

An Immaterial Depth-Fused 3D Display

Cha Lee*

Stephen DiVerdi†

Tobias Höllerer‡

Four Eyes Laboratory
Computer Science Department
University of California, Santa Barbara
Santa Barbara, CA 93106

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 depth-fused 3D principle by using screens in arbitrary configurations and from arbitrary viewpoints. The performance of the generalized DFD effect is established with a user study using the simulator. Based on these results, we developed a prototype display using two immaterial screens to create an unencumbered 3D visual that users can penetrate, enabling the potential for direct walk-through and reach-through manipulation of the 3D scene.

CR Categories: I3.1 [Computer Graphics]: Hardware Architecture—Three-Dimension Displays; H5.2 [Information Interfaces and Presentation]: User Interfaces—Evaluation / Methodology;

Keywords: 3D displays, immaterial displays, depth-fused 3D, user study

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 [Blundell and Schwarz 2000]. 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 towards attaining this ideal.

An interesting unencumbering pseudo-3D display technique is called depth-fused 3D (DFD) by Suyama et al. [Suyama et al. 2004]. 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 lu-



Figure 1: Our prototype immaterial depth-fused 3D display using two FogScreens in an L-shaped configuration, showing a 3D teapot.

minous 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 [Halle 1997]. Today there are desktop sized [Suyama et al. 2004] and handheld size DFD displays [Takada et al. 2004]. 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 [DiVerdi et al. 2006][Rakkolainen and Palovuori 2002]. This screen is a 2.5 x 1.5 meter 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.

Our contributions from this work are twofold. First, we simulate and evaluate a generalized DFD display, and second, we use the generalized DFD technique to develop a prototype immaterial display using FogScreens. 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

*e-mail: chalee21@cs.ucsb.edu

†e-mail: sdiverdi@cs.ucsb.edu

‡e-mail: holl@cs.ucsb.edu

Copyright © 2007 by the Association for Computing Machinery, Inc.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions Dept, ACM Inc., fax +1 (212) 869-0481 or e-mail permissions@acm.org.

VRST 2007, Newport Beach, California, November 5–7, 2007.

© 2007 ACM 978-1-59593-863-3/07/0011 \$5.00

our prototype in two configurations and discuss the results. The advantage of a general immaterial DFD display using FogScreens is that there is the potential for an observer to directly interact with the 3D environment, unencumbered by glasses or headmounted displays. 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, and in Section 6 we discuss the results and implications of this work.

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. [Favalora et al. 2002] project light onto a rapidly spinning screen, carefully timing the projection to illuminate individual voxels. Alternately, Lightspace Technology's DepthCube [Lightspace 2007] projects onto a stack of parallel LCD shutters. More exotic concepts such as Downing et al.'s [Downing et al. 1996] employ infrared lasers to excite points in a rare-earth doped gas, while the lasers in Kimura et al.'s display [Kimura et al. 2006] 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.'s display [Eitoku et al. 2006]. 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 [Grossman et al. 2005]. 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 [Rakkolainen and Palovuori 2005].

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 [Pastoor and Wopking 1997], 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. Autostereoscopic displays [Halle 1997] 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 [Fisher 1982][Paley 1992] 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 [Sutherland 1965] for immersive perception of a 3D scene. Unfortunately, stereo techniques are subject to user fatigue during extended viewing from inaccuracies in the effect [Mon-Williams et al. 1993][Wann et al. 1995] and the encumbrance of glasses.

More recently, an effect called depth-fused 3D (DFD) has been investigated [Suyama et al. 2004] as another technique for simulating




Figure 2: *The Depth-Fused 3D Effect on an L-shaped configuration.*

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 [S. Suyama et al. 2004] and doesn't require any user-worn glasses. This technique has been used for a prototype compact display [Takada et al. 2004], and the interaction between DFD and stereo imaging has been explored [Akeley et al. 2004][Uehira 2005], 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

Figure 2 illustrates the DFD effect on an L-shaped screen configuration. The intensity of pixels on the different screens are chosen to reflect the respective distances of the visible surface points projecting onto them.

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 [Lightspace 2007]. 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 x 0.5 x 0.5 meters in size, with the fog sheet reaching a thickness (depth) of 2 to 8cm, sandwiched in between even thicker sheets of regulating airflow. Airflow interference causes turbulences when another unit is placed alongside of it. This alone imposes a minimum stacking distance of about 1m. 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




Figure 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. The images are left-right reversed for cross-fused stereo viewing.

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 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. Figure 3 and 4 show example stereo images of the DFD effect (Figure 3) and plain 3D stereo (Figure 4). Unlike the image in Figure 3, the DFD images presented to the study participants did not have the semi-transparent screens displayed. The reader of this paper is encouraged to cross their eyes on these figures to experience the DFD effect vs. 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.




Figure 4: Stereo pair for true binocular stereo. Here, the bunny is a 3D model seen by left and right eye. The images are left-right reversed for cross-fused stereo viewing.

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 was reported to not be sufficient to bring out DFD depth impression [Suyama et al. 2004]. On the other hand, we also know that it significantly helps depth impression, when accommodation is in sync with convergence and disparity [Akeley et al. 2004]. We conducted our user study with the simulator in the hope that we would see a significant effect of the generalized DFD conditions on 3D perception, even in absence of correct accommodation. The results in Section 4 indicate that this is the case. Accommodation does play a role in our physical prototype realization of the general DFD display, and anecdotal evidence discussed in Section 6 indicates that it may actually improve the 3D impression for at least some users.

The simulator version used in this work represents screens as simple semi-transparent 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 [Fisher 1982][Paley 1992]. 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 [Suyama et al. 2004], 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 Figure 2. These rendered images are stored to offscreen buffers.

In the final rendering step, we define a normal stereo camera at the user's position, map our rendered images to our transparent screens,



Figure 5: Different Screen Configurations: Cross (discarded), L-Shape, Stack, Triangle




Figure 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, f) unblended off-axis stack (for control purpose). Stereo condition is not shown here. The screens are for clarification purposes only and did not appear in the study conditions.

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 either through a stereo HMD or using an active stereo projector and shutter glasses.

We experimented with a variety of configurations: stacked, crossed, L shaped, and triangle (see Figure 5). The cross configuration was discarded from closer consideration after initial experiments because it was evident that it would perform poorly because of the fact that there is effectively only one transparent screen at the center of the scene. As a result, it creates the smallest 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 Study Design

Our study consisted of sixteen subjects, five female and eleven male, ranging in age between 22-26, 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' white projection screen. Users were instructed to stand on a line approximately 8' 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 sixteen users we began with, one was unable to perceive stereo.

For the remaining users, we displayed a series of different static images (see Figure 6) and asked them to rate how 3D the depicted object appeared on a scale of zero to five, zero being completely flat, and five being completely 3D. We also encouraged users to give feedback on what they perceived. The images users evaluated each showed the Stanford bunny 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 five 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 Figure 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 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 non-overlapping 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

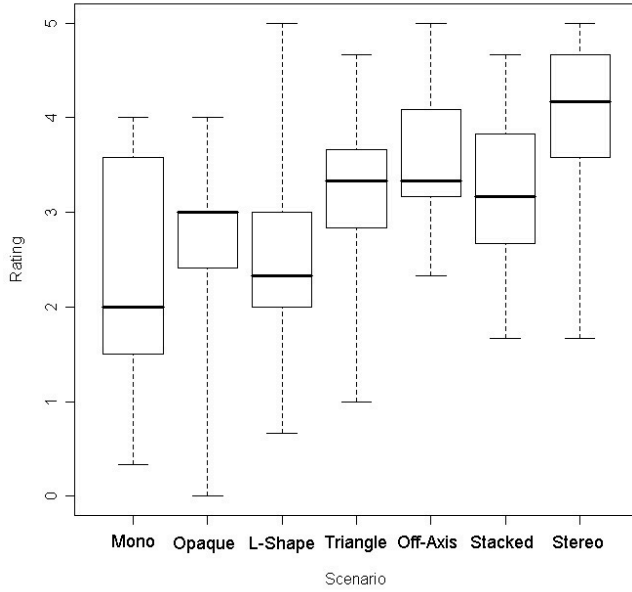


Figure 7: Boxplot of users ratings for each scenario. Each column shows the 0th, 25th, 50th, 75th and 100th percentiles.

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 of the user's ratings versus seven scenario treatments showed a strong statistical significance among the results ($F(6,84) = 12.791$, $p < 0.001$). Figure 7 shows the aggregated ratings for each scenario. We also did a post hoc analysis using Tukey's Multiple Comparisons of Means, for which the results are shown in Table 1. All statistical analysis was performed with the statistical computing environment *R* [R 2007].

Stereo is the highest 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*, *off-axis*, and *stereo* ($p < 0.01682$, $p < 0.00634$, $p < 0.00005$, $p < 0.0000001$ respectively). Finally, *L-shape* and *opaque* were not significantly different from mono ($p < 0.99935$ and $p < 0.75602$ respectively).

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

Table 1: Tukey Multiple Comparisons of Means with 95% family-wise confidence level. The P Adj. column lists the p-value after adjustment for the multiple comparisons. Statistically significant differences are highlighted in bold face.

We expected *stereo* to be rated the highest, and these 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, subjective 3D impression of content presented in a DFD fashion may not be perceived as less three-dimensional than traditional stereo rendering. It is also reassuring to see that the *mono* ratings exhibited the lowest median, even though the high variance did prevent a statistically significant difference to *opaque* and *L-shape*. Some users liked the 3D appearance of 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 3D perception.

The rating of *stack* confirms established results on the DFD effect [Suyama et al. 2004], 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 reinforces our belief that the DFD effect will work well in conjunction with head-tracked rendering for 3D perception from multiple viewpoints. More tests are needed to back up this intuition.

The ratings for *L-shape* are important to consider. The main 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 shows that stacked, opaque displays are




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

not sufficient by themselves to create a 3D perception, and suggests that the DFD technique is critically important to a high quality visual in multi-screen displays.

5 Prototype of a Walk-through 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, so 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 (cf. Figure 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. The L-shape configuration also benefits from the fact that it only requires two FogScreens and does not suffer from any bleed-through effects.

The system we assembled uses two FogScreens, each with their own standard DLP projector, in the stacked and L-shaped configurations. For head-tracking, we use WorldViz's Precision Position Tracker [WorldViz 2006], 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 Figures 8 and 10). Its implementation in our prototype is hindered by the limitations of the FogScreen. Because the




Figure 9: System overview for L-shaped Dual FogScreen setup. No bleed-through, but limited screen overlap.

FogScreen transmits most of the light projected on to it, the screens cannot be mounted too close together, or the image from the rear screen will bleed through and 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 non-zero 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 3m apart (see Figures 8 and 10).

In the L-shaped configuration, the two screens are mounted to form a right angle (see Figures 9 and 10). 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 image does not bleed on to the front image, as occurs in the stacked configuration.

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 the user centered and perpendicular to the screens, the DFD effect was clear, resulting in 3D perception as reported previously [Suyama et al. 2004]. Second, as the user moves around the display, the head-tracked DFD rendering maintains the 3D perception, 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 create the DFD effect.

6 Discussion

The results of our prototype display bear further consideration. In the standard DFD effect, both binocular disparity and ocular accommodation provide strong depth cues. However, in testing our prototype, two users who are not sensitive to binocular disparity and thus are unable to perceive stereo imagery, reported an enhanced 3D effect over a regular 2D image with our physical prototype. This might imply that the ocular accommodation depth cue in a DFD display is strong enough to enable some 3D perception, despite previous claims to the contrary [Suyama et al. 2004]. We




Figure 10: Photographs of physical FogScreen setups with screen areas indicated by outlines.

hypothesize that correct accommodation will only strengthen the results from our user study conducted with the simulator (cf. Section 4), but more formal studies on the interplay of accommodation and vergence are needed.

While the FogScreen's immaterial nature is very appealing for a generalized DFD display, its limitations must be addressed. The fog composition of the screens' projection surface creates distortions from turbulence in the fog flow. This distortion is localized to small areas, so only impacts the display of small objects and fine details. It can generally be ignored for large objects. As the quality of the FogScreen improves, this artifact will be lessened, automatically improving the fidelity of the DFD perception.

The bleed-through effect of the FogScreen also needs to be dealt with. As we demonstrate in our prototype, clever configurations of screens can eliminate the problem, though ideally we would like to be able to experiment with more diverse layouts without this artifact. As the quality of the FogScreen improves, thinner, more stable fog will enable the use of short-throw projectors without sacrificing image quality. This will enable more diverse configurations of real screens for other interesting effects.

Our prototype system does not depend on the FogScreen, but rather demonstrates the concept. However, we currently see the FogScreen or Heliodydisplay [IO2 2007] as the only viable options for actually traversing the screens. The benefit is not only direct "reach-in" interaction (which suffers from calibration errors and pixelation errors when the observer is close to a screen) but also chiefly the ability to walk around in a large 3D image without the need for stereo glasses. If three or more screens were used, the observer could actually stand inside the scene and still perceive the DFD effect. Other immaterial displays could take the place of the FogScreen in the future. Truly volumetric reach- and walk-through displays are not yet possible, but our explorations clearly show progress in that direction. When such technologies first become available, improved DFD effects such as the ones tested here may very well be used with them in combination, for example as a 3D backdrop to a volumetric object in the foreground.

As with any head-tracked display, proper registration of displays and tracker coordinates is a challenge. The overlapped image requirement of the DFD effect increases the importance of proper calibration, which in practice can be difficult to achieve. Tracker errors, calibration errors, and fog turbulence all can cause the images to not overlap perfectly. In our implementation, we found that when an observer was close to a screen small calibration errors be-



Figure 11: Teapot in DFD rendering on L-shaped Fogscreen setup: Two projectors create overlapping image fused in observers brain.

came very evident. Reducing ambient airflow in the environment makes the FogScreen surface more stable, improving registration. Tracker error, in particular lag, can be addressed by using a predictive filter. In general however, exact registration may not be needed - as Suyama stated [Suyama et al. 2004], some small errors in registration produce edge artifacts, but a perception of depth still occurs if these artifacts are small in comparison to the size of the virtual object.

7 Conclusions

We have presented a step towards a room sized walk-through 3D display, using the DFD effect in conjunction with the FogScreen. Our contributions are the demonstration of 3D perception via a generalized DFD technique, and a prototype display system based on this technique and the FogScreen to create an immaterial, unencumbered 3D perception that enables natural interaction through direct manipulation. We showed the effectiveness of the generalized DFD technique to be equivalent to the established DFD results via a formal user study, and our informal testing of the prototype system confirmed the expected 3D perception in a real system. We are enthusiastic about the 3D impression users can get from the DFD principle in these new configurations. For several users it was indistinguishable from true stereo in the simulator. Using the FogScreens, depth is perceived well, but because of registration errors, it does not currently reach the 3D fidelity of stereoscopy.

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 (through online pre-distortion). Further future work includes developing and testing additional screen configurations, both in the simulator and in the real world, including a larger number of stacked screens and square or circular surround-screens, even in multiple layers, with the goal of eventually creating a fully immersive version of our display, which would generate surround-view style visuals without user-worn displays or glasses. Finally, exploration of the possibilities for reach-in user interaction on our prototype display is ongoing.

Acknowledgements

This research was supported in part by NSF grant IIS-0635492, NSF IGERT grant DGE-0221713 in Interactive Digital Multime-

dia, 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

- AKELEY, K., WATT, S. J., GIRSHICK, A. R., AND BANKS, M. S. 2004. A stereo display prototype with multiple focal distances. *ACM Trans. Graph* 23, 3, 804–813.
- BLUNDELL, B. G., AND SCHWARZ, A. J. 2000. *Volumetric Three-Dimensional Display Systems*. Wiley-VCH, Mar.
- DIVERDI, S., RAKKOLAINEN, I., HOLLERER, T., , AND OLWAL, A. 2006. A novel walk-through 3d display. *Proceedings of SPIE* 6005, 1.
- DOWNING, E., HESSELINK, L., RALSTON, J., AND MACFARLANE, R. 1996. A Three-Color, Solid-State, Three-Dimensional Display. *Science* 273 (Aug.), 1185–1189.
- EITOKU, S., HASHIMOTO, K., AND TANIKAWA, T. 2006. Controllable water particle display. In *ACE '06: Proceedings of the 2006 ACM SIGCHI international conference on Advances in computer entertainment technology*, ACM Press, New York, NY, USA, 36.
- FAVALORA, G. E., NAPOLI, J., HALL, D. M., DORVAL, R. K., GIOVINCO, M., RICHMOND, M. J., AND CHUN, W. S. 2002. 100-million-voxel volumetric display. *Proceedings of SPIE: Cockpit Displays IX: Displays for Defense Applications* 4712, 1, 300–312.
- FISHER, S. 1982. Viewpoint dependent imaging: an interactive stereoscopic display. *Proceedings of SPIE: Processing and Display of Three-Dimensional Data* 367.
- GROSSMAN, T., WIGDOR, D., AND BALAKRISHNAN, R. 2005. Multi-finger gestural interaction with 3d volumetric displays. *ACM Trans. Graph.* 24, 3, 931–931.
- HALLE, M. 1997. Autostereoscopic displays and computer graphics. *SIGGRAPH Comput. Graph.* 31, 2, 58–62.
- IO2, 2007. IO2 Technology USA. <http://www.io2technology.com/>, accessed May 2007.
- KIMURA, H., UCHIYAMA, T., AND YOSHIKAWA, H. 2006. Laser produced 3d display in the air. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Emerging technologies*, ACM Press, New York, NY, USA, 20.
- LIGHTSPACE, 2007. Depthcube. LightSpace Technologies Inc., <http://www.lightspacetech.com/>, accessed May 2007.
- MON-WILLIAMS, M., WANN, J., AND RUSHTON, S. 1993. Binocular vision in a virtual world: Visual deficits following the wearing of a head-mounted display. *Ophthalmic and Physiological Optics*.
- PALEY, W. B. 1992. Head-tracking stereo display: experiments and applications. In *Proc. SPIE Vol. 1669, p. 84-89, Stereoscopic Displays and Applications III*, John O. Merritt; Scott S. Fisher; Eds., J. O. Merritt and S. S. Fisher, Eds., vol. 1669 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*, 84–89.
- PASTOOR, S., AND WOPKING, M. 1997. 3-d displays: A review of current technologies. *Displays* 17 (Apr.), 100–110.
- R, 2007. R: A language and environment for statistical computing. The R Project for Statistical Computing, <http://www.r-project.org/>, accessed May 2007.
- RAKKOLAINEN, I., AND PALOVUORI, K. 2002. Walk-thru screen. In *Proc. SPIE Vol. 4657, p. 17-22, Projection Displays VIII*, Ming H. Wu; Ed., M. H. Wu, Ed., vol. 4657 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*, 17–22.
- RAKKOLAINEN, I., AND PALOVUORI, K. 2005. Laser scanning for the interactive walk-through fogscreen. In *VRST '05: Proceedings of the ACM symposium on Virtual reality software and technology*, ACM Press, New York, NY, USA, 224–226.
- S. SUYAMA, Y. I., TAKADA, H., NAKAZAWA, K., HOSOHATA, J., TAKAO, Y., , AND FUJIKADO, T. 2004. Evaluation of visual fatigue relative in the viewing of a depth-fused 3D display and 2D display. In *Proceedings of the International Display Workshops*.
- SUTHERLAND, I. E. 1965. The ultimate display. In *Proceedings of IFIP*, 506–508.
- SUYAMA, S., ISHIGURE, Y., TAKADA, H., NAKAZAWA, K., HOSOHATA, J., TAKAO, Y., , AND FUJIKAO, T. 2004. Apparent 3-D image perceived from luminance-modulated two 2-D images displayed at different depths. *Vision Research*.
- TAKADA, H., SUYAMA, S., DATE, M., HIRUMA, K., , AND NAKAZAWA, K. 2004. A compact depth-fused 3-D display using a stack of two lcds. *The Journal of The Institute of Image Information and Television Engineers*.
- UEHIRA, K. 2005. 3-d display system using two stereoscopic displays at different depths. In *Proceedings of the International Display Research Conference*.
- WANN, J., RUSHTON, S., AND MON-WILLIAMS, M. 1995. Natural problems for stereoscopic depth perception in virtual environments. *Vision Research*.
- WORLDVIZ, 2006. Precision position tracker. WorldViz, <http://www.worldviz.com/>, accessed May 2007.