view, registered objects and overlays will be cropped or will disappear and reappear with head motion. This reduces the effectiveness of the display: the user now must take an active role to discover and keep synthetic overlays in view, may not receive complete information at any instant, and may have a reduced sense that overlays represent synthetic objects that are present in the world. Although the field of view of the human eye is nearly 180°, the field of view achieved through the corrective lenses of ordinary eyeglasses – which generally span a $\geq 100^\circ$ horizontal FOV – suggests a pragmatic target.

Non-Encumbering: A display intended for *casual and extended use* must be ready to use in an instant, must be comfortable, and must not interfere with other tasks when not being used. This requires a non-encumbering design: it must not be too bulky, and must not significantly degrade normal human sight. As with field of view, ordinary eyeglasses or sunglasses demonstrably achieve an acceptable level of encumbrance and provide a target standard. Although

some bulk is generally accepted for research prototypes, it is important to consider the *minimum* size supported by a given design, which often has a hard lower limit due to optical constraints.

Recent developments in *optical see-through near-eye displays*, which superimpose synthetic imagery over a viewer's natural sight, tackle these two key requirements in isolation. Such devices have been demonstrated with a wide FOV (e.g. Cheng et al. [2011]) and in non-encumbering forms (e.g. Google Glass[1], Lumus DK-32[2]). However, to date no practical design has demonstrated both a wide field of view *and* low encumbrance in a single AR device.

The crux of the problem is that these requirements typically place opposing constraints on the optical hardware design. For example, optical see-through devices that place a beamsplitter embedded in a waveguide in front of the eye (e.g. Epson Moverio BT-100[3]) have a field of view that increases with the thickness of the waveguide; in the case of the Epson device, the field of view is limited to 23° diagonally while keeping the device acceptably thin and light.

In contrast, we present a novel optical see-through near-eye display design that provides a wide field of view and supports a compact and lightweight form factor that approaches ordinary eyeglasses. We replace conventional optical approaches, such as waveguides and beamsplitters, with a design that combines simple hardware and lightweight computation. We avoid the need for any optical refractive, reflective, or diffractive components that could limit field of view, and use only two simple and readily manufactured hardware components: an LCD panel and a sparse array of small, bright point light sources, formed on a patterned edge-lit acrylic sheet, that we call *pinlights*. Our core innovation is the use of *defocused* point light sources coded through a transmissive spatial light modulator (SLM) to form miniature, see-through, and therefore *imperceptible* projectors. These miniature projectors direct light into the lens of the eye through a *virtual aperture*, allowing their small image areas to be tiled to create an arbitrarily wide field of view. Software decomposes the target image into a series of tiled sub-images (displayed on the SLM) that each correspond to a miniature projector with a virtual aperture. We demonstrate the feasibility of our approach through a real-time prototype display in the form factor of large glasses that offers a 110° diagonal field of view. We evaluate the prototype display using a camera as a proxy for a human eye and discuss practical challenges to make the device suitable for human viewers.

1.1 Contributions

We present a novel approach to see-through near-eye displays. Specific contributions include:

- the use of point light sources, coded with an LCD and placed *near the eye*, that act as *imperceptible, see-through* projectors
- the use of such projectors in a *tiled* configuration encoded with a *virtual aperture* to expand the FOV of the display
- the use of such tiled projection arrays in an alternative configuration to provide a *light field* over the eye, as a *see-through* alternative to existing near-eye light field displays
- an example *hardware design* for creating transparent tiled projection arrays, described and evaluated in a prototype device

1.2 Benefits and Limitations

The proposed design offers several benefits over existing see-through near-eye displays. A wide FOV is supported with no theoretical upper bound (or approaching 180° if the display components are planar), and a prototype achieving a 110° diagonal FOV is demonstrated. The design also supports a lightweight, eyeglasses-like form factor and uses simple, low cost hardware components.

The proposed design also exhibits several limitations. Precise pupil tracking, while not strictly required, strongly benefits the possible image quality and resolution. (Physical eye tracking was not implemented in this work.) Diffraction limits image quality and provides a challenge to scale beyond modest image resolutions. The brightness of the environment is also dimmed through the display when practical and low-cost LCD spatial light modulators are used, providing the effect of wearing sunglasses. Some lightweight computation ($\approx$1-2 ms on a modern GPU) must be performed to generate an image for the display, and accurate geometric and radiometric calibration of the display is needed to eliminate tiled periodic structures from appearing in the perceived image. See Section 5.3 for additional discussion of current limitations and possible solutions.

2 Related Work

Near-Eye See-Through Displays

Near-eye displays, particularly optical see-through systems, have experienced a groundswell of consumer enthusiasm over the last year, initiated in part by the joint introduction of Google Glass (narrow field of view AR) and Oculus Rift (low-cost yet immersive VR). However, the optical designs underpinning the majority of these commercially-announced devices have been maturing over decades of research. Kress and Starner [2013] provide a concise survey of the state-of-the-art in near-eye displays. We briefly survey these systems, studying their benefits and limitations relative to our new pinlight projection system.

Freeform optics innovate on the traditional beamsplitter approach to achieve optical see-through, leveraging the complex, aspheric surfaces afforded by modern optical fabrication methods [Cakmakci et al. 2010; Cheng et al. 2011]. Rather than employing a planar combiner and separate relay lens elements, a single optical "prism" is manufactured to unify the functions of relaying, magnifying, and combining the virtual image. Aberrations of the virtual and physical environments are corrected by jointly optimizing the optical surfaces.

Freeform optics replace the bulk of head-mounted relay optics with precision-manufactured prisms; however, their FOV remains limited, with the volume of the prism growing proportionally to the target FOV. Waveguide-based designs can overcome this limitation: temple-mounted projectors are used to shine light into a flat waveguide, where it travels by total internal reflection (TIR) to an *out-coupling element* that redirects the image to the viewer's eye. Waveguide approaches are differentiated by the composition of the out-coupling (and associated in-coupling) elements. Systems such as the Lumus DK-32 and the Optinvent ORA-S[4] utilize a set of *cascaded extractors* comprising an array of precisely-aligned semi-reflective mirrors. Alternatively, diffractive optical elements have been successfully employed to achieve out-coupling, including grating patterns and more general volume holograms [Levola 2006]. The FOV of such systems is restricted by the TIR angle supported by the waveguide, practically limiting systems at the moment to less than roughly 60°. We observe that, as with our design, near-eye diffractive elements (e.g., our LCD panels) can also introduce distortion of the see-through environment.

Our use of a direct-view LCD panel, with no additional relay, magnification, or combining optics, is shared by only a handful of near-

[1] http://www.google.com/glass/

[2] http://www.lumus-optical.com/

[3] http://www.epson.com/cgi-bin/Store/jsp/Moverio/Home.do

[4] http://optinvent.com/see-through-glasses-ORA

eye displays. The recently-announced Innovega iOptik display system[5] uses a direct-view display and a custom contact lens. The contact lens employs polarization-selective filters and an embedded microlens to transform the viewer's accommodation range: the central portion of the pupil becomes capable of focusing on the near-eye display (through the microlens), whereas the outer annulus retains the viewer's natural accommodation range. In a closely-related work, Maimone and Fuchs [2013] propose a near-eye display composed of compact stacks of $2-3$ direct-view LCD panels. This multi-layered display presents a series of time-multiplexed attenuation patterns – at greater than the human flicker fusion threshold – creating a near-eye light field display. Similar to our design, eye tracking significantly enhances resolution. We emphasize that the pinlight approach differentiates itself by eliminating the need for contact lenses or complex, multi-layered LCD panels; yet, achieving such hybrid designs will necessitate overcoming the diffraction we observe with direct-view LCDs.

Coded Projections using Defocused Optical Elements

Our development of near-eye coded pinlight projectors is part of a larger body of on-going work revealing diverse applications for defocused projection systems. In the computer graphics community, Mohan et al. [2009] present one of the earliest such designs; their "Bokode" system uses a combination of a backlit, microprinted transparency covered with a microlens to project an image at optical infinity from an aperture of a few millimeters. A coded image is formed when observed with a defocused wide-aperture camera, allowing extremely compact substitutes for traditional 2D barcodes. Hiura et al. [2010] extend this approach to curved arrays of Bokodes, which they dub "Krill-eye beacons" due to the optical similarity with refracting superposition compound eyes in decapods. In contrast to our system, these beacons rely on refractive lenses and are intended to be distantly located from any imager.

Pamplona et al. [2010; 2011] demonstrate near-eye applications of these principles to interactively assess the refractive errors of human subjects. In their system, a near-eye light field display is created by covering a microdisplay with either a microlens array or a pinhole grid to present test patterns containing accommodation cues. Recently, Lanman and Luebke [2013] show that such systems can be optimized for general display applications, demonstrating lightweight, thin-form-factor virtual reality eyeglasses. We emphasize that, while sharing conceptual similarities, neither approach directly facilitates optical see-through applications: refractive microlens arrays irrevocably distort the viewer's perspective.

We do, however, find similarities with a specialized class of light field displays: those exploiting *parallax illumination*. In a related work, Son et al. [2007] substitute a 2D point light source (PLS) array for the microlenses and pinholes (or parallax barriers) routinely found within light field displays. As with our pinlight projectors, multiview imagery is projected by displaying a coded array of elemental images on an LCD placed in front of the PLS array. We also find similarities in our approach to projector-based displays. Jurik et al. [2011] describe a light field display in which each projector in a dense array acts as a single display "pixel" with a high degree of angular variation. We emphasize that, unlike pinlight displays, PLS arrays and projection arrays are designed to function within the accommodation range of the observer (e.g. as a desktop 3D display). Furthermore, we are the first to demonstrate architectures for making these arrays imperceptible in near-eye use. Instead of the LED array of Son et al. [2007], or the pico projectors of Jurik et al. [2011], we use an edge-lit acrylic sheet with a pattern of microdivots.

[5] http://www.innovega-inc.com/new-architecture.php

3 System Overview

3.1 Introduction

Optical see-through near-eye displays place synthetic imagery over one's natural vision using a device placed closer than the eyes can focus. Such displays have two seemingly simple goals: (1) to adjust for the focal range of the eye, while (2) preserving the ability to see the real world through the display. Displays using conventional optical approaches address these functions with additional components, such as lenses and beamsplitters. These components tend to add complexity and bulk to the design and limit the field of view.

We take an optically simpler approach that avoids these components by designing a *highly directional* image source that emits light over a very narrow angle and specific orientation for each pixel. Thus, each image pixel acts as a "ray" source (within an approximation), rather than a point source. With this design, the light emitted from each *individual* pixel of the image source is essentially non-divergent, and can enter the eye directly without corrective optics and regardless of eye focal state. We achieve such an image source by coding the light from a sparse array of point sources with a spatial light modulator. We must, however, account for this different image formation method by preprocessing the image. To preserve the see-through capability, our design uses only transparent components, a see-through point light array and a transmissive SLM, placed directly in front of the eye. This avoids the need for any cumbersome or FOV-limiting components to combine the augmented image and see-through optical paths. In this way, we follow the approach of Maimone and Fuchs [2013], who stacked transmissive SLMs for a see-through near-eye display. We add novelty by achieving transparency using components with small *visible* features ($\leq 100\ \mu m$) that are *imperceptible when defocused*.

3.2 Coded Projection with a Single Pinlight

Image Formation Our core approach is directly encoding light from a *defocused* point source placed near the eye (outside of the accommodation range), which we call a *pinlight*, with a transmissive SLM placed between the eye and the pinlight. This system acts as a miniature projector, which we call a *pinlight projector*, that directs light into the eye as illustrated in Figure 2. Assuming for simplicity that the pinlight and the pixels on the SLM are true mathematical points, each location on the SLM receives light from a single direction (a "ray") that is modulated and then enters the eye. The set of all rays is refracted by the lens of the eye; however, since the rays originate from a single point their angular ordering is preserved, creating a "sharp" copy of the modulated image on the retina. Since the point source is nearer than the viewer's minimum accommodation distance, the image is not flipped by the eye, and thus the modulated image must be inverted along both axes. The image formed on the retina is "sharp" regardless of the focal state of the eye (much as with a camera with an infinitely small pinhole aperture); the eye's focal state changes only the degree of refraction and therefore the scaling of the image formed. Note that the illumination is simply a coded point source; therefore it also captures the point spread function of eye, which has some non-uniformity in shape and intensity. Also note that **the image produced by a pinlight projector is generally round, due to the shape of a human pupil,** which adds complexity to optical design as described in Section 3.3. See Section 3.4.1 for details concerning the creation of real point sources with non-zero extent and handling changes in eye state. See-through ability is achieved through the use of transparent components; see Section 3.4.3 for details.

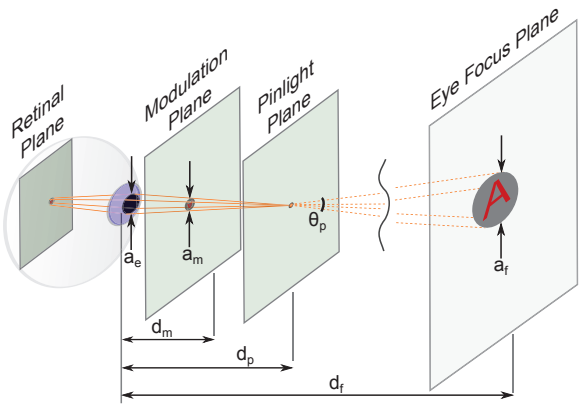

Figure 2: *Pinlight Projection. A defocused point light source, or* pinlight*, placed near the eye is coded with a spatial light modulator to create a narrow field of view image that appears in focus without the use of refractive or diffractive optics.*

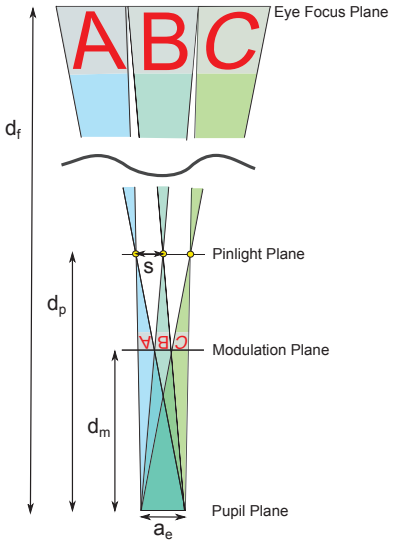

Figure 3: *Ideal tiled pinlight projector geometry. The display is configured so that the pinlight projectors abut on the eye focus plane and on the modulation plane, creating a continuous image and using the SLM efficiently. (Note that the pinlights emit light over a wide angle, but only light entering the pupil is drawn).*

Projection Geometry From Figure 2, it can been seen that the necessary diameter of modulation on the SLM a_m and the diameter of the image on the focus plane a_f for a *single* pinlight projector can be computed as:

$$a_m = a_e\left(1 - \frac{d_m}{d_p}\right),\ a_f = a_e\left(\frac{d_f}{d_p} - 1\right), \tag{1}$$

given eye aperture diameter a_e, pinlight plane distance d_p, modulation plane distance d_m, and eye focus distance d_f. Likewise, the field of view θ_p through the single pinlight can be computed as:

$$\theta_p = 2\tan^{-1}\left(\frac{a_e}{2d_p}\right). \tag{2}$$

The pinlight plane distance d_p and modulation plane distance d_m are the variables under the control of the display designer. From Equation 2, we observe then that the FOV through a single pinlight projector θ_p is increased by decreasing pinlight plane distance d_p. Given a selection of d_p, modulation plane distance d_m can be chosen to select the desired modulation scale on the SLM. However, for eye pupil diameter $a_e = 3$ mm, the pinlight plane set at a practical distance of $d_p = 25$ mm yields a FOV through the pinlight projector θ_p of only $6.9°$. Obtaining a wide field of view of $100°$ would require an impractically close pinlight plane distance of $d_p = 1.26$ mm. However, we explore the use of multiple pinlight projectors to increase the field of view, as described in the following section.

3.3 Coded Projections using Tiled Pinlight Arrays

A single pinlight projector does not alone provide a useful field of view for an augmented reality display. We observe, however, that **multiple pinlight projectors may be *tiled* to significantly increase the field of view**. In this section, we describe various tiling configurations.

Ideal Tiling Geometry When tiling pinlight projectors, we assumed that the pinlight plane now contains an *array* of point light sources, all of which are modulated by a single SLM. Further, we assume that the pinlights emit light over a wide angle so that each is visible to the eye (but we are unconcerned with light rays that do not enter the eye). In this configuration, the pinlights can be tiled to an arbitrarily wide field of view (or approaching $180°$ if the pinlight array and SLM are restricted to planes) where the total display area is approximately proportional to the number of projectors used, subject to practical limitations on the emission angle of the pinlights and the size and supported viewing angles of the SLM. The tiling must satisfy two primary conditions: the eye focus plane must contain a *continuous* tiled image among all the projectors, and the modulated areas on the SLM must be *disjoint* among the projectors. We also aim to maximize resolution by using as much area on the SLM as possible. We first consider an *ideal* one-dimensional case, illustrated in Figure 3. Given an eye aperture a_e, eye focus distance d_f, and pinlight plane distance d_p, the optimal pinlight spacing s and modulation plane distance d_m are computed as:

$$s = a_e\left(1 - \frac{d_p}{d_f}\right) \quad \text{and} \quad d_m = \frac{a_e d_p}{a_e + s}. \tag{3}$$

This spacing ensures that neighboring view cones abut at the modulation plane and focus planes, providing a continuous image at the full resolution of the SLM. Note that the pinlights are placed sparsely (with spacing on the order of the eye aperture size), so that the structure of an *array* of small pinlights will remain imperceptible when defocused. Also note that the display becomes thinner as it is moved nearer the eye (i.e., decreasing d_p decreases $d_p - d_m$).

The effective resolution r in terms of the preserved fraction of modulation plane resolution (by area) can be computed as:

$$r = \left(\frac{d_f(d_p - d_m)}{d_m(d_f - d_p)}\right)^2 \tag{4}$$

This equation provides the ratio of the width of a pixel projected onto the focus plane to the total width of the focus plane, squared to provide a ratio by area. Note that this equation assumes that the pinlight geometry is valid; the eye focus plane must contain a continuous tiled image, and the modulated areas on the SLM must be disjoint among projectors. In the ideal configuration, ratio r equals 1: the entire modulation plane has been used to form an

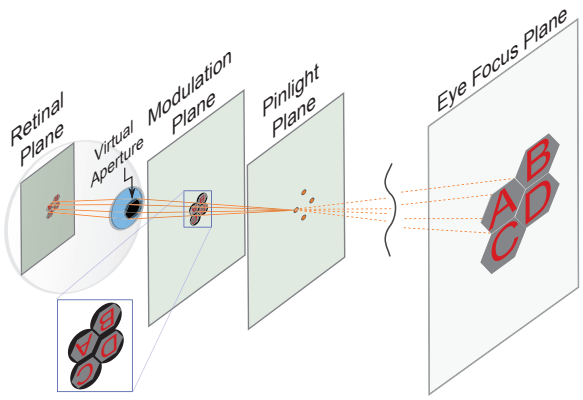

Figure 4: *Tiled pinlights in tracked virtual aperture configuration. An aperture mask is encoded over the desired image on the modulation plane to create a virtual hexagonal aperture over the eye that can be tiled, eliminating toning artifacts in the perceived image. The image on the modulation plane is recomputed based on the tracked eye position to allow eye movement relative to the display. (Note that the pinlights emit light over a wide angle, but only light that enters the pupil is drawn).*

image on the focus plane without redundancy. The horizontal field of view f_h from a point on the eye can also be computed as:

$$f_h = 2\tan^{-1}\left(\frac{c}{2d_m}\right) \tag{5}$$

where c is the width of the modulation plane.

Challenges for Practical Tiling Tiling has the potential to create a wide field of view display. However, in the *ideal* 1D case described above we have not considered several factors which must be addressed to create a practical *human-wearable* display:

- The ideal case can only be directly extended to 2D if the image areas created by the pinlight projectors can be tiled; however, a circular image area is created when projecting into the round human pupil, which cannot be tiled without the inclusion of gaps or overlapping areas.
- The ideal model assumes that the eye position is fixed relative to the display, an invalid assumption for a viewer wearing glasses. If the eye moves the image will circularly shift *within* the tiled sub-images corresponding to each pinlight projector, resulting in a corrupted image.
- A device using the ideal model is only modestly thin. For example, a device designed for eye aperture $a_e = 3$ mm focused at $d_f = \infty$ with the pinlight plane placed at $d_p = 30$ mm yields a device $d_m = 15$ mm away from the eye that is $d_p - d_m = 15$ mm thick. This is thicker than typical eyeglasses, although we note most of the thickness consists of empty space so that the device may remain lightweight.
- The ideal model is affected by changes in pupil size and focal state, which are expected to change over time.

In the remainder of this section, we address items 1-3 under two alternative configurations. Item 4 is addressed in Section 3.4.5.

3.3.1 Tracked Virtual Aperture Configuration

Eye Tracking In one alternative tiled configuration, we allow eye movement relative to the display by assuming that the position of the eye on the pupil plane is known (i.e. tracked) and that the image on the modulation plane is adjusted to account for the new eye position. As the eye moves, the view cones that corresponding to each pinlight projector (illustrated in Figure 3) shift and intersect with a new portion of the intended image on the eye focus plane, which is then flipped and scaled appropriately to be encoded on the corresponding region of the modulation plane (see Section 3.4.4 for details). In this section, we assume theoretical error-free tracking without latency in sensing or display; in Section 3.3.2 we describe how to account for these factors. We did not implement physical eye tracking in this paper (instead relying on a camera at a known location for testing), but discuss the possibility in Section 5.3.

Virtual Eye Apertures Although eye tracking allows compensation for eye movements, it does not resolve the issue of how to tile the circular images formed by individual pinlight projectors due to the round aperture of the eye. In particular, if the circular projectors are tiled without overlap, gaps will remain in the image, and if the circles overlap to fill the plane, the overlapping areas will have greater intensity than the non-overlapping areas (see Figure 6E). **Our solution to create an evenly-toned image is to configure the pinlight projectors so that they *minimally* overlap to fill the focus plane and to encode a *virtual aperture* over the modulation plane so that that light from the overlapping regions does not reach the eye**. This process is illustrated in Figure 4. The virtual aperture has the effect of transforming the viewer's pupil into a shape that can be tiled (e.g. a rectangle or hexagon), as if wearing a contact lens masked with such a shape.

Geometry The ideal display geometry, given by Equation 3, is updated to support a *hexagonal* virtual eye aperture as follows.

$$s_{t_h} = \frac{\sqrt{3}}{2}a_e\left(1 - \frac{d_p}{d_f}\right),\ s_{t_v} = \frac{\sqrt{3}}{2}s_{t_h} \tag{6}$$

$$d_{m_t} = \frac{\frac{1}{2}\left(a_e + \frac{\sqrt{3}}{2}a_e\right)d_p}{\frac{1}{2}\left(a_e + \frac{\sqrt{3}}{2}a_e\right) + s_{t_h}} \tag{7}$$

This geometry assumes that the virtual aperture is encoded as a hexagon with vertical left and right sides and a pointed top, as seen in Figure 10B. Note that the horizontal pinlight spacing s_{t_h} and vertical spacing s_{t_v} are asymmetric and that odd pinlight projector rows should be offset by $\frac{s_{t_h}}{2}$ due to the staggered hexagonal packing of the plane. Also note the similarity to the geometry in the ideal case (Equation 3), except that pinlight spacing s has been decreased to allow the circular image areas of the pinlight projectors to overlap to cover the focus plane, and modulation plane distance d_m has been adjusted to allow space for the virtual aperture mask. In particular, a regular hexagon is inscribed into the circular area on the modulation plane that represents each pinlight projector, and d_{m_t} is set so that the projectors are positioned as closely as possible without intersecting the inscribed hexagons of neighboring projectors. This process causes a resolution loss as some of the modulation plane is now dedicated to providing an aperture mask rather than contributing to the virtual image. The resolution loss can be computed according to Equation 4. Equivalently, the effective resolution (by area) in a tracked virtual aperture configuration can be computed as the ratio of the area of the inscribed hexagon in a unit diameter circle $\left(\frac{3\sqrt{3}}{8}\right)$ with the area of the larger hexagon

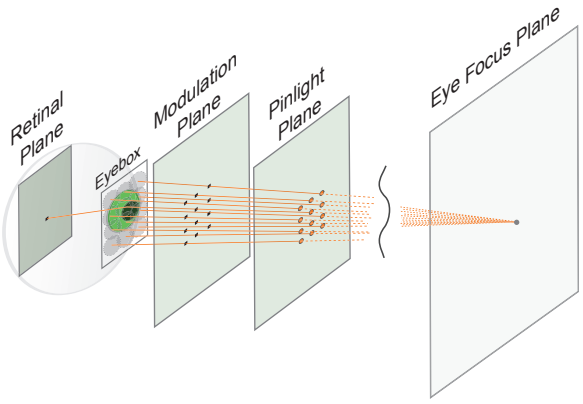

Figure 5: *Untracked light field configuration. For each point on the eye focus plane, a set of rays are created through a tiled set of pinlight projectors that are distributed about an eyebox near the eye, allowing eye movement. The modulated pixels corresponding to the same point in the scene are distributed over the image.*

that would exactly tile the plane without the virtual aperture mask $\left(\frac{7\sqrt{3}+12}{32}\right)$, yielding effective resolution ratio r_t:

$$r_t = \frac{12\sqrt{3}}{12+7\sqrt{3}} \approx 86\%. \tag{8}$$

However, spatial resolution loss is very modest compared to existing near-eye light field displays (e.g. [Maimone and Fuchs 2013]) showing the benefit of a design that allocates nearly of all the display area to forming the image perceived by the user, through tracking. Also note that moving the modulation plane closer to the pinlight plane to accommodate the virtual aperture has the positive side effect of creating a slightly thinner display; e.g., an aperture $a_e = 3$ mm focused at $d_f = \infty$ with the pinlight plane placed at $d_p = 29$ mm now yields a device $d_{m_t} = 15$ mm away from the eye that is $d_p - d_{m_t} = 14$ mm thick (i.e. 1 mm thinner).

3.3.2 Untracked Light Field Configuration

Near-Eye Light Fields In another alternative tiled projector configuration, the display is designed to generate a *light field* near the eye to allow additional capabilities. The display is designed so that the view cones between pinlight projectors *overlap* away from the display, allowing angular variation in the image among projectors. In particular, one can design the display to allow angular variation *around* the eye to create an *untracked* configuration, which we explore in this section. It is also theoretically possible to design the display to emit a light field with angular variation *over* the eye (creating depth of field), but due to the very high display panel resolution requirements, we leave this as a topic for future work.

Untracked Configuration The tracked display configuration offers high spatial resolution, but the need for pupil tracking adds engineering complexity. An alternative display can be defined with an *eyebox*, a region in which the eye can move around while still seeing the complete intended image. Here, the display is configured to emit multiple light rays that appear to originate from each point on the eye focus plane, each of which is directed towards different regions of the eyebox, as illustrated in Figure 5. To maximize resolution, the display geometry minimizes the number of modulated rays emitted from each point on the eye focus plane with assurance that an eye placed anywhere in the eyebox will receive one such ray; Equations 6 and 7 take the following form in this configuration:

$$s_{u_h} = \frac{\sqrt{3}}{2} a_{e_m} \left(1 - \frac{d_p}{d_f}\right),\ s_{u_v} = \frac{\sqrt{3}}{2} s_{u_h},\ d_{m_u} = \frac{a_{e_b} d_p}{a_{e_b} + s_{u_h}} \tag{9}$$

Horizontal pinlight spacing s_{u_h} and vertical spacing s_{u_v} are defined with respect to a minimum eye aperture a_{e_m}; a viewer with a smaller aperture will perceive gaps in the image. Modulation plane distance d_{m_u} is defined with respect to a constrained window around the eye a_{e_b}, creating an eyebox of size $a_{e_b} - a_e$. Unlike a conventional display, each ray in the eyebox can be modulated individually, allowing different images to be perceived with eye movement and without tracking. The display is considerably thinner in this configuration; an eye with minimum aperture $a_{e_m} = 3$ mm and eyebox of $a_{e_b} - a_e = 7$ mm (when $a_e = a_{e_m}$) focused at $d_f = \infty$ with the pinlight plane at $d_p = 18.9$ mm yields a device $d_{m_u} = 15$ mm from the eye that is $d_p - d_{m_u} = 3.9$ mm thick.

Limitations The untracked display configuration described above offers two notable challenges, however. First, the displayed image contains unevenly toned areas where the eye receives light from multiple pinlight projectors. This is a consequence of distributing rays in the eyebox with assurance that the *round* human pupil receives light from every point in the scene; however, the inability to exactly tile the eyebox plane with circles causes overlapping regions. This effect cannot be directly compensated for without knowledge of the eye position; however, approaches for mitigating the effect are discussed in Section 5.3. Another limitation of the untracked configuration is a significant loss in spatial resolution. In the example above, the preserved SLM resolution computed using Equation 4 is approximately 6%, which is comparable to other near eye light field displays [Pamplona et al. 2010; Pamplona et al. 2011; Lanman and Luebke 2013]. Resolution loss may be minimized, however, by creating a tracked display with a small eyebox to compensate for tracker error and latency.

3.4 Practical Implementation Details

In this section, we address issues of realizing a display based on the tiled pinlight projector method that is practical for human viewers.

3.4.1 Creating Point Light Sources

Requirements To create point light sources for pinlight projectors, we note three primary requirements. First, the pinlight sources should be very bright, i.e. the total light emitted should be on par with a normal display panel, but concentrated into the much smaller area of a sparse grid of dots. Second, the emission area of each pinlight should be very small (i.e. on the order of the size of a pixel on the SLM) as the image formed is essentially convolved with the pinlight shape. To maximize resolution, the effective aperture size of the SLM pixels (with consideration for the pixel fill factor) plus the size of the pinlight should be less than or equal to the pixel pitch of the SLM. Finally, the pinlight array should be highly transparent and any of the structures composing the pinlight array should be small enough as to be imperceptible when defocused.

Implementation Our prototype implementation is a waveguide (an acrylic sheet) that is etched with tiny divots and edge illuminated with LEDs, as illustrated in Figure 6. Light from the LEDs is channeled through the waveguide to the opposite side except where it encounters the divots, which scatter light and cause bright spots to appear. Note that a similar configuration is commonly found in the backlights of LCD panels. The divots were etched using a sewing needle attached to the moving robotic platform of a 3D

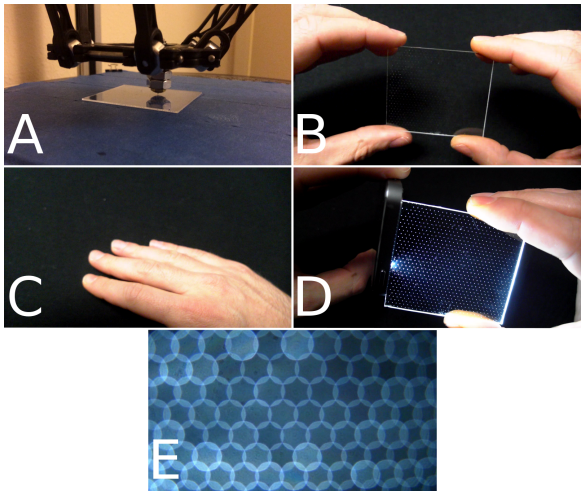

Figure 6: *Waveguide-based pinlight array. A) Array manufactured by creating small divots in an acrylic sheet using a needle attached to the robotic arm of a 3D printer. B) The waveguide contains a sparse array of visible divots. C) When viewed with a camera with an aperture and focus similar to a human eye, the divots are imperceptible and see-through. D) When illuminated from the side with an LED, bright spots appear in the waveguide. E) When defocused near the camera, the pinlights form discs that tile the image plane.*

printer, which was programmed to create a staggered pinlight array according to Equation 6. The etched features are visible, but the pinlight array appears completely clear when held near the eye. This method produced cone shaped divots with a maximum diameter of ≈ 100 μm, potentially limiting resolution; a future version will use more precise laser drilling. Also note that this implementation is light inefficient since much of the light escapes through the waveguide edges; however a bright array can be achieved with a modest number of LEDs (six per eye in our prototype display).

Note other possibilities for creating an array of transparent light sources: transparent emissive displays (e.g. transparent OLEDs), LED or laser chips mounted on a transparent conductive substrate (e.g. ITO coated glass), fiber optics, holograms, and quantum dots.

3.4.2 Modulating Light Sources

An SLM intended for a pinlight projector display must work in a transmissive mode and should have high transmittance, a wide viewing angle, high contrast, and should minimally diffract light. In our implementation, we selected an LCD microdisplay intended for use in a projector due to its relatively high transmissivity and pixel density. To minimize the diffraction and light loss through the display caused by color filters, a monochrome panel was selected and operated in conjunction with a color sequential pinlight array. Hardware implementation details are described in Section 4.1.

3.4.3 Optimizing See-Through Capability

To achieve a see-through capability, it is assumed that the SLM and pinlight array are effectively transparent when defocused near the eye. Note, however, two complicating factors. First, light from the environment may permeate the SLM in addition to illumination from the pinlights, causing a soft defocused glow around the synthetic image (see Figure 10C). Second, light from the environment only reaches the eye through the defocused mask on the SLM, causing a soft, uneven coloring of the environment and allowing little light to reach the eye in areas where there are no synthetic objects (see Figure 10D). To mitigate these issues, we rapidly alternate between displaying an augmented image with the pinlights on (see Figure 10D), and displaying a occlusion mask of the augmented image with the pinlights off (which appears defocused, see Figure 11D). This allows light from the environment to reach parts of the display where no augmented imagery is shown and reduces the apparent soft glow around the augmented images (see Figure 10E).

3.4.4 Creating Modulation Masks

To create the modulation masks to display on the SLM, the virtual scene is projected through the pinlights onto the modulation layer with respect to the eye. We implement this process in software by rendering the virtual scene with a camera that uses an off-axis projection frustum defined by the eye and each pinlight; this process is repeated for each pinlight. If the scene is a simple 2D plane at the eye focus distance (rather than an arbitrary light field), this process is performed more efficiently by finding the intersection of the camera frustum and the focus plane and transferring a flipped and appropriately scaled copy of the image region onto the correct region of the modulation plane. For images representing occlusion masks, we simply draw an unmodified copy of the image on the modulation plane, which appears defocused when viewed. Example modulation and occlusion masks are shown in Figure 11.

3.4.5 Changes in Eye State

Although accommodating the display for eye movement was discussed in Section 3.3, the eye may change in other ways that may affect the display; notably the pupil size or focal state may change.

Handling Change in Pupil Size In the tracked configuration (Section 3.3.1), we assumed that the eye aperture size was fixed in Equation 6. We can account for a dynamic aperture size using one of four methods. First, the display geometry can be configured for the maximum allowable aperture size and a *virtual aperture* (see Section 3.3.1) can be created with the minimum allowable size. This allows only a smaller central area of the pupil to receive light as the pupil size changes, but will cause a resolution loss proportional to the ratio of the largest and smallest apertures. Second, the display could be configured with a small eyebox sufficient to allow the pupil size to expand beyond a minimum size. This approach also results in a similar loss in resolution but the additional of a small eyebox also helps compensate for tracker error and latency. Third, the display could be outfitted with a *dynamic* pinlight array (e.g. a transparent OLED display) that can adjust the spacing of the pinlights according to tracked pupil sizes. Finally, the amount of light reaching the eyes could be controlled in an active loop by the SLM and/or pinlight array in an attempt to control pupil size to a predetermined size. The display could also simply be set sufficiently bright to force the pupils to contract to near their minimum size. However, note other factors than lighting may affect pupil size. In the untracked configuration (Section 3.3.2), variable aperture sizes are already considered in Equation 9.

Eye Focus In a tiled pinlight projector display, changes from the expected eye focus do not cause the image to become appreciably "blurred", but rather change the scaling of the tiled sub-images among the various projectors so that they expand or contract, causing tile gaps or overlaps to occur in some cases. This is expected to appear less natural to the user than normal focal blur. However, the change in scaling is small unless the viewer is focused at very close range; e.g. the change in scaling from a distance of 1 m to infinity is $\leq 3\%$ in a typical display (see Equation 6). Gaps caused

by nearer than expected focus can be avoided by configuring the pinlight projectors so that they slightly overlap.

3.4.6 Diffraction

Light passing through an aperture diffracts (expands angularly) with a magnitude inversely related to the aperture size. This effect is troublesome in the proposed design as the viewer sees the augmented image and real environment through a series of small apertures on the SLM, degrading both views.

We approximate the diffraction-limited resolution of our display by assuming that the image is viewed at infinity (i.e. Fraunhofer diffraction) and that two adjacent points are no longer resolvable when the diffraction maximum of one point coincides with the first minimum of the other (i.e. Rayleigh criterion). The angular limits of the display are then computed as:

$$\theta_d = \frac{\lambda}{w}, \quad \theta_g = \frac{w}{2(d_p - d_m)} \tag{10}$$

where θ_d is the angle of the first diffraction minimum, λ is the wavelength of the light, w is the SLM aperture size, θ_g is half the angular extent of a pixel defined by the geometry of the display, and $d_p - d_m$ is the distance between the pinlight plane and modulation plane. The maximum angular resolution is then calculated by equating these angles:

$$w_{opt} = \sqrt{2(d_p - d_m)\lambda}, \quad r = \frac{1}{\left(\frac{180}{\pi}\right)\left(\frac{\lambda}{w_{opt}}\right)} \tag{11}$$

where w_{opt} is the optimal aperture size (i.e. when θ_d and θ_g are equal) and r is the corresponding angular resolution in pixels/degree. Among the display geometries discussed in this paper and the wavelengths of visible light, our display design is capable of angular resolutions in the range of $2-5$ pixels/degree. These resolutions are very modest; however they are sufficient for a range of possible augmentation as shown in experimental results of Section 5.2. The possibility of further increasing resolution is discussed in Section 5.3. Diffraction of the see-through view of the environment can cause visible rainbowing effects, especially in bright and high contrast areas. See Figure 12 for examples of the environment seen through a prototype display device. We note that see-through diffraction is not a unique problem to our approach and must be mitigated in other see-through designs (e.g. waveguide designs using holographic out-couplings), providing motivation to mitigate this problem in future work.

4 Implementation

4.1 Hardware

To test our design experimentally, we created a prototype device that operates in the "tracked" configuration (see Section 3.3.1), but with the use of a camera in a known location rather than a human eye. Our prototype device consists of two main optical components: LCDs (Sony LCX017, 36.9 x 27.6 mm active area, 1024x768 pixels, 705 dpi, 60 Hz, monochrome) and a waveguide pinlight array that was constructed as described in Section 3.4.1 (1 mm thick, 1.8 mm horizontal pinlight pitch) with RGB LED color sequential illumination. The modulation plane and pinlight plane were spaced at a distance of $d_p - d_m = 13.5$ mm, creating an optical assembly with a total thickness of 15.5 mm including component thickness and 10.5 mm of empty space. The components were mounted in a 3D printed plastic eyeglasses frame designed in OpenSCAD and printed by Shapeways. See photos in Figure 1B and Figure 7. The bulky components on the far left and right of the glasses house an LCD adapter board that can be removed in a future device.

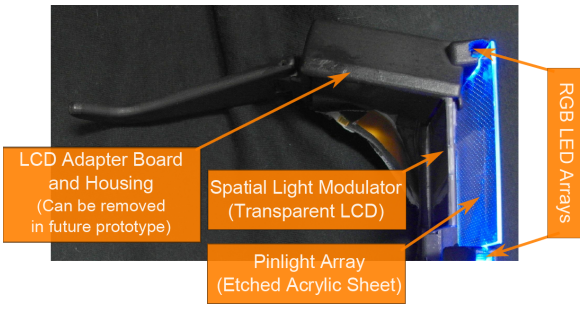

Figure 7: *Display components of one side of prototype display.*

The display was tested using a camera that approximated a human eye (PointGrey Flea2 camera body with a Fujinon F1.2, 2.3 mm aperture lens). The camera was mounted behind the display (see Figure 8) at a distance that approximated a human wearing eyeglasses (i.e. a camera center of projection to modulation plane distance of $d_m = 16$ mm). We did not implement eye tracking, but assumed the camera was in a known position behind the display. The field of view of the display through the camera is $\approx 110^\circ$ diagonally, limited by the camera's field of view and a cropping of the top of the camera image to remove portions of the LCD with poor viewing angles. The FOV was measured by placing a protractor at the camera's center of projection as shown in Figures 8 and 1D; this method was inspired by similar measurements taken in William Steptoe's AR Rift System[6].

An Arduino microcontroller board and transistor circuit were used to drive the RGB LEDs in a color sequential mode that was synchronized to the monochrome display. Since the utilized LCD panels have a very slow response time (greater than the 16.67 ms duration of a 60 Hz frame), we doubled each color sub-frame and illuminated the LEDs for a short period at the end of the second subframe to reduce switching blur. A occlusion mask sub-frame was also included to improve the see-through ability (see Section 3.4.3). With the doubled subframes, the effective framerate was 7.5 Hz. We also experimented with a bright monochrome color mode with three duplicated color subframes that operated at 12 Hz; the LEDs were turned on the entire third subframe after the LCDs had fully switched during the first two subframes. All photographs of the display were taken from real-time video capture at the rates listed above with the exposure time set to the entire frame time (to integrate over all sub-frames). The LCD panels were controlled by an externally housed driver board connected by a DVI link to an Intel Core i7 PC with an NVIDIA GTX 580 GPU running Linux.

4.2 Software

Real-time tiled imagery for the prototype display was generated with an OpenGL/GL Shader Language based program using the fast image transfer method described in Section 3.4.4. Computation times for generating the tiled images given the target image were $1-2$ ms per frame. The program also decomposed the image into color channels for color sequential operation and adjusted the image to compensate for poor LCD contrast at wide angles. Simulated images were generated by drawing the modulation plane image and a grid of pinlight dots at their apparent positions with

[6] http://willsteptoe.com/post/66968953089/ar-rift

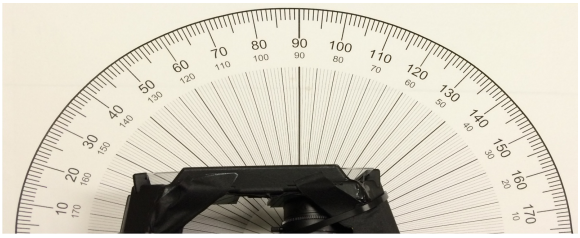

Figure 8: *Camera mounting and display measurement. A camera is mounted behind the display prototype that approximates a human viewer wearing eyeglasses. A protractor is placed with the origin at the camera's center of projection to measure field of view; an image taken through the camera is shown in Figure 1D.*

respect to a virtual camera placed at the eye position. Images were summed over a densely sampled set of positions over a round camera aperture that matched the diameter of the intended viewer.

5 Experimental Assessment

5.1 Simulated Results

To evaluate the theoretical performance of the proposed display design, we simulated various tiled pinlight projector configurations as described in Section 4.2 with a modulation plane that matched the specifications of our prototype LCD panel (see Section 4.1).

The following configurations were tested: (1) a tracked virtual aperture configuration of $d_p = 29$ mm and $d_m = 15$ mm, (2) an untracked, small eyebox configuration of $d_p = 22.8$ mm, $d_m = 15$ mm, and $a_{e_b} = 5$ mm, and (3) a untracked larger eyebox configuration of $d_p = 18.6$ mm, $d_m = 15$ mm, and $a_{e_b} = 11$ mm. All configurations were simulated with an eye pupil size of $a_e = 3$ mm focused at a distance of $d_f = \infty$. Note that the tracked configuration is similar to that used by our prototype display. The untracked, small eyebox configuration provides a miniature eyebox of $a_{e_b} - a_e = 2$ mm that could be used to compensate for tracker error and latency and pupil size variation.

Results are shown in Figure 9. Note that resolutions follow predictions of Sections 3.3.1 and 3.3.2 and that the untracked configurations exhibit uneven image toning. The tracked virtual aperture simulation may be compared to the result achieved on our prototype display in Figures 12C and 12D. Note that these results do not simulate diffraction, which will cause additional image degradation.

5.2 Prototype Display Results

To evaluate real-world performance, we also tested our tiled pinlight projector design using a hardware prototype with a camera placed behind the display. The hardware prototype was configured in the tracked virtual aperture configuration (see Section 3.3.1) with the assumption that the eye position was known using a camera placed in a fixed position. See Section 4.1 for full specifications.

Figure 10 shows the steps of image formation on our prototype. The image begins as overlapping defocused discs (Figure 10A) from the pinlight array which are converted to a plane of abutting hexagonal tiles through the encoding of a virtual aperture on the modulation plane (Figure 10B). An augmented image is then encoded in the modulation plane, which appears defocused when viewed with unstructured illumination from the scene (Figure 10C). When illuminated, the strongly directional light from the pinlight array causes a focused augmented image to appear, although the dark regions of

Figure 9: *Display simulations. Blue outlined inset images show magnified regions. Top Left: Reference image, used as target image for simulation. Top Right: Tracked configuration simulation. Bottom Left: Untracked configuration simulation (small eyebox). Bottom Right: Untracked configuration (larger eyebox).*

the modulation mask cause the background to appear dark, except for a glowing region around the virtual image (Figure 10D). When an occlusion mask sub-frame is included with the pinlight array off, the see-through ability is improved (Figure 10E).

Figure 11 shows the inputs to the prototype display and the captured outputs when the display is operated in a color sequential mode with an occlusion mask sub-frame. During each color sub-frame, an encoded image is sent to the display's LCD panel (Figure 11A), which appears as Figure 11B when viewed through the display. During each occlusion mask sub-frame, the backlight is disabled and an occlusion image is sent to the LCD (Figure 11C), which appears defocused through the display (Figure 11D). The final perceived image is the integration of the color and occlusion mask sub-frames (Figure 11F), which can be compared to the theoretical performance of the display through simulation (Figure 11E).

Figure 12 shows the display generating imagery for a variety of AR scenarios: a gesture based interaction (Figure 12A), visualization of a detailed model (Figure 12B), and display of overlaid text (Figure 12C). Note that Figures 12B and 12C were taken in the high brightness mono mode and that Figure 12C shows the display utilizing its entire field of view of approximately 110° diagonally. Figure 12D shows a magnified region of Figure 12C which can be compared to the similar simulated case in Figure 9 (top right).

5.3 Assessment and Discussion

We have demonstrated a real-time, compact, see-through AR display with a very wide field of view that can display imagery with a modest resolution sufficient to display medium sized text. However, we observe several practical issues to address in future work to create a human wearable display with optimal image quality:

Eye Tracking Eye tracking was not implemented in our prototype. Possible approaches include placing a camera at the edge of a tapered backlight that can see the eyes through total internal reflection [Travis et al. 2013], or placing a camera highly off-axis [Świrski et al. 2012]. A configuration with tracking and a small eyebox could be used to compensate for tracker error and latency.

Improving Resolution The size of the light scattering divots on our needle etched pinlight array backlight were $2\times - 3\times$ larger than

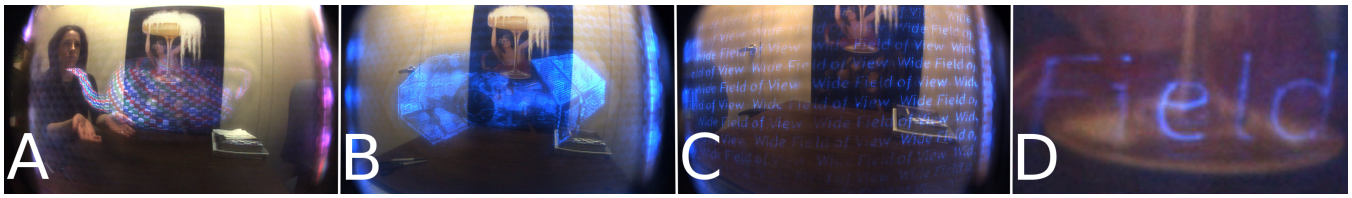

Figure 12: *Sample results from prototype display. A) User interacting with teapot. B) Model visualization. C) Text displayed to fill entire field of view. D) Magnified region of image C which represents a horizontal FOV of approximately* 12°*. Ship model by Staffan Norling.*

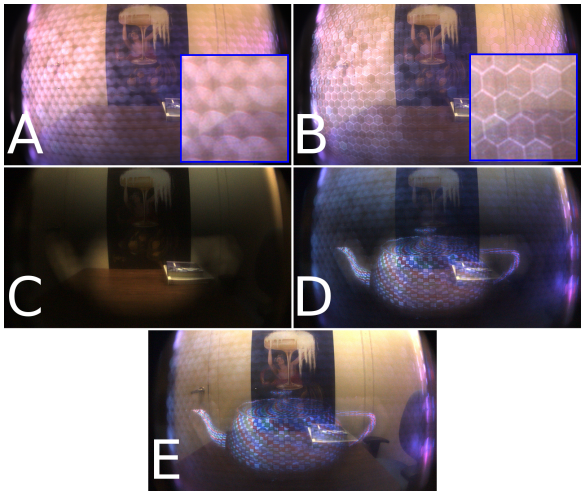

Figure 10: *Image formation process. Blue outlined inset images show magnified regions. A) Defocused pinlights form overlapping discs on image plane. B) A hexagonal virtual aperture is encoded in a modulation pattern. C) A modulation pattern is displayed without the pinlights illuminated. D) Pinlights are illuminated, causing the augmented image to appear. E) The image is displayed with a periodically displayed occlusion mask, improving see-through ability.*

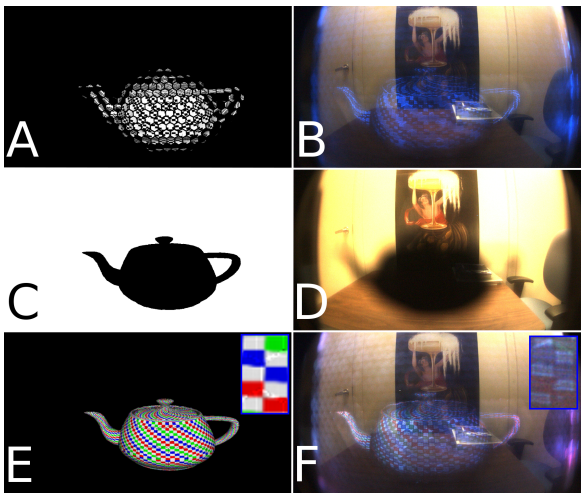

Figure 11: *Sample display inputs and results. Blue outlined inset images show magnified regions. A) Modulation pattern for augmented image sent to display for blue color channel. B) Photograph of display while blue color channel displayed. C) Modulation pattern sent to display as occlusion mask. D) Photograph of display while occlusion mask is shown. E) Simulated reconstruction of augmented image. F) Photograph of actual augmented image and background, taken over complete color and occlusion cycle.*

the pixel pitch of our LCD panel, potentially limiting resolution. Much smaller divots ($\geq 1\ \mu m$) can be created using a laser drilling process. The full resolution of an LCD microdisplay may not be achieved due to diffraction; however, we believe there is utility in a lower resolution, very wide FOV display for many spatially registered applications such as navigation and object identification.

Improving Contrast and Frame Rate Our display contrast depends on the contrast of the pinlight array (i.e. the ratio of the light emitted from the pinlights to that emitted or transmitted through the remaining pinlight layer area) and the contrast of the SLM (i.e. the ratio of transmitted light between the most transparent and opaque states). Our prototype suffers from low contrast; Figures 12A and 12B show that areas of the display without augmented imagery still show a visible pinlight projector pattern, especially near the corners. In the case of our etched waveguide, pinlight array contrast can be improved by taking care to keep dust and scratches off the array (which scatter light), to mask side light leakage, and to make the pinlights as large as possible while still preserving the maximum display resolution (see Sections 3.4.1 and 3.4.6). The LCD panels in our prototype, chosen in part for their availability of an off-the-shelf driving board, also suffer from low contrast, especially off axis, limiting the image quality and FOV. The high response time and low refresh rate of the panels has also made the prototype's color sequential mode too slow to be perceptually integrated by the eye. Similar LCD panels with higher contrast ratios (1000:1), higher frame rates (120 Hz) and faster response times are available from vendors such as Epson and Sony (e.g. Sony LCX086).

Improving See-Through Transparency The transmittance of our spatial light modulators is 23%; while efficient for an LCD, the see-through view is similar in transparency to medium sunglasses even when the panels are set to their most transparent state. It is expected that transparency could be improved with more light efficient SLM technologies, e.g. MEMS microshutters [Hagood et al. 2007].

Improving Color Quality The prototype display exhibited "washed-out" and non-uniform color. Color quality could be improved by adjusting the individual intensities of the RGB LEDs and by performing radiometric calibration over the display.

Mitigating Diffraction Diffraction appears to limit the resolution of our display to modest values below those of available SLMs and causes some degradation of the see-through view. We expect that this effect may be mitigated with LCDs designed to reduce diffraction (e.g. Chaing et al. [2005] and Benoît-Pasanau et al. [2010]). Major LCD manufactures (e.g. Samsung) are also now manufacturing LCDs designed for transparent applications, providing an industry motivation to reduce diffraction effects. A complementary approach is to attempt to predistort the image displayed on the SLM so that it more closely matches the intended image when diffracted.

Decreasing Image Toning Our prototype display produced a tiled effect, especially in large areas of uniform intensity, due to varying light intensity and slight misalignment between the pinlight projectors. Quality could be improved with better geometric calibration

and the inclusion of radiometric calibration and blending between neighboring projectors (Brown et al. [2005]). In the untracked configuration, image toning caused by overlapping projectors is a fundamental issue predicted by theory; however, the use of a dynamic pinlight array to allow the positions of the projectors to be swept or randomized over time may reduce perception of the effect.

Decreasing Aliasing Artifacts Some of the results exhibited aliasing artifacts, most evident on the ship model in Figure 12B and in the supplemental video. The effects could be reduced by using more careful sampling (e.g. anisotropic filtering) when projecting the scene from the eye focus plane to the modulation plane through the pinlights with respect to the eye.

Decreasing Device Thickness In the tracked configuration, the optical stack is typically thicker than ordinary glasses. However, most of the thickness consists of empty space; it is possible that this space could be collapsed by placing a small refractive or diffractive element over each pinlight, causing it to appear more distant. Such elements should be small relative to the pupil size to avoid appreciable distortion of the see-through view. The use of a tracked configuration with a small eyebox, a likely candidate for a practical device, would also result in a thinner device.

6 Conclusion

Spatially registered augmented reality has shown promise to enhance the capabilities of human vision: reading a sign far down the street, seeing a name to every forgotten face, never being lost again. We imagine graphics becoming a part of our eyes for a brief moment and fading unobtrusively away the next. For this dream to become a reality, we argue, will require an augmented reality display with the wide field of view and non-encumbering form factor of a pair of eyeglasses.

We describe a new optical see-through AR display that makes advances toward this goal. Our approach combines optical simplicity and software flexibility, enabling us to demonstrate an unprecedented field of view even in an early prototype display. The prototype has a competitively thin form factor and the approach should scale down to even thinner final systems. Some hurdles remain to achieve full practicality, but we continue to be inspired by the prospect of future AR displays indistinguishable from eyeglasses and just as indispensable to those who wear them.

Acknowledgements

We thank the reviewers for their valuable feedback. We also thank Jim Mahaney for engineering support and Don Miller for assistance with 3D printing. This work was supported in part by the National Science Foundation (award #1319567).

References

BENOÎT-PASANAU, C., GOUDAIL, F., CHAVEL, P., CANO, J.-P., AND BALLET, J. 2010. Minimization of diffraction peaks of spatial light modulators using Voronoi diagrams. *Opt. Express 18*, 14, 15223–15235.

BROWN, M., MAJUMDER, A., AND YANG, R. 2005. Camera-based calibration techniques for seamless multiprojector displays. *Visualization and Computer Graphics, IEEE Transactions on 11*, 2 (March), 193–206.

CAKMAKCI, O., THOMPSON, K., VALLEE, P., COTE, J., AND ROLLAND, J. P. 2010. Design of a free-form single-element head-worn display. *Proc. SPIE 7618, Emerging Liquid Crystal Technologies V*, 761803–761803–6.

CHAING, H.-C., HO, T.-Y., AND SHEU, C.-R., 2005. Structure for reducing the diffraction effect in periodic electrode arrangements and liquid crystal device including the same. US Patent. US 6977705.

CHENG, D., WANG, Y., HUA, H., AND SASIAN, J. 2011. Design of a wide-angle, lightweight head-mounted display using free-form optics tiling. *Opt. Lett. 36*, 11 (Jun), 2098–2100.

HAGOOD, N., BARTON, R., BROSNIHAN, T., FIJOL, J., GANDHI, J., HALFMAN, M., PAYNE, R., AND STEYN, J. L. 2007. 35.5l: Late-news paper: A direct-view mems display for mobile applications. *SID Symposium Digest of Technical Papers 38*, 1, 1278–1281.

HIURA, S., MOHAN, A., AND RASKAR, R. 2010. Krill-eye: Superposition compound eye for wide-angle imaging via grin lenses. *IPSJ Transactions on Computer Vision and Applications*, 186–199.

JURIK, J., JONES, A., BOLAS, M., AND DEBEVEC, P. 2011. Prototyping a light field display involving direct observation of a video projector array. In *IEEE International Workshop on Projector-Camera Systems (PROCAMS)*.

KRESS, B., AND STARNER, T. 2013. A review of head-mounted displays (HMD) technologies and applications for consumer electronics. In *Proc. SPIE*, vol. 8720.

LANMAN, D., AND LUEBKE, D. 2013. Near-eye light field displays. *ACM Trans. Graph. 32*, 6 (Nov.), 220:1–220:10.

LEVOLA, T. 2006. Diffractive optics for virtual displays. In *Journal of the Society for Information Display*.

MAIMONE, A., AND FUCHS, H. 2013. Computational augmented reality eyeglasses. In *Mixed and Augmented Reality (ISMAR), 2013 IEEE International Symposium on*, 29–38.

MOHAN, A., WOO, G., HIURA, S., SMITHWICK, Q., AND RASKAR, R. 2009. Bokode: Imperceptible visual tags for camera based interaction from a distance. In *ACM SIGGRAPH 2009 Papers*, ACM, New York, NY, USA, SIGGRAPH '09, 98:1–98:8.

PAMPLONA, V. F., MOHAN, A., OLIVEIRA, M. M., AND RASKAR, R. 2010. Netra: Interactive display for estimating refractive errors and focal range. In *ACM SIGGRAPH 2010 Papers*, ACM, New York, NY, USA, SIGGRAPH '10, 77:1–77:8.

PAMPLONA, V. F., PASSOS, E. B., ZIZKA, J., OLIVEIRA, M. M., LAWSON, E., CLUA, E., AND RASKAR, R. 2011. Catra: Interactive measuring and modeling of cataracts. In *ACM SIGGRAPH 2011 Papers*, ACM, New York, NY, USA, SIGGRAPH '11, 47:1–47:8.

SON, J.-Y., SAVELJEV, V. V., KIM, D.-S., KWON, Y.-M., AND KIM, S.-H. 2007. Three-dimensional imaging system based on a light-emitting diode array. *Optical Engineering 46*, 10, 103205–103205–4.

ŚWIRSKI, L., BULLING, A., AND DODGSON, N. A. 2012. Robust real-time pupil tracking in highly off-axis images. In *Proceedings of ETRA*.

TRAVIS, A., LARGE, T., EMERTON, N., AND BATHICHE, S. 2013. Wedge optics in flat panel displays. *Proceedings of the IEEE 101*, 1, 45–60.