3D Visualisation Techniques for Multi-Layer Display™ Technology

Vijay Prema[1], Gary Roberts[1] and Burkhard Wünsche[2]

[1] Dept. Electrical & Electronic Engineering, University of Auckland, New Zealand.
[2]Department of Computer Science, University of Auckland, New Zealand.
Email: burkhard@cs.auckland.ac.nz

Abstract

Traditional computer monitors offer limited depth perception due to their 2D nature. The multi-layer display technology uses two or more display layers stacked in parallel and separated physically by depth. When viewing a Multi-Layer Display (MLD™) objects displayed on the front layer appear closer than objects on the back layer, and when moving the head while viewing the display objects on the front and back layer move relative to each other. However, it is not clear how complex 3D scenes can be rendered effectively using two physically separated view planes. We have experimentally analysed differences in perception when using single and Multi-Layer Displays and used the results to develop novel rendering techniques for MLD™. We found that perception of scenes can be improved by emphasizing important objects by displaying them on a different layer, by separating datasets on different layers, by extruding objects across layers, by transitioning objects smoothly between layers and by making use of the transparency of the front layer. As a result of our user studies we present a set of guidelines for the most effective use of Multi-Layer Display technology for rendering 3D scenes.

Keywords: Multi-layer displays, 3D displays, visual perception, human-computer interfaces, visualization

1 Introduction

Consumer level display technology has advanced dramatically with the advent of plasma and LCD displays. One important feature that has yet to reach the mainstream consumer is real 3D depth in images. There are many display technologies which can achieve varying levels of 3D depth, however most are expensive, inconvenient or have depth limitations. The Multi-Layer Display (MLD™) developed by PureDepth functions similarly to a conventional LCD monitor except that it features a second screen directly behind the transparent front screen. Images can be rendered on either of these two layers which are separated by a small space, conveying a limited amount of 3D depth [1].

3D display technology has a wide range of applications in entertainment, advertising, medicine, military and other fields. Examples include animated signs, video games, television, heads-up-displays and design visualizations [2]. 3D depth in display technology allows images to be interpreted faster, with more clarity and more realism. The MLD™ is one of the most accessible depth limited displays because of its compact size, low cost and compatibility with common PCs. Traditionally its dual layers are used as discrete surfaces for displaying overlaid information and highlighting objects by making them appear physically closer to the user.

In this paper we present a number of novel rendering techniques for harnessing the power of MLD™ technology and creating more effective visualizations. Section 2 summarises results about human depth perception. Section 3 introduces 3D display technologies and the PureDepth MLD™ technology. In section 4 we analyse differences in perception of single- and multi-layer displays and use the results in section 5 to develop more effective rendering techniques for a MLD™. Section 6 presents the summary of results obtained by performing user testing for our novel rendering methods. In section 7 we draw conclusions to our research and suggest directions for further studies.

2 Depth Perception

Human beings perceive depth using a combination of depth cues. *Psychological depth cues* are attributes of a physically flat image which are interpreted by the brain as 3D distance information and are hence extensively used when rendering 3D scenes on conventional 2D displays. *Linear perspective* is the recognition of parallel lines converging towards a point in the distance. Closer objects appear larger than distant ones and the known sizes of recognized objects can also be recalled from memory in order to make accurate distance estimations. *Occlusion* is the overlapping of objects and gives some idea of the order of objects in a scene. Depending on light sources, *shadows and shading* can provide clues as to where objects are with respect to the ground plane. *Atmospheric perspective* is the blurring and blue tinting of distance objects due to scattering light. Some psychological depth cues involve motion. These include *motion parallax*, where nearer objects move faster than further ones; and *optic flow*, where the scene seems to expand from the point that the camera is moving towards [3][4].

Physical depth cues rely on the fact that humans have two eyes, and cannot be utilized by ordinary 2D displays. *Binocular disparity* is the main physical depth cue which involves the brain processing the images from both eyes. Since the eyes are some distance apart they capture slightly different images with a large overlap. The differences in the overlapping region can be perceived as 3D depth. *Vergence* is the movement of both eyes in opposite directions as they focus on an object which is moving towards the viewer. This can be used by the brain to very accurately judge distance [3][4].

The brain uses a weighted combination of all depth cues to perceive 3D depth. Physical cues are weighted more heavily at closer distances and psychological cues (particularly motion based cues) at long distances [3]. Gestalt psychology states that an important part of visual perception involves grouping parts of geometry in a scene into recognizable objects, e.g. by similarity, continuation, proximity, and common fate. Gestalt does not refer to depth perception in particular but we utilise the brains ability to perceive Gestalt when making objects appear to be continuous across both layers of the display [5].

3 3D Display Technologies

Artists have exploited size, shape, overlay, linear perspective and shadows to add depth to an image [6]. 3D displays are designed to utilise as many of the depth cues covered in section 2 as possible [7]. 3D capable technologies include anaglyphs, stereoscopic displays and autostereoscopic displays which use goggles or other tools to generate different images for both eyes [8]. All of these techniques suffer from user discomfort and eyestrain. A hologram records the intensity and the phase of the wavefront emanating from an objects surface but at present the images are fixed in film and cannot be manipulated. Volumetric displays illuminate points in 3D space but are very expensive [8].

Multi-Layer Displays do not have the same issues with discomfort, are smaller than other displays, can be easily installed on most computers (they require a dual head graphics card), and are cheaper than most alternatives. A Multi-Layer Display blends the colours of pixels rendered on the front and back layer together. This means if a dark pixel is rendered on the back layer, then the corresponding pixel on the front layer will also be dark. What is rendered on the front layer must therefore take into account the colour of the scene behind it.

In our research we use a 17 inch MLD™ prototype, which consists of 2 LCD layers, with the back layer 7 mm behind the front layer. The display is connected to the computer via a dual head graphics card. The resolution of the screen is set to 2560x1024. The first 1280 pixels correspond to pixels on the front layer, the rest are for the back layer. We use OpenGL for rendering because it is platform independent, easily portable, offers fast real-time 3D graphics, has a stencil buffer for rendering silhouettes, and includes a shading language for implementing per-pixel operations. In OpenGL an easy way to render on the MLD™ is to create 2 viewports, one for the front layer and the other for the back, and render in each viewport separately.

4 Perceptional Differences for MLD™

We have performed and analysed a series of experiments in order to better understand how perception of the MLD™ varies from that of single layer displays (SLDs). Details of the experimental set-up and results are described in [9,10]. We found that users sitting within 0.5m from the screen in most cases could determine what was on the front layer and what was on the back layer. Reference objects helped with this which indicates that binocular disparity is an important depth cue in the MLD™. Similarly being able to move the head improved perception when using reference objects (motion parallax). Performance was further improved when the objects were overlapping.

5 Rendering Techniques on MLD™

We have developed various techniques to improve depth perception when rendering 3D scenes on a MLD™. The following subsections introduce these techniques and discuss their advantages, disadvantages and limitations.

5.1 Emphasising objects by putting them on a different layer

Since depth is more powerful than colour to help find an object [11], objects can be emphasized by putting them on a different layer, usually the front layer. Care must be taken when choosing colours for the emphasised object and the background scene. We found that the technique works best if the background has light colours, and the foreground has dark colours. If the background is dark then foreground objects are hard to see and if the foreground object is light it appears transparent (because of the physical makeup of the front layer) and the background shines through. The first problem can be alleviated by rendering a white silhouette of the foreground object onto the back layer. We achieve this in OpenGL by drawing the background into the stencil buffer, then drawing the foreground objects in white where stencil values are non-zero, and finally drawing the foreground objects onto the front layer.

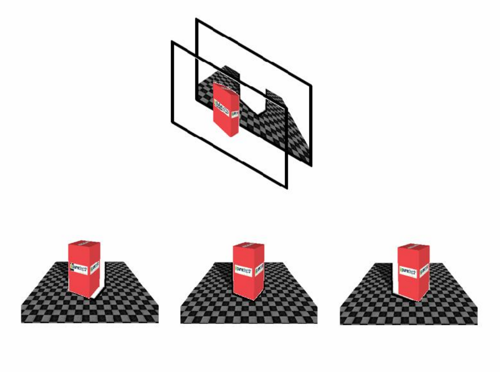

Figure 1: Emphasizing important objects.

We tested the scene displayed in figure 1 and found that most users perceived the red object as more accentuated when using the above described technique. One problem is that the white silhouette becomes visible when the user moves the head. This can be alleviated by fading the silhouette similar to the technique explained in the next subsection.

5.2 Determining layers by object depth value

The Z-value technique splits the entire scene by its Z-value (depth buffer value) and renders each half on a separate layer. All parts of the scene with a Z value greater than a certain threshold distance are rendered on the back layer and everything else on the front layer. An example is shown in the top row of figure 2, which shows a scene consisting of a rotating cube suspended in space and casting a shadow onto a platform below it. As the camera moves towards an object which is on the back layer, the object will eventually cross the threshold distance and gradually move to the front layer. The faces of any 3D object which intersect the threshold plane will be cut accordingly and the object will be partially rendered on both layers.

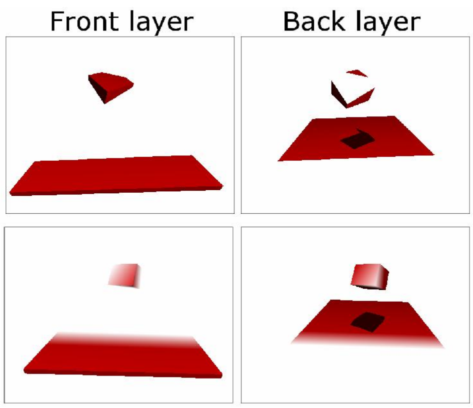

Figure 2: The Z-value technique using hard edges (top) and continuous shading (bottom).

The initial implementation of this technique was not very effective because of the discontinuity in the image caused by objects crossing the Z threshold. The main problem is that objects cut by the threshold plane appear to be unnaturally discontinuous or overlapping, particularly when the viewer moves the head.

The Z-value technique was dramatically improved by using *continuous shading* rather than discretely splitting a scene and rendering each half on a separate layer. In this implementation, each pixel in the scene is rendered with an independent alpha which depends on its Z-value and four other constants. These constants are *fMinZ*, *fMaxZ*, *bMinZ* and *bMaxZ*. The alpha of each pixel is calculated using equation (1) for the front layer and equation (2) for the back layer.

$$\mathrm{fAlpha}(z) = \begin{cases} 1 & z < fMinZ \\ 0 & z \geq fMaxZ \\ 1 - \dfrac{fMinZ - z}{fMinZ - fMaxZ} & otherwise \end{cases} \quad (1)$$

$$\mathrm{bAlpha}(z) = \begin{cases} 0 & z < bMinZ \\ 1 & z \geq bMaxZ \\ \dfrac{bMinZ - z}{bMinZ - bMaxZ} & otherwise \end{cases} \quad (2)$$

The constants can be adjusted in order to provide enough overlap of the shading to naturally blend the layers. The demo application uses a custom OpenGL fragment shader to adjust the alpha for each pixel in real time [12]. The default values used are *fMinZ* = 0.91, *fMaxZ* = 0.96, *bMinZ* = 0.88 and *bMaxZ* = 0.93. Figure 2 shows a comparison of using hard edges on silhouettes compared to continuous smooth shading. Although this technique can be applied to any 3D scene with Z-values available, the amount of depth added to the scene is limited, as the distance between the layers is small. Another limitation is that the continuous smooth shading is less effective when rendering more complex objects with detailed surfaces (textures), particularly if the gradients overlap significantly.

5.3 Gradients

Simple view plane aligned static objects can be rendered effectively by splitting them and rendering the outer part on the front layer and the inner part on the back layer and fading the parts at the contour where they were split using the OpenGL smooth shade model.

We tested this technique by showing the scenes displayed in figure 3 and figure 4 to users. The three images at the bottom of each figure show the perceived scene from a view point to the left, in front and to the right of the monitor. The viewers were able to tell when the scene was rendered on only one layer and all users agreed that using two layers improved depth perception. Further user studies showed that the technique is most effective when the width of ring object is small (the size of the two rings are around the same size) and the length of the gradient is a long. When the width of the ring is large, viewers can't see any difference from the equivalent single layer technique. The technique is only effective when the ring appears facing up. The most suitable background colour for ring area (which determines the colour of the highlight) is white or a colour lighter then the colour of the ring. When it's dark it makes the part rendered on the front layer hard to see. Possible explanations for these observations are that the ring appears less flat since when is rendered on both layers and that the whitish region where the rendered parts overlap moves as a viewer moves their head, which is consistent with how a specular reflection on a ring would behave when being viewed.

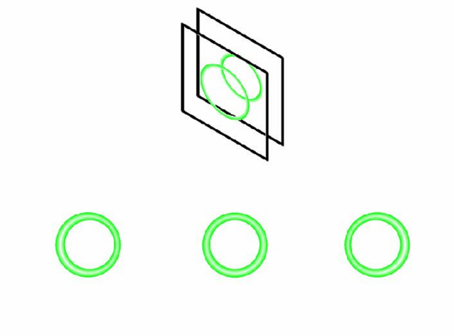

Figure 3: Effective use of gradients for ring objects.

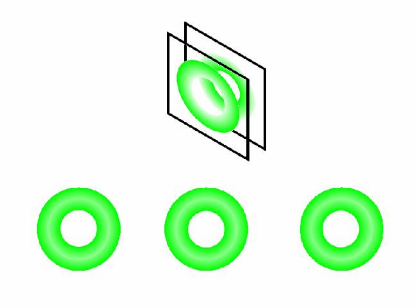

Figure 4: Ineffective use of gradient.

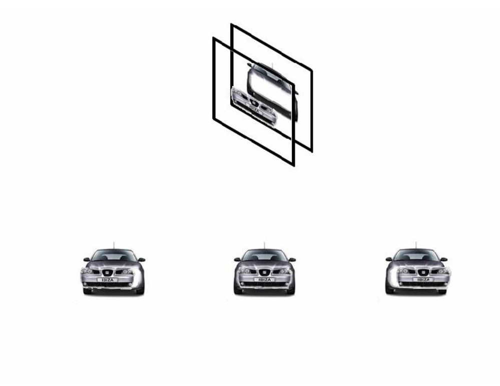

Figure 5: Ineffective use of gradient on other objects.

We applied this technique to other objects such as the car in figure 5, but found it to be ineffective. The most likely explanation is that the depth of the object is too large, i.e. when the viewer moves the head the headlights move independently of the rest of the car, which is unnatural. In addition the shape and behaviour of the whitish region where the scene components meet is inconsistent with that of a specular highlight.

5.4 Transitioning between layers

Objects can be made to appear between the display layers by rendering a percentage of the object on each layer. The object appears closer to the layer which has the higher percentage of the object rendered on it. We implemented this technique in OpenGL using alpha blending.

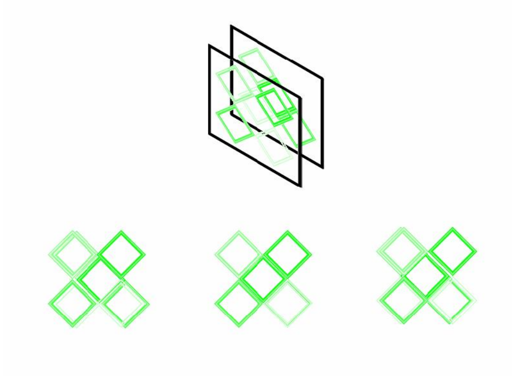

Figure 6: Transitioning objects between layers.

Two versions of the scene displayed in figure 6 were shown to users who were asked to order the squares by depth away from them. For the first version five out of six viewers agreed on the expected ordering, with the last viewer disagreeing in 3 positions. For the second version 3 out of six agreed on the expected ordering, with 2 disagreeing on 2 positions and one disagreeing on 3 positions. Note that is possible, that viewers use the apparent size of the gap between objects displayed on both layers to order the squares. However, overall the technique is effective for moving objects between layers. Possible problems are that objects change their perceived colour when moving between layers.

5.5 Calibrating objects to be viewed from a particular position

One way to render on the MLD™ is to assume that the viewer will only view the scene from a particular angle and to render the scene from that viewpoint such that both layers show the correct projection. However, it is very hard for the user to keep the head completely fixed and since we would loose depth perception due to motion parallax this technique does not seem suitable.

5.6 Grey scale depth map to determine layer

One of the most promising techniques we developed utilizes two images: the image to be displayed and a greyscale depth map which dictates how to render the image on each layer. A white and black pixel in the depth map results in the corresponding image pixel to be displayed on the front and back layer, respectively. For gray scale pixels we blend the images between the layers by using the gray scale value as the weighting factor for alpha blending. Viewers found this technique to be very effective for an image of a brick wall (figure 7). The bricks appear closer and the grouting appears to recede behind the bricks.

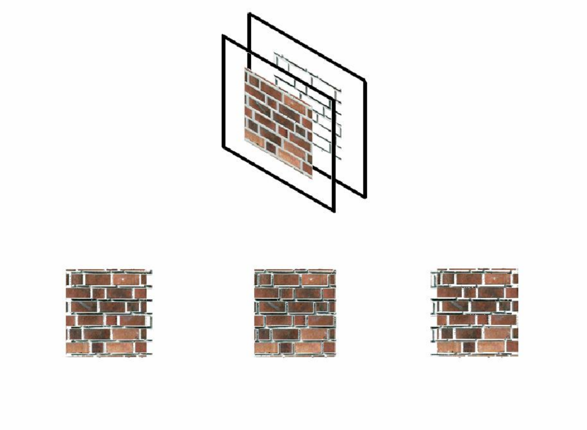

Figure 7: Depth Map to determine layer.

5.7 Visualising two data sets simultaneously separated by physical depth

Two data sets with matching domain (independent variable) can be effectively visualised on the MLD™ by displaying the data sets on a different layers on top of each other as illustrated in figure 8. The main advantage over a single layer display is that both datasets are physically close on the display which makes it easy to compare values for the same point in the domain. When points on one dataset block points on the other set the user can simply move the head to see the missing points.

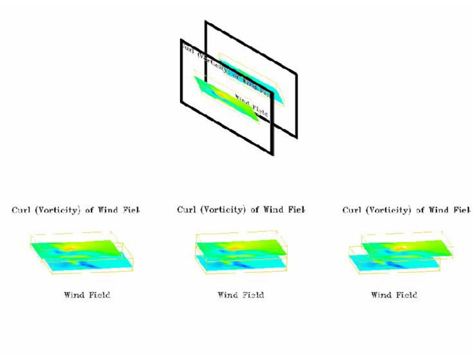

Figure 1: Comparing two datasets.

5.8 Moving objects on two layers with different speeds

Moving the scene on the front layer at a different rate to the scene displayed on the back layer gives the impression of a moving camera. An example is to have moving stars on the back layer and a stationary spacecraft on the front layer. However, user testing indicates that the technique is equally effective for single layer displays.

5.9 Transparency

The transparent front layer can be utilised to display semi-transparent materials such as glass, water and fog. The scene rendered on the back layer then appears to be physically behind the semi-transparent material. This is in particularly the case if the semi-transparent material is textured, e.g. slight waves, in which case motion parallax enhances depth perception. While this technique seems promising we did not have time to explore it in more detail.

6 Results

6.1 Analysis of Experimental Results

The perception of a MLD™ differs from a SLD in two ways. Firstly what is displayed on the front layer appears closer and separated from what is displayed on the back layer; this is due to binocular disparity. Secondly what is displayed on the front layer moves relative to what is displayed on the back layer when a viewer moves their head; this is due to motion parallax. Techniques that utilize either or both of these two properties to their advantage are more effective on the MLD™ than for a SLD.

For example, visualising two data sets simultaneously (figure 8) effectively makes use of both properties and works well. The gray scale depth map technique (figure 7) effectively makes use of binocular disparity and also works well. Techniques that don't make use of these two properties look identical on the MLD™ and SLD. Techniques where these depth cues interfere with the displaying data reduce the perceived information. An example is figure 5 which makes poor use of motion parallax and therefore appears confusing to the user.

6.2 Rules for creating effective 3D displays on MLD™

Our research found no general technique that works well for all applications. A developer must make intelligent decisions about what to render on the front and back layer to produce an effective scene. The following rules are compiled from our experiences will help to make this decision.

Emphasize important objects
Rendering a scene on the back layer and putting selected objects onto the front layer emphasises them. Other techniques such simulating depth using gradients also accentuates objects. This is useful in applications such as advertising, where the advertised product can be accentuated, and visualization applications such as satellite information where the

designer wants to emphasise GIS information or military activity.

Making use of layer separation to separate information
Putting different datasets on different layers clearly shows that the datasets are separate but at the same time enables the user to read and compare both datasets.

Extruding objects across layers
Rendering an object over two layers, as explained in subsection 5.2 and 5.3, can give the illusion of physical depth and makes the scene more eye-catching.

Transition objects between layers
When moving objects between layers it is best to fade them between the two layers to give a continuous movement. This is useful when animating an object in 3D and a gradual movement between layers is required in order to emphasise its motion towards or away from the camera.

Making use of transparency
The transparency of the front layer can be used to render semi-transparent materials, such as glass, water and fog. The objects rendered on the back layer appear to be physically behind the semi-transparent material.

Avoiding visual discontinuity
When rendering a scene on the MLD™ it is important to take into account user head movements and that multiple users might view the display at the same time. In particular visual discontinuities as illustrated in figure 2 (top row) and figure 5, must be avoided.

7 Conclusion

Binocular disparity and motion parallax are the main depth cues users employ in order to determine which objects are displayed on the front and the back layer of the MLD™. Binocular disparity makes objects on the front layer appear closer and separated from what is displayed on the back layer. Motion parallax causes objects displayed on the front layer to move relative to objects displayed on the back layer when the viewer moves the head.

These depth cues cannot be depicted on SLDs and we have used them to develop effective rendering techniques for the MLD™. Gradients are useful for both reducing discontinuity caused by the physical gap between layers and for making objects appear continuous across layers. An effective general technique is to split a scene by Z-value to add a limited amount of physical depth to the scene. Important objects or objects that are closer to the viewer should be rendered on the front layer. In general, areas of an image can be made to appear some distance between layers by rendering them with appropriate transparency values on both layers. The example with the brick wall in figure 7 demonstrated that this works best if only a relatively small depth is simulated.

Care must be taken that the physical separation between layers does not lead to unnatural effects such as gaps between layer images and unrealistic motion parallax (see figure 5 where the car's head lights move in an unnatural way).

In future research we want to develop an OpenGL style graphics library for use with MLD™. This might involve the development of special graphics card drivers to make full use of hardware acceleration.

References

[1] PureDepth. *PureDepth Multi-layer Display*. Retrieved from http://www.puredepth.com on 15 March 2006.

[2] R. Delaney, Forget the Funny Glasses. *IEEE Computer Graphics and Applications*, 25(3), May/June 2005, pp. 14-19, 2005.

[3] G. Mather, *Foundations of Perception*. Psychology Press, USA, 2006.

[4] D. McAllister, *Display Technology: Stereo & 3D Display Technologies*. Retrieved from http://research.csc.ncsu.edu/stereographics/wiley.pdf on 15 March 2006.

[5] C. Pedroza, *Visual Perception: Gestalt Laws*. Retrieved from http://coe.sdsu.edu/eet/Articles/visualperc1/start.htm on 28 March 2006.

[6] W. Ittelson, *Visual Space Perception*. Springer Publishing, USA, 1960.

[7] T. Widjanarko, Brief survey on Three-Dimensional Displays: from Our eyes to Electronic Hologram, 2001. Retrieved 23 April 2006, from http://filebox.vt.edu/users/twidja na/cv/holopaper.PDF.

[8] Sullivan, A (2004) 3-Deep: New displays render images you can almost reach out and touch. *IEEE Spectrum*, May 2004, p 30-35.

[9] V. Prema, 3D Visualization Techniques with Multi-Layer Display Technology, SE Project Report, University of Auckland (submitted for publication).

[10] G. Roberts, 3D Visualization Techniques with Multi-Layer Display Technology, SE Project Report, University of Auckland (submitted for publication).

[11] Nakayama, K. & Silverman, G.H. (1986). Serial and parallel processing of visual feature conjunctions. *Nature*, 320, pp. 264-265.

[12] Kessenich, J. (2006), *The OpenGL Shading Language*. Retrieved 1 September 2006 from http://www.opengl.org/registry/specs/ARB/GLSLangSpec.Full.1.20.6.pdf.