# Depth-Fused 3D Imagery on an Immaterial Display

Cha Lee, *Student Member*, *IEEE*, Stephen DiVerdi, and Tobias Höllerer, *Member*, *IEEE*

**Abstract**—We present an immaterial display that uses a generalized form of depth-fused 3D (DFD) rendering to create unencumbered 3D visuals. To accomplish this result, we demonstrate a DFD display simulator that extends the established DFD principle by using screens in arbitrary configurations and from arbitrary viewpoints. The feasibility of the generalized DFD effect is established with a user study using the simulator. Based on these results, we developed a prototype display using one or two immaterial screens to create an unencumbered 3D visual that users can penetrate, examining the potential for direct walk-through and reach-through manipulation of the 3D scene. We evaluate the prototype system in formative and summative user studies and report the tolerance thresholds discovered for both tracking and projector errors.

**Index Terms**—Three-dimensional displays, immaterial displays, virtual reality, user studies, depth-fused 3D.

## 1 Introduction

As computational power and the interest in 3D graphics have increased dramatically in recent years, 3D display technology has become an active field of novel systems capable of creating real 3D images, where light is emitted from the actual 3D position within the viewing volume [1]. These displays create a realistic 3D perception because all depth cues are faithfully recreated, but so far every such display is limited to creating a visual in an enclosed volume the user cannot penetrate, hindering intuitive interaction. The ideal 3D display would create a 3D image without this limitation, enabling users to directly select and manipulate virtual 3D objects in a natural and intuitive manner, without the need for encumbering user-worn glasses. In this paper, we present a display system that takes a step toward attaining this ideal.

An interesting unencumbering pseudo-3D display technique is called depth-fused 3D (DFD) by Suyama et al. [2]. DFD perception occurs when two 2D images are displayed such that they are superimposed on two transparent screens with varying luminance and the observer perceives a 3D image. The image appears closer to the observer if the front screen is more luminous and farther away if the back screen is more luminous. In Suyama's original display, only a single view was possible but it could simulate a 3D scene with no eyewear, similar to autostereo displays [3]. Today, there are desktop-sized [2] and handheld-sized DFD displays [4]. We call these standard DFD displays, consisting of two or more screens stacked parallel to one another, and restricting the observer to a single viewpoint. We extend this principle to arbitrary viewpoints and screen configurations to create and evaluate a general DFD display.

The emergence of immaterial displays has created a great opportunity for direct interaction techniques. Immaterial displays are displays which allow the user to occupy the same space as the image. We have experimented with a large-scale immaterial display, the FogScreen [5], [6]. This screen is a $2.5 \times 1.5$-m projection surface, which consists of a thin, stable sheet of fog. The fog scatters rear-projected light to create an image that floats in thin air. Because of its immaterial composition, users can touch and even walk through the fog and, with adequate tracking, interact directly with the displayed virtual objects.

The contributions this paper presents are threefold. First, we simulate and evaluate a generalized DFD display; second, we use the generalized DFD technique in developing a prototype immaterial display using FogScreens (Fig. 1) and third, we evaluate aspects of the prototype in formative and summative evaluations and report error thresholds associated with projector and tracking devices. The purpose of our generalized DFD display is to demonstrate that multiple transparent screens, in arbitrary configurations and with arbitrary viewpoints, can still achieve the DFD effect, extending the current established DFD results. This is confirmed in a formal user study using the simulator. Using this result, the prototype display uses two FogScreens and an optical tracking system to create an immaterial DFD display. We tested our prototype in two configurations and discuss the results. Our results demonstrate that observers can indeed perceive 3D objects as having real depth with our system.

The rest of this paper is organized as follows: In Section 2, we survey established results pertaining to 3D display technologies. Section 3 describes the design of the simulator, while Section 4 details the user study that measured the simulator's performance. In Section 5, we describe the design of our display prototype. In Section 6, we discuss the general problems which are inherent to this type of display. In Sections 7 and 8, we evaluate and report the projector

• *The authors are with the Department of Computer Science, University of California, Santa Barbara, Santa Barbara, CA 93106-5510. E-mail: {chalee21, sdiverdi, holl,}@cs.ucsb.edu.*

*Manuscript received 3 Mar. 2008; revised 20 May 2008; accepted 27 May 2008; published online 5 June 2008.*
*Recommended for acceptance by W. Purgathofer.*
*For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org, and reference IEEECS Log Number TVCGSI-2008-03-0032.*
*Digital Object Identifier no. 10.1109/TVCG.2008.83.*

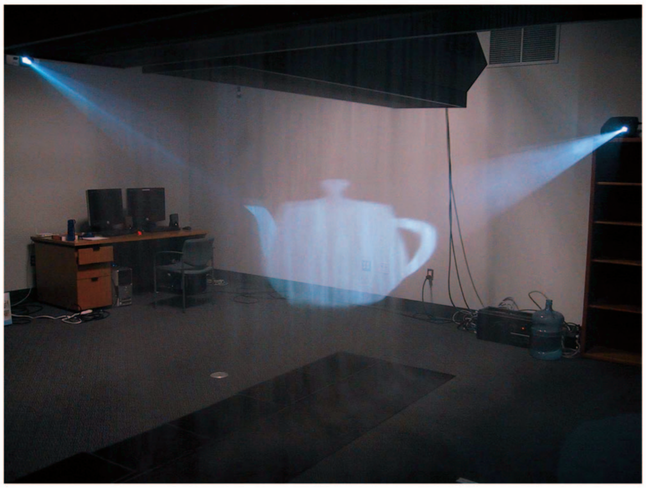

Fig. 1. Our prototype immaterial DFD display using two FogScreens in an L-shaped configuration, showing a 3D teapot.

and tracking thresholds we discovered in the course of an expert pilot study and a formal user study.

## 2 Related Work

Many different technologies have been pursued to create the perception of a 3D scene in an audience. The most appealing notion is to simply create points of light in a 3D volume, effectively scanning a 3D image one voxel at a time. Achieving this result has required some ingenuity. Favalora et al. [7] project light onto a rapidly spinning screen, carefully timing the projection to illuminate individual voxels. Alternately, Lightspace Technology's DepthCube [8] projects onto a stack of parallel LCD shutters. More exotic concepts such as Downing et al. [9] employ infrared lasers to excite points in a rare-earth doped gas, while the lasers in Kimura et al. display [10] create light-emitting plasma out of the air. Carefully controlled falling water droplets have even been used to scatter projector light in a 3D volume, as in Eitoku et al. display [11]. While each of these technologies is a novel approach to the 3D display problem, they are subject to some fundamental limitations. The nature of 3D data, being one dimension higher than a traditional raster display, means there is a tremendous amount of data that must be processed and transferred by the computer, often necessitating custom hardware. From a user interface perspective, each display creates its image in an enclosed volume that the user cannot penetrate without risking the display or their health. This limits the intuitive interaction a 3D scene affords, instead requiring additional work into user interfaces tailored to 3D displays [12]. One of the primary advantages of our use of the FogScreen is that its immaterial nature does not in any way prevent users from inserting their hands directly into the scene to select and manipulate objects naturally.

A popular alternative to volumetric 3D displays is to approximate the effect with a 2D display designed to augment the image with additional synthetic depth cues for increased 3D perception. The most common way to do this is stereoscopic imaging [13], possibly in surround-view projection environments, in which user-worn glasses enable the display of separate images to the left and right eyes, simulating binocular disparity. Autostereocopic displays [3] remove the need for glasses by using a lenticular lens or parallax barrier to separate images along different viewing directions. Stereo and autostereo displays both have particular ideal viewing locations where the effect is most distinct. Head-tracked rendering [14], [15] is often used in conjunction with stereo rendering to expand the ideal viewing region and provide an additional depth cue via motion parallax. These techniques are combined in head-mounted displays [16] for immersive perception of a 3D scene. Unfortunately, stereo techniques are subject to user fatigue during extended viewing from inaccuracies in the effect [17], [18] and the encumbrance of glasses.

More recently, an effect called DFD has been investigated [2] as another technique for simulating depth cues with 2D imagery. By rendering the same image on two overlapping screens at different depths, the binocular disparity and ocular accommodation at the two screens are fused into a single 3D perception in between. In addition to the simulation of multiple depth cues, the main advantage of DFD is that it avoids the fatigue problems of stereo displays [19] and does not require any user-worn glasses. This technique has been used for a prototype compact display [4], and the interaction between DFD and stereo imaging has been explored [20], [21], but always with two or three parallel screens and a single viewing location. One of our contributions is to show that DFD is still effective for arbitrary screen configurations and viewing locations.

## 3 Simulation of a General DFD Display

The generalized DFD principle is an important intermediate result on our path to the long-term goal of a truly volumetric walk-through display, using FogScreens as an enabling technology. There are many challenges to reaching that goal. Consider a stacked volumetric configuration of multiple FogScreens, in the spirit of the DepthCube display or volume rendering using axis-aligned textured rectangles [8]. One physical limitation is imposed simply by the dimensions and the operating mode of the FogScreen. The main generator unit of one FogScreen is about $2.0 \times 0.5 \times 0.5$ m in size, with the fog sheet reaching a thickness (depth) of 2 to 8 cm, sandwiched in between even thicker sheets of regulating airflow. Airflow interference causes turbulence when another unit is placed alongside of it. This alone imposes a minimum stacking distance of about 1 m. Even if the FogScreens were to become "thinner," there is no straightforward way to project a separate image onto each transparent screen plane. As the fog scatters incoming light, depending on the chosen fog density, a high percentage of the projected light gets transmitted through the screen and only a small portion gets reflected. This transparency is a necessary effect for the volumetric composition of a 3D image, but unlike the DepthCube display, we cannot time-multiplex the image creation. Hence, we have the problem of projector bleed-through onto nearby screens. One option we explored was the use of short-throw projectors to bring the image in at a very acute angle. But because the fog has thickness, this solution introduces smearing as light traverses the screen diagonally

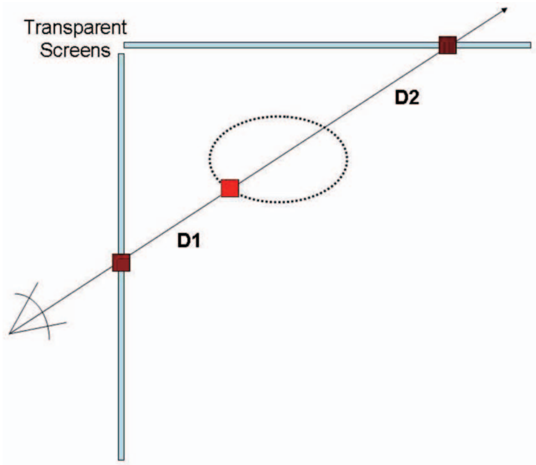


Fig. 2. The DFD effect on an L-shaped configuration.

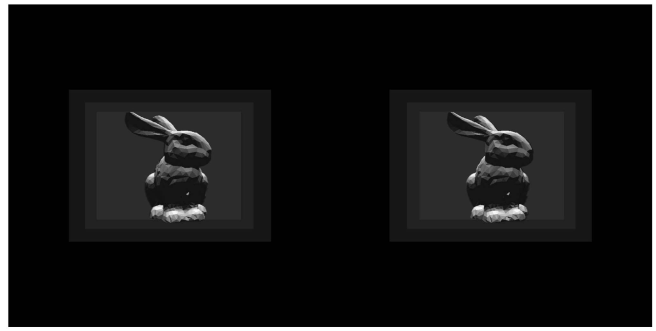

Fig. 3. Stereo pair for the DFD effect. There is no 3D model in the scene, but 2D textures on 3D screen planes. Planes are depicted for clarification purposes only.

and the image appears quite blurry to an observer with a viewing direction perpendicular to the screen. To minimize the bleed-through effect, we placed the FogScreens further apart (in one configuration) and at an angle to each other (in another) and used the DFD principle to achieve a 3D effect.

In this section, we demonstrate the feasibility of the DFD principle with arbitrary screen configurations and arbitrary user position using a stereoscopic 3D graphics simulator we implemented. Virtual transparent screens are observable to the user in different configurations. Each of these screens show a specifically calculated contribution of the whole 3D scene in between the screens using per-pixel accurate intensity values. When the individual screens overlap with the other screens, a 3D image impression is created in the visual system of the observer. Note that this still allows the user to freely move in and interact directly with the virtual scene, but several requirements and limitations of the DFD technique need to be mentioned: First, we need to track the user's head pose, since the 2D images displayed on each screen are dependent for the user's specific viewing direction and are computed in real time, and second, a 3D impression occurs only when the user looks in a direction where two or more screens overlap each other and depict objects in between them.

We evaluated the 3D perception users felt from the DFD-rendered images as compared to standard stereo and monoscopic rendering in a controlled user experiment, described in the next section. Figs. 3 and 4 show example stereo images of the DFD effect (Fig. 3) and plain 3D stereo (Fig. 4). Unlike the image in Fig. 3, the DFD images presented to the study participants did not have the semitransparent screens displayed. The reader of this paper is encouraged to cross their eyes on these figures to experience the DFD effect versus true binocular stereo.

Using the simulator, we can change the number and configuration of the employed transparent screens at will, and choose arbitrary vantage points without having to worry about tracking accuracy and physical screen limitations, enabling us to experiment with various setups, including configurations that are currently infeasible in the real world.

We used the simulator to explore what an observer could see when using the general DFD display in different configurations in real life. Each image that appears on a virtual screen has to be computed on the fly in 2D, and the final scene has to be rendered in stereo. Because of the stereo rendering of the texture mapped screen polygons, binocular disparity is accurately represented by the simulator, as is convergence, occlusion, perspective, motion parallax, height in the visual field, and, depending on the realism of the depicted 3D geometry, shading and possibly aerial perspective (or the scattering effect due to fog particles from our simulated display). Accommodation, however, is not accurately reflected, since the focus plane is fixed in both the head-worn display and the stereo projector we used to observe the simulator results. Accommodation is not a very strong depth cue, and by itself is not sufficient to bring out DFD depth impression [2]. On the other hand, we also know that it significantly helps depth impression, when accommodation is in sync with convergence and disparity [20]. The simulator version used in this work represents screens as simple semitransparent polygons onto which the projected images are applied using 2D textures calculated on the fly in offscreen buffers. To do this, we render the geometry from the virtual user's point of view using head-tracked rendering [14], [15]. We do this once per screen using a standard offscreen rendering technique. In the first rendering pass, we calculate the luminance of each pixel on each individual screen. Using the DFD principle [2], we cast a ray from the user through the geometry to each pixel to determine the object's depth at that pixel. The brightness of each pixel is the distance ratio of the object (at that pixel) to its neighboring screens as shown in Fig. 2. These rendered images are stored to offscreen buffers.

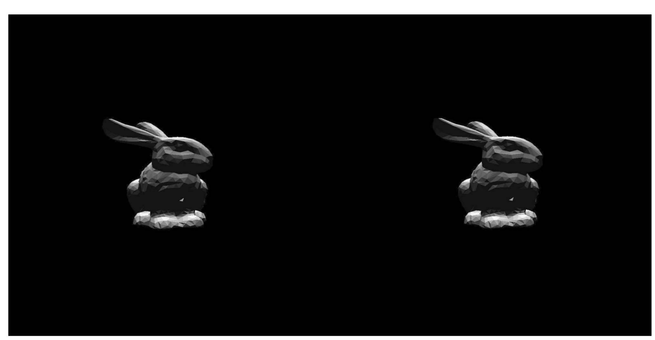

Fig. 4. Stereo pair for true binocular stereo. Here, the bunny is a 3D model as seen by the right and left eye.

Fig. 5. Different screen configurations: cross (discarded), L-shape, stack, and triangle.

In the final rendering step, we define a normal stereo camera at the user's position, map our rendered images to our transparent screens, and render the whole scene in stereo. This accurately simulates what would occur on the real display assuming perfect tracking. The user views this simulated environment through either an HMD or stereo projector with shutter glasses.

We experimented with a variety of configurations: stacked, crossed, L shaped, and triangle (see Fig. 5). The cross configuration was discarded because there would be effectively one transparent screen at the center of the scene. As a result there would be the least DFD effect at the most critical part of the scene. The remaining configurations were evaluated in a user study.

## 4 Evaluation of the Simulator

To evaluate the effectiveness of a general DFD display, we conducted a study comparing the 3D perception of different display configurations within our simulator.

### 4.1 Design

Our study consisted of 16 subjects, 5 female and 11 male, ranging in age between 22 and 26, all familiar with computers and computer games, but only a third with any experience with stereo imagery. The study used a within subjects design. The evaluation system was a DepthQ stereo projector and a standard 6 ft. white projection screen. Users were instructed to stand on a line approximately 8 ft. from the screen, wearing active shutter glasses. To test users' ability to perceive stereo images, we first presented each with a random dot stereogram. Users who were unable to describe the object in the stereogram were eliminated from the study. Of the 16 users we began with, one was unable to perceive stereo.

For the remaining users, we displayed a series of different static images (see Fig. 6) and asked them to rate how 3D the depicted object appeared on a scale of zero to five, zero being totally flat, and five being totally 3D. We also encouraged users to give feedback on what they perceived. The images users evaluated each showed the same 3D object in the same orientation, in different display technique scenarios. There were seven scenarios total, each shown three times, in random order. Between each trial, the screen was blanked for 5 seconds, to avoid direct comparisons. To ensure consistency across different users' experiences, no user interaction was possible. The particular scenarios that were tested are as follows (see Fig. 6 for images).

The stack scenario has three screens arranged in a stacked, parallel configuration with the images on each screen rendered using the DFD technique. The screens are then rendered in stereo. The user is located centered in front of and perpendicular to the screens, so they all overlapped providing three planes for the DFD effect. This scenario tests the established DFD results in our simulator, to evaluate how well our system mimics a true DFD display's qualities.

The off-axis scenario uses the same stacked configuration as the stack scenario, but the user's position is moved off center, so the screens are viewed from an angle. This tests the perception of the DFD effect for parallel screens with head-tracked rendering, which we predict will match the results of the regular DFD display in the stack scenario.

The triangle scenario is the first scenario to test a novel DFD display configuration. Three screens are arranged to form a triangle, with the user centered in front of one side. Images for the screens are rendered using the DFD technique and the screens are rendered in stereo. As our hypothesis is that general DFD displays perform as well as the traditional case, we predict this scenario's ratings will be

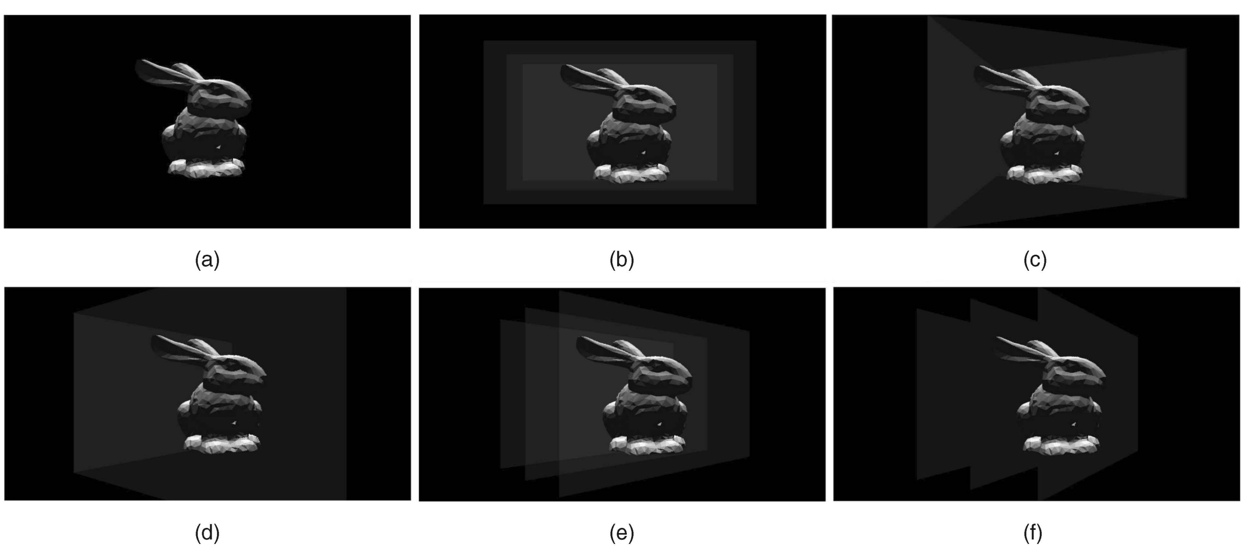


Fig. 6. Scenes used in controlled generalized DFD user study: (a) mono, (b) stacked planes, (c) triangle shape, (d) L-shape, (e) off-axis stacked planes, and (f) unblended off-axis stack (for control purpose). Stereo is not shown here and these screens are for clarification purposes only.

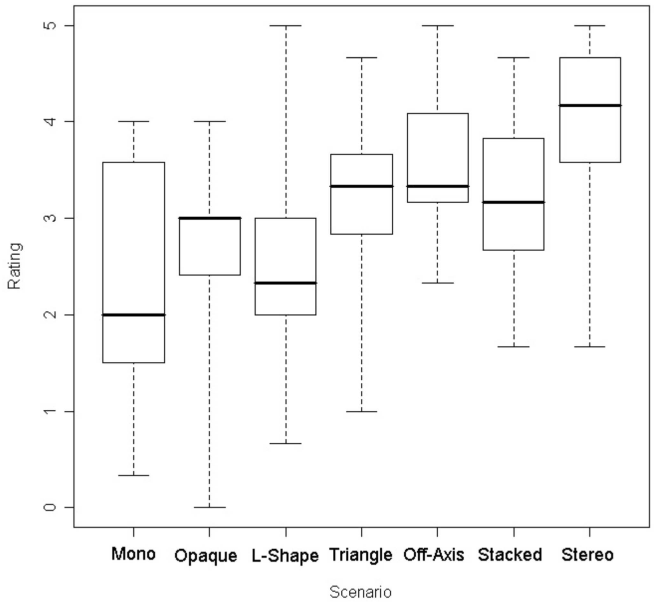


Fig. 7. Boxplot of users' ratings for each scenario. Each column shows the 0th, 25th, 50th, 75th, and 100th percentiles.

TABLE 1
Tukey Multiple Comparisons of Means with 95 Percent Family-Wise Confidence Level across All Scenarios

| Scenario | P Adj. |
|---|---|
| Mono vs L Shape | 0.9993517 |
| Mono vs Opaque | 0.7560168 |
| Mono vs Stacked | 0.0168194 |
| Mono vs Triangle | 0.0063395 |
| Mono vs Off Axis | 0.0000477 |
| Mono vs Stereo | 0.0000000 |
| L Shape vs Opaque | 0.9452921 |
| L Shape vs Stacked | 0.0587850 |
| L Shape vs Triangle | 0.0249445 |
| L Shape vs Off Axis | 0.0002712 |
| L Shape vs Stereo | 0.0000001 |
| Opaque vs Stacked | 0.4680931 |
| Opaque vs Triangle | 0.2850091 |
| Opaque vs Off Axis | 0.0097243 |
| Opaque vs Stereo | 0.0000093 |
| Stacked vs Triangle | 0.9999183 |
| Stacked vs Off Axis | 0.6449223 |
| Stacked vs Stereo | 0.0111768 |
| Triangle vs Off Axis | 0.8286272 |
| Triangle vs Stereo | 0.0283513 |
| Off Axis vs Stereo | 0.4971054 |

similar to those of the stack scenario. The L-shape scenario tests the effect of an edge artifact with two screens in an L configuration, oriented so that the overlapping region only covers half of the 3D object. We call this type of depth disparity an edge artifact. Images on the screens are rendered using the DFD technique and the screens are rendered in stereo. Because of the edge artifact, we predict users will perceive a 3D image of low quality, and that the overall rating will be less than the other DFD scenarios, but still higher than a 2D display.

The opaque scenario is a more extreme case than the off-axis scenario, with the user's position far enough off center that portions of the model are on nonoverlapping portions of the screens. Also, the virtual screen images are not rendered transparently, so there is no DFD effect. The purpose of this scenario is to see what effect, if any, the use of stacked screens has on 3D perception without the influence of the DFD technique. Since some 3D information is available, we expect it will be rated higher than a 2D display, but less than scenarios with the DFD technique. The stereo scenario is normal stereo rendering of the model geometry without any DFD effect. The purpose of this scenario is to provide a measurement of the best possible 3D perception result on our display, and so we predict it will have the highest rating in the study. The mono scenario is the same as the stack scenario except the final image is displayed without stereo. Therefore, there are no extra depth cues to be perceived and the user should see a flat 2D image. This provides a baseline measurement of the worst possible 3D effect on our display, and we expect it to have the lowest overall rating.

## 4.2 Results

We generated a single rating by each user per scenario by averaging the user's ratings on the three trials. A one-way within-subjects ANOVA [22] of the user's ratings versus seven scenario treatments showed a strong statistical significance among the results ($F(6, 84) = 12.791,\ p < 0.001$). Fig. 7 shows the aggregated ratings for each scenario. We also did a post hoc analysis using Tukey's Multiple Comparisons of Means and the results are shown in Table 1. Stereo is clearly the best rated and is significantly different from every other scenario except for off-axis. Off-axis is not significantly different from stereo ($p < 0.49711$). The next highest rated are stack, off-axis, and triangle, which are not significantly different with respect to each other. Mono was rated the lowest, significantly different from stack, triangle and off-axis, and stereo ($p < 0.01682$, $p < 0.00634$, $p < 0.00005$, $p < 0.00001$, respectively). Finally, L-shape and opaque were not significantly different from mono ($p < 0.99935$ and $p < 0.75602$, respectively).

We expected stereo to be rated the highest, and our results confirm that expectation. What is somewhat surprising is that off-axis is not significantly different from stereo. This demonstrates that under particular viewing conditions, DFD viewing can be similar in quality to traditional stereo rendering. It is also reassuring to see that mono is the lowest rated, though with a high variance. Some users liked the plain 2D image the best, describing it as very clear. We suspect this is partially due to unfamiliarity with stereo and DFD viewing, and the observer is confusing proper lit shading with stereoscopic 3D perception.

The rating of stack confirms established results on the DFD effect [2] stating that the 3D perception on a standard DFD display is improved over standard 2D displays, but not as high fidelity as good stereo techniques. Our prediction of little difference between triangle and stack is also confirmed, which supports the idea that the DFD effect will work in conjunction with head-tracked rendering for 3D perception from multiple viewpoints.

The ratings for L-shape are important to consider. The difference between L-shape and triangle is the large edge artifact in the middle of L-shape, and the result this artifact has on the perception is clearly reflected in the ratings. Users also commented on the image being blurry and disjoint. While this result appears to show the poor performance of an L-shaped configuration, the triangle configuration is very similar and performs well. The outcome of this result is to underscore the importance of proper positioning of the screens and user to ensure the maximal region of screen overlap in a general DFD display.

Finally, the opaque rating verifies that stacked, opaque displays are not sufficient to create a natural 3D perception, even when seen from a side angle, and suggests that the DFD technique with its coordinated pixel intensities on transparent screens is critically important to a high-quality 3D visual on multiscreen displays.

## 5 Prototype of a DFD Display

The choices for the configurations we tested with our prototype were based on the results of the user study and the fact that we had access to only two FogScreens. The stacked configuration performed very well, in both the on- and off-axis positions, and was chosen for this reason. Even though the stacked configuration in the simulator used three screens and the actual prototype only uses two, there should be no significant difference in the DFD effect perceived by users. The DFD principle does not rely on the number of screens. As long as there are at least two screens, the virtual object appears to exist continuously within the space enclosed by them. The L-shaped configuration was chosen because it was the closest feasible physical representation of the triangle configuration, which performed second best in the user study. In the user study, we had intentionally positioned the user and bunny in the L-shaped configuration such that a part of the bunny was perceived in mono (see Fig. 6d), in order to evaluate that effect. Participants were able to perceive the edge artifacts and from their comments gave a lower ranking due to these artifacts, and not due to the configuration itself. Unlike the triangle configuration, the L-shaped configuration only requires two FogScreens and unlike the stacked configuration, it does not suffer from any bleed-through problems since the images are projected orthogonally to each other.

The system we assembled uses two FogScreens, each with their own standard DLP projector, in the stacked and L-shaped configurations. We also evaluated a setup using as display areas one FogScreen and a wall. For head-tracking, we use WorldViz's Precision Position Tracker [23], which tracks the 3DOF position of an infrared LED inside our viewing volume using four infrared cameras placed around the display system. The displays are driven by a single desktop computer with a Quadro FX 4500 graphics card. The images on the screens are generated using the same DFD technique implementation as in the simulator, to ensure visuals are consistent across the two systems.

In the stacked configuration, the two screens are parallel to one another (see Fig. 8). Its implementation in our prototype is hindered by the limitations of the FogScreen. Because the FogScreen transmits most of the light projected

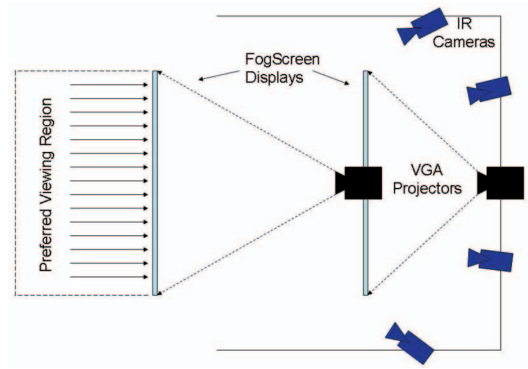


Fig. 8. System overview for stacked Dual-FogScreen setup. Distance between screens is large in order to avoid bleed-through from angular projection.

onto it, the screens cannot be mounted too close together, or the image from the rear screen will bleed through to obscure part of the front screen. We experimented with using short-throw projectors that project from a very steep angle to allow mounting the FogScreens closer together, but the nonzero thickness of the fog plane creates a significant pixel smearing effect for off-axis projection that seriously reduces image quality. Our final configuration compromises among these limitations and places the screens 2 m apart, which allows us to project onto the screens from the top back at an angle of about 26.5 degrees to the horizontal, without incurring a bleed-through overlap.

We also experimented with a stacked configuration which consists of a single FogScreen and a wall. This setup was prompted by two factors: there is considerably less turbulence with just one FogScreen (alleviating such problems as depicted in Fig. 13), and, furthermore, it is a simpler and cheaper approach, acknowledging that people would more easily get access to a single immaterial screen than two. For this setup, it is a little harder to calibrate the brightness levels on the two screens, since the perceived brightness of projection onto a white wall is higher and, unlike projection onto a FogScreen, mostly independent from viewing angle. The image quality the two screens provide is obviously different (the front screen being a dynamic fog layer, the back screen a stable opaque wall), but after brightness adjustment, our test users were able to fuse the DFD imagery. With this setup, it is possible to use a short-throw projector on the back (wall) screen and it, thus, allows more freedom in the placement of the FogScreen (see Fig. 9). A study to quantify which solution (two FogScreens versus FogScreen and wall) yields a better DFD effect is planned as future work. First, qualitative comparisons indicated that, after calibration, the two solutions afforded similar levels of depth perception, but the smaller amount of turbulence favors the FogScreen/wall approach.

In the L-shaped configuration, the two screens are mounted to form a right angle (see Figs. 10 and 11). Proper selection of the viewing location to the region where the virtual geometry is contained within overlapping regions of the screens alleviates this artifact and is more similar to the results from the triangle configuration in the user study. The advantage of the L-shaped configuration is that it places the screens and projectors in such a way that the rear

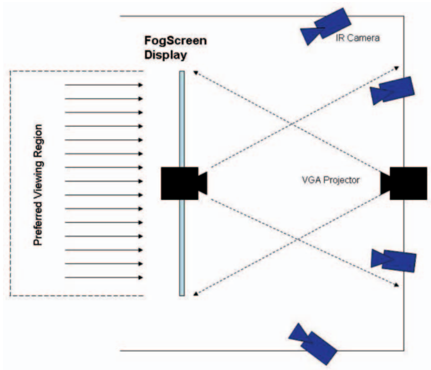


Fig. 9. System overview for single-screen stacked setup. One FogScreen and a wall is used to create a stacked configuration.

image never bleeds onto the front image, as occurs in the stacked configuration when projection angles and screen distances are not set up carefully.

For each of these configurations, we informally evaluated the quality of the DFD perception. What we found confirmed the results from our user study. First, in the stacked configuration—with either one or two FogScreens—and with the user centered and perpendicular to the screens, the DFD effect was clear, resulting in 3D perception as reported previously [2]. Second, as the user moves around the display, the head-tracked DFD rendering maintains the 3D perception where tracking is good, confirming our result that the DFD effect continues to work for arbitrary viewpoints. Finally, in the L-shaped configuration, users are still able to perceive the correct 3D image, demonstrating that arbitrary screen configurations can still support the DFD effect.

## 6 General Discussion

While a generalized DFD display using FogScreens is very appealing, we need to consider its limitations. This section discusses the main challenges to an unencumbered 3D imaging experience with our prototype. We separate these limitations into three categories associated with the nature of the technologies used: fog projection, fog dynamics, and tracking limitations. We do this with the intention to define metrics that determine the working parameters for successful depth perception.

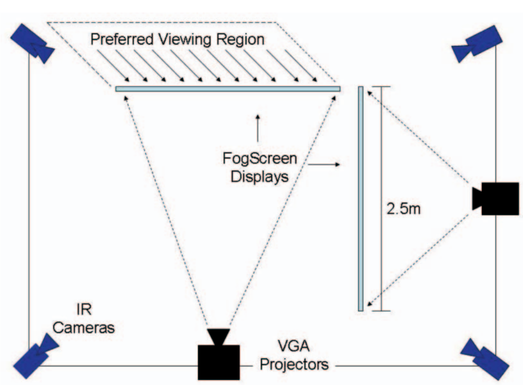


Fig. 10. System overview for L-shaped Dual-FogScreen setup. No bleed-through, but limited screen overlap.

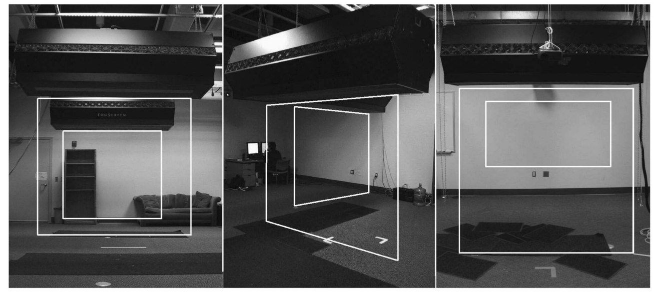

(a) (b) (c)

Fig. 11. Photographs of physical setup with FogScreen dimensions indicated by outlines. (a) Stacked configuration. (b) L-shaped configuration. (c) Single-screen stacked configuration.

### 6.1 Projection onto Fog

Obviously, not all projectors are equivalent. The first major problem is due to the intensity differences between different projectors. Because the DFD effect is produced by manipulating the intensity of corresponding pixels on each transparent layer, we need to calibrate the projectors. The base intensity of all projectors must be leveled for the DFD effect to work as expected. Another intrinsic problem is due to color discrepancies between projectors. We found this problem to be minor in comparison to intensity issues. For our prototype, we manually calibrated both projectors for both intensity and color. Another major problem has to do with a nonlinearity in perceived intensity on FogScreen displays. For projected imagery that is dimmer than a certain threshold, we witnessed a sharp perceived intensity drop-off of the final images as displayed by projectors. For example, an image which has a 40-percent-intensity value may be seen using our simulator but may not be perceived by users of our real life prototype, while, say, a 50-percent-intensity image could be perceived. In the following sections, we will report controlled experiments to determine the respective intensity thresholds beyond which the DFD effect breaks down. The last major limitation is angle of projection. As mentioned above, projection onto a FogScreen at steep angles compared to the observer's viewing direction leads to low-quality washed-out images, since the screen is transparent and has a certain thickness. At the same time, we must carefully place projectors so no bleed-through occurs. Optimizing projector placement, orientation, and observer's viewing angle is a challenging task. Ideally, the image generation engine for the FogScreen DFD effect would automatically correct for all of these influence factors and adapt pixel intensities accordingly. While this is theoretically possible, we do not yet have a complete model of the fog's scattering properties and visual behavior with respect to projection angles and other influence factors.

### 6.2 Fog Dynamics

With the FogScreens, we found turbulence, bent screen geometry, and inconsistent fog density to be the most noticeable problems. Normally, the FogScreen has a small amount of turbulence due to the corrective air flow that sandwiches the fog layer to create a stable sheet of fog. However, when environmental factors are introduced, turbulence and screen deformations can become a problem.

Fig. 12. The stacked Dual-FogScreen setup can provide stable imagery when ventilation in the room is controlled and the airflow for the two screens is adjusted to avoid cross-interference. Even standing within a screen is then possible without major disturbance.

Normal air flow created by air conditioning vents, people moving around the screens, and air flow from other screens can cause the fog sheet to bend, waver, and fluctuate. Such higher turbulence can easily be detected by users and break the DFD effect (see Figs. 12 and 13). Fog density is manually adjustable, but the same setting on different FogScreens is often not consistent because of wear and tear due to the age of each FogScreen. In our setup, we manually adjust the screens so the screen qualities are subjectively the same, as independently judged by two screen experts.

## 6.3 Tracking

Tracking accuracy is another limiting factor. In our prototype, we used an optical tracking solution. The WorldViz PPT tracker we used produces errors of less than 0.5 cm over a $3.0\ \mathrm{m} \times 3.0\ \mathrm{m} \times 3.0\ \mathrm{m}$ viewing volume. This accuracy drops off when a bigger room, fog occlusion, and camera locations are factored in. For the prototype display, our software needs to register the position of the screens with the position of the user. We found that there was some jitter associated with the user's registered position, and at times, it was not possible to fuse the images or see the DFD effect. We call this a tracker-induced *alignment error* and we derived a metric which measures the alignment error associated with registration problems over the entire image. For this metric, we assume that the user is focusing on a certain pixel on the front screen, and measure the angular distance of the corresponding pixel on the back screen to the pixel that is *actually* hit by the viewing ray. In other words, the alignment error, as illustrated in Fig. 14, is the angle $\theta$ between the pixel on the back screen which is actually aligned with the front screen pixel (head of blue arrow) and the real pixel it should have been aligned with respect to the front pixel (head of red arrow). The metric for evaluating the error is then computed as the root mean of the squares of the alignment error over every pixel within a certain object boundary (e.g., all pixels that constitute the bunny image). Initial evaluations indicate that this angular error metric is reasonably independent of actual screen dimensions and distances, but somewhat dependent on how big the depicted object is with respect to these distances. The same alignment error value can represent displacement of a significant percentage of a very small object or just a small percentage overlap of a very large object. For objects of comparable size, this metric is expected to be a good predictor of a user's ability to fuse a given DFD image pair. Going forward then, while this metric gives us a measure of the error as it relates to misregistration, we also wanted a quantitative measure of how sensitive the DFD effect was to this particular limitation. In the next sections, we discuss user evaluations we performed to explore both the acceptable thresholds of screen brightness and the sensitivity of the DFD effect to misregistration.

## 6.4 Binocular Parallax

All our real-life prototype setups are subject to an alignment problem due to horizontal binocular parallax. There is an inherent error due to the fact that the images rendered on the front and back screens are rendered assuming a pinhole camera situated between the focal centers of our eyes (Fig. 15). Since our eyes are set apart horizontally, horizontal alignment of two corresponding images at different depths

Fig. 13. Under unfavorable conditions, turbulence effects triggered by air conditioning and cross-interference among the airflow from the two screens can be quite pronounced and do disrupt the DFD effect. These situations are much reduced for a 1-screen/1-wall setup.

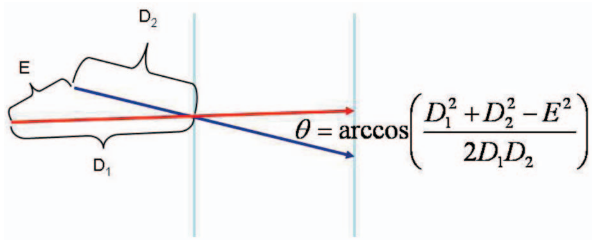


Fig. 14. Alignment error $\theta$. Distance error E is the real-world distance between the actual user's position and the user's position assumed by our software. D1 and D2 are the distances between the actual and assumed user positions to a specific pixel on the front screen. $\theta$ is the alignment error angle for this one pixel. We refer to the root mean square of this value over all pixels defining the 3D object as the aggregated alignment error $\theta_{agg}$.

with both eyes open is always a compromise. The error seen by each eye can be computed in the same manner as the alignment error. Users cope with the binocular parallax problem by centering the object on the front screen with the outline on the back screen, accepting blurred edges in the horizontal dimension. We explored the possibility of correcting this effect by adjusting the software so that the back screen's image would show the two images rendered with cameras set in the left and right eye. We tested image combination through either a maximum intensity projection or by first rendering the nondominant eye and then painting it over with the image for the dominant eye. The reasoning for altering the back image and not the front image was that most viewers naturally choose to focus on the front image and then align with the back image. We also considered rendering the object on both screens as if vergence occurs at the real object distance in between the two screens, which would result in an object with blurred boundary on *both* screens. This did not improve the user's ability to fuse the images. In general, our attempts to minimize the horizontal binocular parallax problem by two-eyed rendering were not successful. On the contrary, they were clearly a detriment to the DFD effect for small objects since these methods introduced blurred object boundaries. It should be noted that this problem does not prevent the fusion of a 3D image in the viewer's brain. However, it makes it somewhat more difficult to arrive at a fused image. In order to de-emphasize the horizontal parallax effect in the controlled user evaluations we will describe in the next sections, we decided to place the user about 4 m away from the front screen in those experiments. This is further away than is desirable for a technology with reach-through potential, but it was done to isolate and quantify other influence factors, in particular intensity mismatch and tracking error. By increasing the distance of the observer to the front screen relative to the fixed distance between the two screens, we can diminish the alignment errors due to horizontal parallax. This is not a prerequisite distance to experience the DFD effect. Even at less than half that distance, some of our users were still able to fuse the 3D image. The amount of adaptation to binocular parallax appeared to be user-dependent.

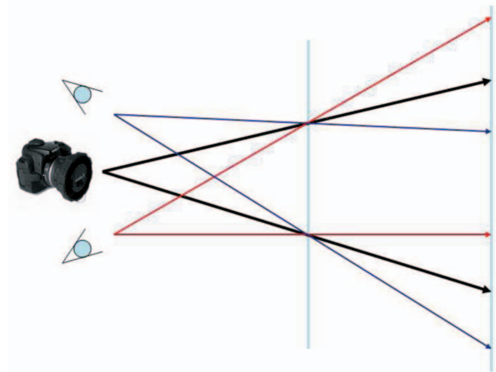

Fig. 15. Alignment error due to horizontal parallax. Images on front and back screen are rendered using pinhole camera, but seen with two eyes. The intersecting rays show an exaggerated error that is inherent to every user.

## 7 Expert User Pilot Study

We conducted a formative pilot study with researchers who were experienced with perceiving the DFD effect. The goal was to discover consistent trends toward intensity and alignment error thresholds that would make or break the DFD effect among our users. The participants in this evaluation were four members of our research lab. Each user was familiar with the DFD display and could successfully fuse the images in a calibrated setup. All users were also familiar with VR and AR software. We decided on using the stacked configuration setup for a number of reasons. We were limited by the room we were using and the number of available FogScreens (two), and compared to the L-shaped configuration, the stack configuration provides for more even intensity among the two screens as the user can view the setup from an orthogonal viewing axis. We placed the screens approximately 2 m apart and placed the user approximately 4 m away from the front screen. This setup reduces the alignment error associated with horizontal parallax. Users were asked to confirm a DFD effect before proceeding with the study. Fig. 16 shows the setup.

### 7.1 Intensity Thresholds

For the first task, we manually calibrated both projectors for an optimal DFD effect. This initial state was the same for all users and each user was asked to confirm if they could successfully fuse the images into a 3D image. Once confirmation was received, one screen's intensity was set to 0 in our software, leaving a single 2D image on the other screen. The intensity of the first screen was then gradually increased and the user was asked to report when he could successfully fuse the images. Once this was confirmed and the threshold intensity logged as a percentage of the original base intensity, the intensity (for the same screen) was gradually increased again. The user was asked to confirm when he could no longer fuse the images and we again logged the result. Once both observations had been made, we reset the intensities to the initial state, asked for confirmation of the user's seeing a DFD image, and repeated the process with the second screen. Each screen was tested three times, and we recorded low and high thresholds for both screens, amounting to 12 observations

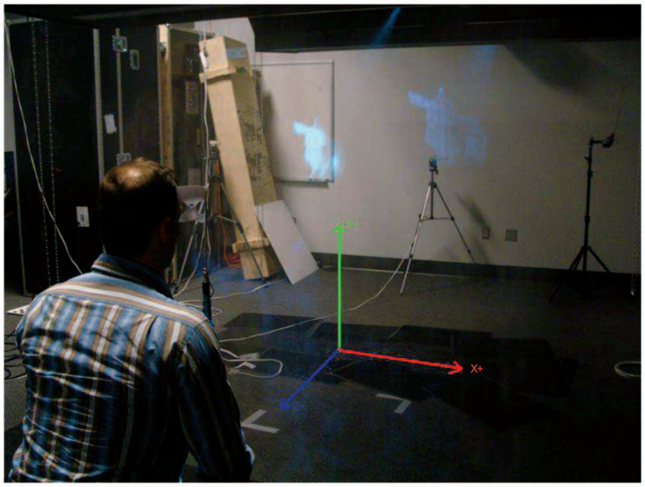


Fig. 16. User looking through a rigid mask. $X, Y, Z$-axes are marked with red, green, blue, respectively. The lights were turned off in the actual test. From the user's point of view, the two bunnies visible from this vantage point were aligned and fused into a single bunny.

per user. It is important to note that only one screen was adjusted at a time and the respective other screen would remain at its original calibrated intensity. As an example for the percentage values, logged intensity thresholds of 75 percent and 125 percent for the back screen would mean that, at 75 percent of the original back screen intensity (as calibrated before the study), the user began to see the DFD effect while the front screen was at 100 percent of the calibrated value, and when the back screen reached 125 percent of the calibrated intensity, the DFD effect broke down for the user because the back image became too bright.

### 7.2 Alignment Error Sensitivity

For the second task, we again manually calibrated the projectors to optimal levels as before the previous task. Each user was then asked to look through a rigid mask on a tripod onto the screens (see Fig. 16). No head tracking was used. Instead, the user location (position of the tripod) was calibrated so that the images could be perfectly fused while looking through the mask. Once a DFD effect was confirmed, we asked the participant to close their eyes and a small random error was introduced to the user position. The error was random along all three axes and measured between 0 and 5 cm in length. We chose 5 cm because that capped any expected errors in registration for a properly calibrated PPT system. The user was then asked to open their eyes and confirm whether DFD was observable. The reported result and calculated error (as the root mean square over all object pixels of the value $\theta$ from Fig. 14) were automatically logged by the software. We repeated this process for 100 observations per user and the user was required to confirm DFD with the initial state and close their eyes every time a new random error was introduced.

### 7.3 Expert User Pilot Study: Results

The results from the intensity thresholds test were very consistent. The low and high means of the thresholds for the front screen were 79 percent and 138 percent, respectively. The back screen's low and high thresholds were 90 percent and 157 percent, respectively. The variance in the responses of the four users for the low and high means of the front screen were 1 percent and 2 percent, respectively. For the back screen, the variance was 0.4 percent and 12 percent, respectively. These results implied to us that the responses from users are generally similar with respect to the low and high thresholds of the front screen and the low threshold of the back screen. The high threshold for the back screen was not quite as consistent, but, arguing conservatively, we hypothesized for a follow-up summative user study (see Section 8) that a variance of a few percent in the intensities of either screen should be well within the limits of continued DFD experience, all other things being equal.

We designed the alignment error sensitivity test in the hope to arrive at a error metric threshold that would define the neighborhood of a cutoff point between DFD fusion and nonfusion of the two layered images, for a general expert observer. The results from the pilot study were sufficiently structured to provide insight in this direction. Four hundred observations were made with the same four expert users as in the intensity study, and random errors were introduced. From the 400 observations, 162 were reported as "unfused," where no DFD effect was perceived by the user, and 238 were reported as "fused," where a DFD effect was perceived. The mean value of all per-observation total error values (root mean of the squares of the error theta) was 0.44 degrees and 0.25 degrees for unfused and fused, respectively. The medians were 0.44 degrees and 0.19 degrees. The variance was 0.04 and 0.04, respectively. With these results, we performed a t-test of two samples, one of all fused observations and the other of all unfused observations, assuming equal variance and a hypothesized mean difference of 0.0 and $\mathrm{alpha} = 0.05$. We found there to be a significant difference between the fused and unfused observations $(\mathrm{dF} = 398,\ \text{t-Stat} = 9.16,\ \text{two-tail t-Crit} = 1.966,\ \mathrm{p} < 2.85\text{e-}18)$. As far as a single error threshold is concerned, we are aware that in spite of the statistically significant differences between the fused and unfused sample sets, there exists no such value that would separate the observations perfectly. As an approximation to such a single threshold, we present the mean of the medians for the fused and unfused samples: 0.31 degrees (aggregated as the root mean square angular distance over all bunny pixels.

Fig. 17 shows the distributions of all 400 fused and unfused observations from the pilot study. It also indicates the medians for all fused observations (0.19 degrees), of the unfused ones (0.44 degrees), as well as the mean of the medians (0.31 degrees).

During the course of our study, we had to make major efforts to control air turbulence as an impact factor. As illustrated in Fig. 13, with multiple FogScreens, open doors, and/or air conditioning, the shape of the FogScreens sometimes wavered to a point that it clearly influenced the perception of the DFD effect. In general, while the pilot study indicated interesting trends and hypotheses, in order to come closer to our goal of identifying and quantifying at least two influence factors for the successful fusing of DFD images in our prototype, namely intensity and alignment thresholds, we needed a user study with more participants and more rigorously controlled airflow.

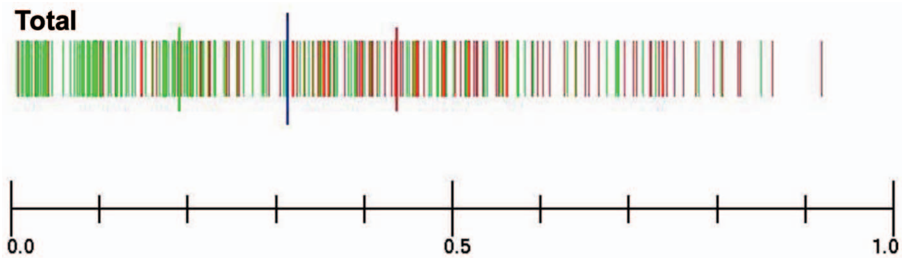


Fig. 17. Plot of pixel-aggregated alignment error ($\theta_{agg}$ in degrees for all 400 observations from the pilot study. Green vertical line segments represent fused observations, red line segments indicate unfused ones. The blue bar represents the mean of the two medians.

## 8 General User Study with FogScreen/Wall Configuration

We decided to run the same experiments from our pilot study with more users, but with an altered setup; a single FogScreen setup instead of two FogScreens, illustrated in Fig. 18. We verified that there would be considerably less turbulence in absence of a second FogScreen unit. For the second image plane, we used one wall from our lab space. The FogScreen was parallel to the wall and set approximately 4 m from the wall. Users were also positioned approximately 4 m from the FogScreen although each user's actual position was adjusted for the best possible DFD effect considering their individual heights. In this study, we had 14 volunteers, drawn from the UCSB student body, providing no payment or other incentives. Users ranged from 22 to 27 years old, 4 females and 10 males. The users were varied in their experience with computer graphics, from none to expert. None of them was initially familiar with the DFD effect. As in the pilot study, we took 12 observations regarding image intensity and 100 observations regarding alignment error from each user.

Before each subject was tested, we performed an initial training period. All users were tested for stereo-viewing abilities (using random-dot stereograms) and all passed. Each user was then shown a set of 2D, stereo, and DFD images on a head-mounted display to form an impression on what constitutes a 3D imagery. We then manually calibrated the projectors for the best possible color consistency between the two final images, and in our software, manually adjusted the intensity of the final images rendered so that the intensity of the perceived images was equivalent. Each user was asked to confirm DFD perception before each observation, just as in the pilot study.

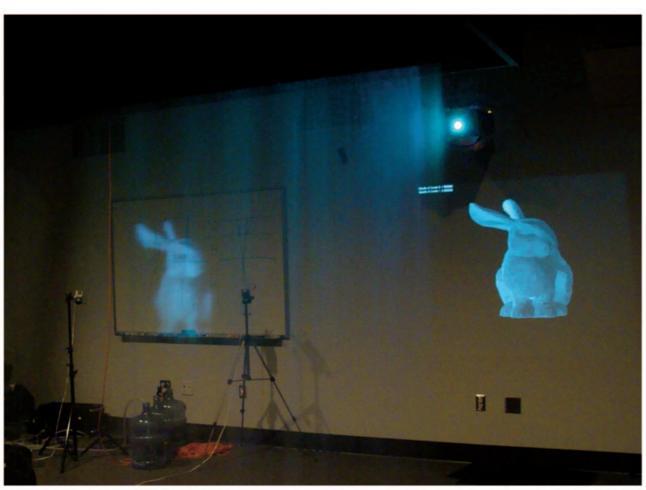

Fig. 18. Screen configuration for summative user study with one FogScreen and a white wall.

### 8.1 General User Study: Results

The results for the intensity thresholds showed the low and high means of the intensity of the back screen (the wall) to be 55.63 percent and 150.03 percent of the initial intensity (Fig. 19). For the front screen, the low and high means were 63.83 percent and 155.53 percent. The variance in the responses from the 14 users for the front screen's low and high means was 1.36 percent and 3.91 percent, while the variance for the low and high means of the back screen were 3.82 percent and 9.70 percent, respectively. While the absolute numbers were different from our first pilot study, and the variance was also higher, we see one particular pattern confirmed, namely that the high threshold for the back screen exhibits a much higher variance than the other three threshold values. Overall, users' sensitivity levels are not uniform and any application might benefit from being calibrated for the individual user. However, in view of large intervals of nondisrupted DFD effects, and comparatively small variances, our conservative hypothesis, that "a variance of a few percent in the intensities of either screen should be well within the limits of continued DFD experience" holds, even for variances around 10 percent.

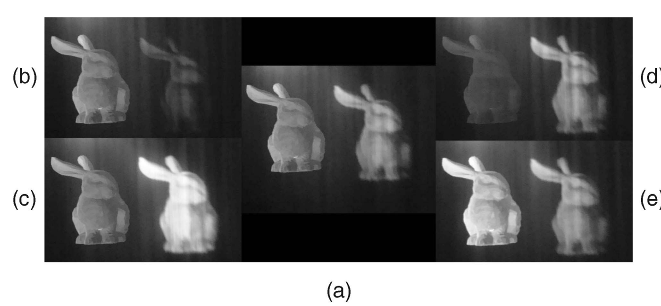


Fig. 19. Images of image intensity threshold means observed. (a) Manually calibrated system for optimal DFD effect. (b) Front plane at 63.8 percent of optimal and back plane at 100 percent. (c) Front plane set at 155 percent and the back at 100 percent of optimal. (d) 100 percent of optimal in the front and 55.63 percent of optimal at the back. (e) 100 percent of optimal in front and 150.03 percent optimal in back.

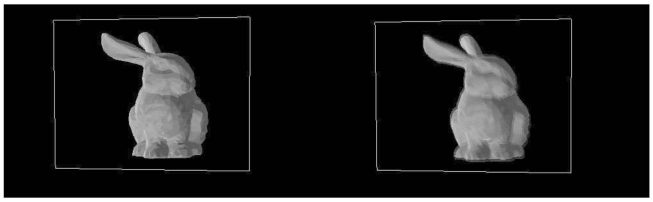

Fig. 20. Comparison of a perfectly aligned DFD bunny with one that exhibits the mean of medians threshold determined from our 14 user, 1,330 sample user study (0.35 degrees). The direction of the original translation error leading to the 0.35 angular alignment error was arbitrarily chosen. Screen outlines are presented solely as a reminder that there are two screen planes involved.

For the alignment error, the results were again significantly different for fused and unfused responses over all users. From the 14 users, we collected 1,330 observations (50 samples from one user had to be discarded due to miscommunication, and 20 from the rest of the users were lost due to logging errors), resulting in 564 unfused observations and 766 fused observations (Fig. 20). The means of the unfused and fused observations were 0.45 degrees and 0.25 degrees and the variance was $< 0.044$ degrees and $< 0.033$ degrees, respectively. We performed a t-test of the two samples which confirmed a highly statistically significant difference between the two samples ($\mathrm{dF} = 1{,}328$, $\mathrm{t\text{-}Stat} = 18.30$, $\mathrm{t\ Critical\ two\text{-}tail} = 1.96$, and $\mathrm{p} < 6.92\mathrm{E}\text{-}67$). Our approximation of the cutoff threshold between fused and unfused responses, the mean of the medians of the response sets, was 0.35 degrees, as can be seen in Fig. 21. These results confirm the findings from our pilot study. Among the four expert users, we had arrived at a cutoff threshold of 0.31 degrees. This is even more reassuring insomuch as the distances between the screens were different as compared to the pilot study (4 m between FogScreen and wall versus 2 m between the two FogScreens). This is supporting evidence that our metric $\theta_{agg}$ is indeed likely independent of the screen configuration, as planned. We seem to zero in to something like a threshold value (0.35 degrees) around which the likelihood of fusing and nonfusing 3D imagery gets higher or lower, depending on which side of the threshold you are.

The unfused responses from users were significantly different along the three viewing axes (assuming the user is looking in the $Z$-axis, with $Y$ in the up direction). Taking only the responses for which the user position error was greater along one axis than the other axes by 150 percent for all three axes, we did indeed find the results to be statistically significant with a single-factor ANOVA ($\mathrm{F}(2{,}287) = 31.90, \mathrm{p} < 3.10\mathrm{E}\text{-}13$). The means were 0.51 degrees, 0.49 degrees, and 0.28 degrees on the $X$, $Y$, and

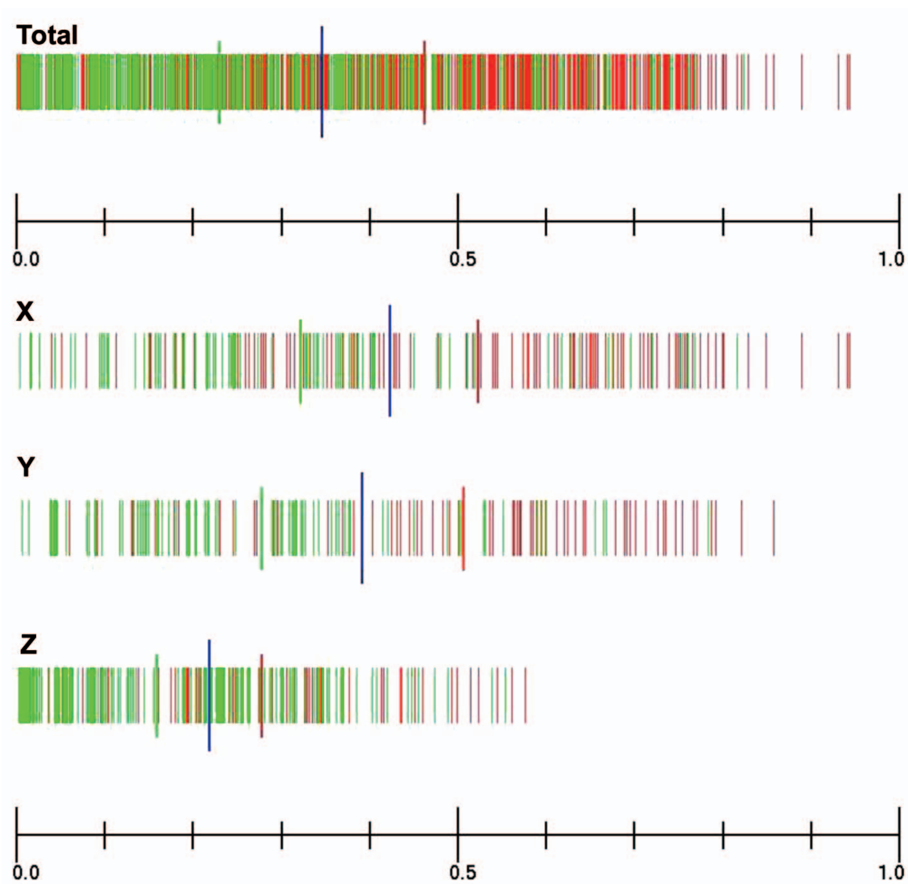


Fig. 21. Plot of alignment error $\theta$ in degrees for 1) all observations (TOTAL), 2) observations where the distance error is predominantly along the $X$-axis ($X$), and 3) predominantly along the $Y$-axis ($Y$), and predominantly along the $Z$-axis ($Z$). Green vertical line segments represent fused observations, red line segments indicate unfused ones. The blue bar represents the mean of the two medians. $X, Y, Z$ plots only show observations in which the position error contribution along the main axis is higher than for both other axes by 150 percent or more. For each plot, the left, right, and middle long bars represent the fused median, unfused median, and mean of fused and unfused medians.

$Z$-axes, respectively. The mean for nonfused samples along the $Z$-axis was lower than the means for the nonfused samples along the $X$ and $Y$-axes as seen in Fig. 21. These results suggest that users are more sensitive to alignment errors along the $Z$-axis. We also found confirmation to our related prediction that the same absolute tracking errors along the $Z$-axis will produce much smaller error in angular alignment than distance in $X$ or $Y$. This can be gleaned from the much shorter interval that observations cover in the $z$-biased plot of Fig. 21 (bottom row).

## 9 Conclusion

We have presented a step toward a room-sized 3D display, using the DFD effect in conjunction with the immaterial FogScreen. Our contributions are the demonstration of 3D perception via a generalized DFD technique, a prototype display system based on this technique and the FogScreen, and an evaluation of the display and its limiting factors and their influence. We showed the effectiveness of the generalized DFD technique to be equivalent to the established DFD results via a formal user study, and our testing of the prototype system confirmed the expected 3D perception in a real system. We are excited about the 3D impression users can get from the DFD principle in these new configurations. Using the FogScreens, depth is perceived well, but because of registration errors and fog turbulence, it does not currently reach the 3D fidelity of stereoscopy. In practice, we found a walk-through display is still currently out of reach because alignment errors and induced turbulence are more obvious when users are close to a screen. This can be mitigated with better tracking and closer screen placement. We listed and characterized the most important hindrance factors for perceiving a true DFD effect, and analyzed the user's tolerance regarding two of them, intensity mismatches and observer-position-induced alignment errors.

We are currently working on simulating the fog sheet screens from our physical system more accurately in our simulator using particle systems and GPU-accelerated flow simulations. This will enable the development and testing of algorithms to optimize the visual appearance of our projection. Further work includes experimenting with environmental factors to reduce fog turbulence and increase tracking accuracy. From the pilot study, we found turbulence due to other screens to be a major factor and having ground vents would remove this problem as well as improving the accuracy of optical trackers by removing fog build up. Finally, exploration of the possibilities for reach-in user interaction on our prototype display is ongoing.

## Acknowledgments

This research was supported in part by US National Science Foundation (NSF) Grant IIS-0635492, NSF IGERT Grant DGE-0221713 in Interactive Digital Multimedia, and a research contract with the Korea Institute of Science and Technology (KIST) through the Tangible Space Initiative project. Special thanks to FogScreen Inc. and WorldViz Inc. for their extensive hardware support, and to Thomas Klemmer for hardware and programming support.

## References

[1] B.G. Blundell and A.J. Schwarz, *Volumetric Three-Dimensional Display Systems.* Wiley-VCH, Mar. 2000.

[2] S. Suyama, Y. Ishigure, H. Takada, K. Nakazawa, J. Hosohata, Y. Takao, and T. Fujikao, "Apparent 3-D Image Perceived from Luminance-Modulated Two 2-D Images Displayed at Different Depths," *Vision Research,* 2004.

[3] M. Halle, "Autostereoscopic Displays and Computer Graphics," *ACM SIGGRAPH Computer Graphics,* vol. 31, no. 2, pp. 58-62, 1997.

[4] H. Takada, S. Suyama, M. Date, K. Hiruma, and K. Nakazawa, "A Compact Depth-Fused 3-D Display Using a Stack of Two LCDs," *J. Inst. Image Information and Television Engineers,* 2004.

[5] S. DiVerdi, I. Rakkolainen, T. Höllerer, and A. Olwal, "A Novel Walk-Through 3D Display," *Proc. SPIE,* vol. 6005, no. 1, 2002.

[6] I. Rakkolainen and K. Palovuori, "Walk-Thru Screen," *Proc. SPIE,* Projection Displays VIII, M.H. Wu, ed., ser. presented at the Soc. of Photo-Optical Instrumentation Engineers (SPIE) Conf., vol. 4657, pp. 17-22, Apr. 2002.

[7] G.E. Favalora, J. Napoli, D.M. Hall, R.K. Dorval, M. Giovinco, M.J. Richmond, and W.S. Chun, "100-Million-Voxel Volumetric Display," *Proc. SPIE,* D.G. Hopper, ed., vol. 4712, no. 1, pp. 300-312, 2002.

[8] Lightspace Technologies Inc., lightSpace Technologies, http://www.lightspacetech.com, Aug. 2007.

[9] E. Downing, L. Hesselink, J. Ralston, and R. Macfarlane, "A Three-Color, Solid-State, Three-Dimensional Display," *Science,* vol. 273, pp. 1185-1189, Aug. 1996.

[10] H. Kimura, T. Uchiyama, and H. Yoshikawa, "Laser Produced 3D Display in the Air," *Proc. ACM SIGGRAPH '06,* p. 20, 2006.

[11] S. Eitoku, K. Hashimoto, and T. Tanikawa, "Controllable Water Particle Display," *Proc. ACM SIGCHI Int'l Conf. Advances in Computer Entertainment Technology (ACE '06),* p. 36, 2006.

[12] T. Grossman, D. Wigdor, and R. Balakrishnan, "Multi-Finger Gestural Interaction with 3D Volumetric Displays," *ACM Trans. Graphics,* vol. 24, no. 3, p. 931, 2005.

[13] S. Pastoor and M. Wopking, "3-D Displays: A Review of Current Technologies," *Displays,* vol. 17, no. 11, pp. 100-110, 1997.

[14] S. Fisher, "Viewpoint Dependent Imaging: An Interactive Stereoscopic Display," *Proc. SPIE,* vol. 367, 1982.

[15] W.B. Paley, "Head-Tracking Stereo Display: Experiments and Applications," *Proc. SPIE,* Stereoscopic Displays and Applications III, ser. presented at the Soc. of Photo-Optical Instrumentation Engineers (SPIE) Conf., J.O. Merritt and S.S. Fisher, eds., vol. 1669, pp. 84-89, June 1992.

[16] I.E. Sutherland, "The Ultimate Display," *Proc. Int'l Federation for Information Processing (IFIP) '65,* pp. 506-508, 1965.

[17] M. Mon-Williams, J. Wann, and S. Rushton, "Binocular Vision in a Virtual World: Visual Deficits following the Wearing of a Head-Mounted Display," *Ophthalmic and Physiological Optics,* 1993.

[18] J. Wann, S. Rushton, and M. Mon-Williams, "Natural Problems for Stereoscopic Depth Perception in Virtual Environments," *Vision Research,* 1995.

[19] Y.I.S. Suyama, H. Takada, K. Nakazawa, J. Hosohata, Y. Takao, and T. Fujikado, "Evaluation of Visual Fatigue Relative in the Viewing of a Depth-Fused 3D Display and 2D Display," *Proc. Int'l Display Workshops (IDW),* 2004.

[20] K. Akeley, S.J. Watt, A.R. Girshick, and M.S. Banks, "A Stereo Display Prototype with Multiple Focal Distances," *ACM Trans. Graphics,* vol. 23, no. 3, pp. 804-813, 2004.

[21] K. Uehira, "3-D Display System Using Two Stereoscopic Displays at Different Depths," *Proc. 25th Int'l Display Research Conf. (Eurodisplay),* 2005.

[22] R.D.C. Team, *R: A Language and Environment for Statistical Computing,* R Foundation for Statistical Computing, ISBN 3-900051-07-0, http://www.R-project.org, 2006.

[23] *Precision Position Tracker,* worldViz, http://www.worldviz.com/, Sept. 2006.

**Cha Lee** received the BS degree in computer science from California State University Chico. He is currently working toward the PhD degree in computer science in the Department of Computer Science, University of California at Santa Barbara. He currently works at WorldViz LLC. His research interests include novel display systems. He is a student member of the IEEE.

**Stephen DiVerdi** received the PhD degree in computer science at the University of California at Santa Barbara, in 2007, where he was a member of the "Four Eyes" lab. He is a research scientist at Adobe Systems Inc. His work is focused on applying computer graphics and computer vision to create new modes of interaction between humans and computers.

**Tobias Höllerer** received the degree in informatics from the Technical University of Berlin and the MS and PhD degrees in computer science from Columbia University. He is an associate professor of computer science in the Department of Computer Science, University of California at Santa Barbara, where he leads the "Four Eyes" research group, conducting research in the "Four I's" of Imaging, Interaction, and Innovative Interfaces. He is a member of the IEEE and the IEEE Computer Society.