

Volumetric Hyper-Reality, A Computer Graphics Holy Grail for the 21st Century?

Gavin Miller,
Apple Computer, Inc.

Abstract

This paper starts by examining the issues related to integrating real and virtual objects in a virtual reality environment. In particular it discusses the problems of occlusion, shadows and specular reflections. As an alternative to immersive virtual reality, this paper considers existing 3-D display technologies. It then discusses the limitations of these approaches and presents a thought experiment to design a completely general volumetric display. Such a display would convincingly create the illusion of objects with arbitrary optical properties. A metallic object depicted using the display would reflect the visual surroundings of the display. Dielectric materials would show correct refraction and reflection effects. Light shone on the display would illuminate the virtual objects within it. When programmed to depict empty space, the display would, for all practical purposes, disappear, rendering the contained volume invisible. Incremental steps towards such a device are discussed.

Keywords

3D display, volume rendering, invisibility, virtual reality, hyper-reality, hype.

1.0 Introduction

Three-dimensional computer graphics has had two major goals since its invention as an area of endeavor. The first goal is to create ever more realistic depictions of objects, often termed "photo-realism". (This has led to the development of techniques for hidden surface elimination, accurate surface shading and lighting, global illumination, and high complexity modeling.) The second goal is to provide systems rapid enough to display real-time imagery. Usually retracing the steps towards photo-realism, real-time systems have been expensive and somewhat domain specific. In recent years, however, the real-time systems have become cheaper and more powerful. In a few years, many desk-top computers will have the power of multi-million dollar flight simulators from a decade ago. As photo-realism matures and real-time systems catch up, it is a good opportunity to examine the goals of computer graphics, and suggest a new goal, or holy grail, which may

be appropriate as a technological research direction for the 21st Century.

1.1 Properties of Useful Holy Grails

To come up with a new holy grail for computer graphics it is fun to think about what attributes make a holy grail useful.

Successful holy grails have a number of properties:

1. The goal must be inspiring.

A holy grail, if attained instantly, would be sufficiently revolutionary that the results would resemble magic or at least inspire astonishment. (The airplane is a good example.) Alternatively they may be so obviously useful as to be wanted without question. (Free energy with no pollution.)

2. The goal must, at least in principle, be attainable.

A holy grail which defies the laws of physics, as we know them, might be inspiring, but impossible to accomplish. (A perpetual motion machine or faster than light travel are examples of these types of problems.)

3. Incremental steps towards the goal should be useful.

To keep enthusiasm for the mission, shorter term rewards are desirable. These help fund the endeavor and such partial solutions may be very useful in their own right.

4. The goal should be slightly vague.

By giving latitude in the problem, certain approximations are allowed. (An anti-gravity device which only supports an object above a special track allows the use of magnetic levitation.)

5. Once attained, the goal should not eat you

Some things seem great in anticipation, but once attained may be so suitable for adverse purposes that it would be better not to make them. (Pick your own examples.)

2.0 Real Objects inside Virtual Reality

Immersive virtual reality holds the promise of being a good holy grail. A different image is fed to each eye by a head-mounted display and the head is tracked to compute the appropriate view of the computer-generated world. Such a system was described in Sutherland '68. The display in that case used a half-silvered mirror so that synthetic objects were visible superimposed on the real objects surrounding the user. A more general way of integrating real and synthetic objects is to place a camera in front of each eye and then combine the synthetic and real images using the computer. Figure 1 shows such an augmented reality system. We might term this approach "Composite Virtual Reality", since real and imaginary objects are composited together. Such an arrangement of cameras and head-mounted display was used to merge ultrasound data with the body of a real patient in Bajura et al. '92. Compositing real and imaginary objects in a seamless way has been a long time preoccupation of the special effects community. The step we are taking here is to require the techniques to work for arbitrary scenes, in real time.




Fig. 1. Composite Virtual Reality Architecture.

While the results from Bajura et al. '92 were impressive, more seamless integration might be possible. To fully integrate physical and virtual objects, it is necessary to compute the effects of hidden surface elimination, the casting of shadows and the appearance of specular surfaces. Techniques which might be appropriate for achieving this in the near future are discussed in the next few sections.

2.1 Hidden Surface Elimination

An important prerequisite for realism is to determine the problem of how the real and imaginary objects obscure each other. The traditional approach for special effects shots is to have the objects ordered in back to front order with an accompanying mask or "matte". The objects are then composited in the usual way using an "over" operation (Porter and Duff '84).

Trying to do this in real-time for arbitrary real and imaginary objects is especially challenging. Compositing with mattes will not be a general solution to the hidden surface problem, since the ordering of the scene can change on the fly. In addition, real and imaginary objects may interpenetrate, leading to the need for computing visibility in the region of the intersections.

A more robust way to combine real and imaginary objects is to use Z-buffer compositing in which the nearest pixel is chosen based on both the measured depth for the real scene, and the computed depth for the synthetic scene. To avoid aliasing artifacts, we can use alpha masks or super-sampling (Duff '85). Such super-sampling would require very high resolution depth cameras. While challenging, they are a conservative assumption compared to devices described later in this paper. We must remember to match the lighting to the real scene. Matching the lighting will only be truly effective if we also take shadows into account.

2.2 Shadows For Real and Virtual Objects

Shadows will affect the appearance of the scene in two ways:

1. Synthetic objects cast shadows onto the real objects.

The synthetic-to-real shadows may be computed using traditional graphics techniques. We know the 3-D position of the real surface point. This can be tested against the synthetic objects using ray-tracing or shadow Z-buffering (Reeves et al. '87). For a real surface point computed to be in shadow, we then have to falsely shade it without the direct illumination. Rather than trying to estimate the contribution from the light source, we could use video projectors as the light sources. By projecting shadow mattes onto the real objects in the scene, virtual objects may cast shadows onto real objects. (This is much in the spirit of the use of shadow mattes in Reeves '85.) Alternatively, the surface could be illuminated by one light source at a time. The real surface intensity could then be computed for any set of light sources visible from that point.

2. Real objects cast shadows onto the synthetic objects.

This may be computed if we place a depth camera at every light source shining on the scene. This gives us the ability to compute shadows using shadow Z-buffers (Reeves et al. '87). We have to fake the ambient term for the synthetic objects since the true ambient term depends on interreflections from the real scene in a way which we cannot compute.

2.3 Specular Reflections of Real Objects

An equally difficult problem is if the synthetic objects are specular. We should see reflections and refractions of the surrounding real objects. In particular we should probably see reflections of parts of the real objects which are not

visible in the original image. Consider a synthetic flat mirror. It should reflect the back side of objects in front of it, which are clearly hidden from the user's point of view. What we need is a way to capture the appearance of objects seen from the location of the synthetic objects.

In film production it is possible to capture an environment map around the location of a synthetic object using a shiny ball or reflective sphere placed in the real scene. For a real-time interactive system this is more problematic. Portions of the shiny ball may be obscured by real objects in the scene. In particular, a user's hand reaching out to touch the shiny object would obscure parts of the ball. These parts may correspond to orientations needed to display the synthetic object at that location. Also, more pedantically, the approximations of environment mapping would break down as the real and imaginary objects came into close proximity. A way to avoid the occlusion problem is to replace the reflective sphere with a full-sphere panoramic camera. This camera would capture the appearance of the scene in every direction and the result could be processed on the fly to produce an environment map for use by the environment mapping hardware.

Environment mapping can be used to approximate ray-tracing for two different optical situations. The first is shown in Figure 2 and consists of a small object surrounding the environment map origin. Reflections for the small object are computed using the reflected ray direction in the environment map. For the approximations to remain accurate, the origin of the environment map must move to track the center of the virtual object. In this case, the environment camera must be physically moved to track the motion of the virtual object.




Fig. 2: Approximate Ray-Tracing for Small Objects

A second use of environment mapping is shown in Figure 3 in which an environment map is used to compute reflections for a planar virtual surface. For a real surface which is not completely smooth and flat, surface reflections might be blurry or show ripples in the surface. To create these effects for imperfect synthetic reflective surfaces, the reflections can be defocused by filtering the environment map, or small ripples may be rendered convincingly using bump mapping (Blinn '78). In such a system, a real object could be seen reflected in an imaginary rippling lake, in real time. The catch for the planar approximation is that the environment map origin must be at the mirror reflection of

the user's head position. A virtual reality mirrored wall would require an environment camera moving around in a physical space which was the mirror image of the user's physical space. If the user tried to move his or her head through the virtual mirror, he or she would collide with the motorized environment camera coming in the opposite direction! Similarly, a reflective virtual floor would require a transparent physical floor with the motorized camera moving around under the user's feet.




Fig. 3: Environment Mapping for Planar Objects

Despite these objections, the virtual reality approach to interactive graphics will probably prove extremely fruitful in the next decade, since the display and computational requirements are well matched to existing and near-future technology.

3.0 Virtual Objects in Physical Reality

An alternative approach to immersing the user inside a display, is to create a device which exists inside the user's real world. Most displays are currently like this, and they have a number of natural advantages:

1. They do not interfere with interactions between people in the same room.
2. They are less intrusive on the user's experience. The user can look at the display when required, or just get up and walk away.
3. They are not tiring. Immersive displays create fatigue and nausea (Pausch et al. '92).
4. They typically supply high resolution experiences, since they subtend less of a user's field-of-view than immersive displays. Also, the rest of the environment is at "full resolution" since it is viewed directly by the user, rather than through a camera-display combination. (Deering '92.)

Physical displays which attempt to create a three-dimensional illusion inside the user's real world are discussed in the following sections.

3.1 Additive Volumetric Displays

One approach to rendering a convincing three-dimensional illusion is to super-impose a number of planes of emissive light in such a way that they stack up to form a volume of glowing elements. Such a display is often called "volumetric" or a direct volume display. It has the advantage of creating full motion parallax and not requiring any special glasses. We shall call this approach "additive volumetric" since the planes of light do not obscure each other, and merely add intensity in that volume of space. The view through such a display is the sum of the intensities of the voxels along the view direction.

Many mechanisms have been proposed for additive volumetric displays:

1. A Spinning L.E.D. array

Jansson and Goodhue '81 proposed using a rapidly rotating 2-D array of light emitting diodes. An angular shaft encoder addressed the appropriate planar slice of data from the image memory. By spinning the array fast enough (above 30 revs/second), visual fusion occurs, giving a glowing volumetric display. A disadvantage of the scheme was the variable sampling of the volume, the center having higher resolution than the edge.

2. Vibrating L.E.D. array

Kameyama et al. '93, used a 2-D array of L.E.D.s which was displaced along a linear path to create a volumetric effect. This had the advantage of keeping the mapping between memory and physical space straightforward, but somewhat reduced the viewing angle, since the rear side of the display was opaque.

3. Projected Light onto a 1-D Stack of Electrically Switchable Mirrors

Buzak '85 described the use of electrically switchable mirrors in a one-dimensional stack. At any given time interval, one mirror becomes reflective while the others are all transparent. During that interval, an image is seen reflected in the mirror from a traditional 2-D display device such as a cathode ray tube. By activating the mirrors in sequence and simultaneously displaying the corresponding image on the 2-D display, a 3-D illusion can be created. Disadvantages of this particular approach arise from the fact that it is not easy to scale. The mirrors need to be extremely transparent when "off". Even with nearly perfectly transparent materials, stacking several hundred such planes together would reduce the contrast to almost zero. (This will be referred to as the "stacking problem". It arises whenever large numbers of sheets of material are stacked on top of each other.)

4. Orientation Selective Holograms

Bartelt and Steibl '85, proposed using orientation selective holograms. Such a hologram only reflects light shone onto it from a particular direction. A display may be

created by making a stack of holograms each of which selects a different orientation. A set of projectors is placed relative to the display one at each of those orientations. The image from each projector would only be reflected by the corresponding layer in the display. The net result would be an additive volumetric image. This approach has the advantage that the projectors could even be slide projectors, since the layers are added together without time slicing. The disadvantage of this approach is the stacking problem. Bartelt and Steibl '85 only demonstrated a 2-layer display.

5. A Varifocal Mirror

Harris et al. '86 described using a varifocal mirror to modulate the apparent depth of a virtual image. The mirror consisted of an aluminized membrane which vibrated using sound waves. As the membrane changed curvature, the location of the reflection of a cathode ray tube was displaced in depth. The cathode ray tube was driven in vector-scan fashion from a frame-buffer with appropriate slices of the volume data displayed in synchronization with the vibrating mirror. The field of view of the display was somewhat limited since the display had to be seen reflected in the mirror.

6. Laser Projection onto a Spinning Double Helix

Soltan et al. '92 described using a scanning laser system to project onto a spinning double helix. The helix sweeps out the volume as it spins and a laser is deflected using acousto-optical scanners to create points of light on the surface at different depths. The laser deflection system allows between 4,000 and 40,000 points to be illuminated in 1/20th of a second. The authors also mention the possibility of shining a laser into a solid matrix of non-linear optical material, but the details are sketchy.

All of the additive volumetric displays have the same drawbacks. The displays are unable to take account of occlusion and surface shading. Most of the usual visual cues have been eliminated to allow for a multi-user full-motion parallax display without special glasses. The shortcomings of this approach will become more apparent as the capability of the displays increase. As the volume images become more detailed and complex, the resultant cognitive load on the user will become more taxing. Such displays may be important for certain niche applications, such as air-traffic control, where it is important to resolve the spatial relationship of a few point-like objects in a diagrammatic format. However, they are not the ideal solution for a completely general computer graphic display.

3.2 Stereoscopic Displays

A second approach to providing three-dimensional pictures is the traditional stereoscopic display of a left and right image, one for each eye. While this does allow for hidden surface elimination and realistic shading, it has a number of limitations. Stereo-pair displays suffer from the fact that as the user moves his or her head, the motion parallax is incorrect, leading to an unstable perception of the three-

dimensional scene. To counter this effect, the user can be tracked, and the graphics display updated appropriately. (Deering '92).

One requirement of all stereo displays is the need to supply different images to the left and right eyes. A major technical challenge is to minimize "cross-talk" in which the image for one eye is fed inadvertently, and in an attenuated form, to the other eye.

Stereo display may be achieved in a number of ways:

1. Red-Green Stereo

Color filters are placed over each eye and the corresponding images are displayed in the appropriate color channel on the display. This approach can only display monochrome images and can lead to strange perceptual color effects.

2. Active Shutter Glasses.

A light valve on each eye is opened and closed appropriately as the corresponding image is displayed on the screen. To avoid flicker, the images are usually displayed at 120 Hz. Such light valves have limited transparency when "open" and users may find the wearing of glasses burdensome. A very rapid response display is required to prevent cross-talk.

3. Active Polarizing Screens.

A polarizing shutter is placed in front of the screen, and the user wears passive glasses to direct the polarized light to the appropriate eye. Again, cross-talk can be a problem.




Fig. 4: Stereo Display using Lenticular Screens

4. Lenticular Stereo Screens

To avoid having to wear glasses, it is possible to use an array of cylindrical lenses placed in front of a raster display. This "lenticular" system directs alternate pixels to the left and right eye. See Figure 4. The pixels are displayed over a certain solid angle determined by the

position of the cylindrical lenses relative to the display. By displacing the lenses horizontally it is possible to change the ideal viewing location. By tracking the approximate location of the user, it is possible to make the lenticular display effective over a large viewing volume, (Shiwa et al. '94).

5. L.C.D. Screens with Directional Backlighting.

One problem with lenticular screens is the loss of horizontal resolution for the sake of stereo. An alternative approach described in Eichenlaub '93 uses vertical strips of backlighting to control the directionality of the L.C.D. Two sets of vertical strips are illuminated sequentially, one for each viewing direction. By changing the display at twice the usual frequency, a stable stereo effect is achieved. Fine control of the backlighting direction also allows head-tracking for the display, with no moving parts. This arrangement is shown in Figure 5. By having more than two sets of lines, more than two viewing zones may be displayed. The more zones, the faster the L.C.D. must be switched between images. This leads us to the topic of integral displays.




Fig. 5: Integral Display using Backlighting Modulation.

3.3 Integral Displays

Integral displays attempt to account for motion parallax by showing many images simultaneously, projected into a spread of orientations.

Integral displays fall into two broad categories:

1. Unidirectional Integral Displays

Unidirectional Integral Displays have a 1-D array of images which project onto a horizontal spread of orientations.

2. Omnidirectional Integral Displays

Omnidirectional Integral Displays have a 2-D array of images which project onto a two-dimensional solid angle of

orientations. Okoshi '76 describes the history of these ideas for static photography as well as static holography.

Integral displays have been implemented in the following ways:

1. Virtual Vertical Slits using a Cylindrical Holographic Optical Element

A cylindrical holographic optical element is used to create a virtual slit which passes in front of the viewer as the cylinder rotates. By projecting a sequence of images onto the hologram as the cylinder rotates, a unidirectional integral display is achieved (Shires '93).

2. 2D to 1D Interleaving Using a Planar Holographic Optical Element.

With a standard lenticular display, horizontal resolution is sacrificed for orientation dependence. In the case of a stereo display this loss of a factor of two in resolution may be acceptable. However, as the number of images increases, the disparity between vertical and horizontal resolution becomes unwieldy. Shires '94 describes using a flat plate holographic optical element to transpose a square array of screen pixels into a linear set of orientations. Thus a display of 10 times the linear spatial resolution allows for a hundred orientations of viewing. Of course this approach does not support vertical parallax.

3. Omnidirectional Integral Display using Retro-reflectors and Multiple Projectors

An omnidirectional integral display has been described using a retro-reflective screen in which a 2-D array of projectors created a set of images which gave both vertical and horizontal motion parallax (Borner '93). The number of projectors was limited to four rows of six, but could, in principle be increased.

4. Omnidirectional Photographs using Fly's Eye Lens Array

Okoshi '76 describes various people's photographic schemes for using a 2-D array of lenses in the image plane to create 3-D photographs with both horizontal and vertical motion parallax. He refers to this structure as a fly's eye arrangement, since it resembles the arrays of lenses in a fly's eye.

5. Synthetic Aperture Holography

A synthetic holographic method has been described which uses traveling sound waves in an acousto-optical modulator to create diffraction patterns. Currently the field of view and resolution is limited, but may be improved over time (St. Hilaire et al. '92).

4.0 Volumetric Hyper-Reality

It is instructive to step back from the technological details of existing 3-D displays, and examine the ultimate goal

towards which these technologies aspire. Optical holography has often been held up as the ideal towards which 3-D display technology strives. Recreating the appearance of objects in the vicinity of an image plane with full motion parallax, hidden surface elimination and shading cues, etc. seems like the ultimate display technology. As a goal it is yet to be achieved. However, as a goal it is also limited to some extent. Since the Renaissance, the idea of depicting a scene with an image has become the dominant metaphor for pictorial representation. The Nineteenth Century invention of photography seemed to enshrine images as *the* depiction of optical truth. The invention of holography and integral photography seems to have provided the final missing cue, namely motion parallax.

It is worth considering that, even in two-dimensions, a photograph is a limited subset of the optical possibilities. Surfaces may be reflective, with metallic or glossy optical properties. They may also be transparent. Similarly, in three dimensions, objects can be reflective, or refractively transparent. A photograph of a mirror merely shows a reflection of the camera. A photograph of a glass object refracts the light from that object's surroundings. The idea of putting your hand behind a photograph of a glass and seeing your hand refracted by the glass seems inappropriate or even amusing. We have been subconsciously trained to accept the limitations of the medium.

As a thought experiment let's specify an ideal volumetric computer graphic display. It would contain a volume of space within which objects could be defined with arbitrary optical properties:

- A metallic object depicted using the display would reflect the visual surroundings of the display. (A user would see his or her face reflected in a mirror specified inside the display.)
- Dielectric materials would show correct refraction and reflection effects. (A person could put their hand on the other side of the display and see it refracted and reflected by objects within the display.)
- Light shone on the display would illuminate the virtual objects within it, including the effects of interreflection, and caustics.
- When programmed to depict nothing, the display would, for all practical purposes, disappear, rendering the contained volume invisible.

A display with all these properties would be said to embody the idea of volumetric hyper-reality. It could be called a hyper-display for short. A synthetic object within such a display would be indistinguishable from a corresponding real object at the same location.

As a second thought experiment, let's try to design such a device, to be built some time in the Twenty-First Century. We will make two assumptions about the progress of technology:

- Computational and communications resources will become inexpensive, both in terms of space and power consumption.
- The surface micro-structure of the display may be arbitrarily complex. This includes computational elements, light emitting devices, light sensitive devices and optical elements such as lenses.

4.1 The Null Set and Invisibility

The first question to address is how to come up with a display technology which, when it depicts empty space, disappears. In "The Invisible Man" by H.G. Wells, a person is rendered invisible by first eliminating all light absorbing chemicals, (the man is an albino) and secondly by reducing the refractive index of the materials of the body to unity. Unfortunately, there is no reason to think that refractive indices can be reduced to unity for useful electronic materials. Rather than modify the bulk properties of materials, we might concentrate instead on making a surface micro-structure which creates the illusion of these bulk properties.

Given that the outer surface of a display has to be made of a transparent material such as glass we have three problems. The first is reflections at the surface of the glass. No matter what is inside the display, light will still reflect off this boundary. A second problem is the refraction of the light at the interface. The final problem is the unencumbered passage of optical information through the volume in every direction.

A solution to these problems may be found in the form of a hyper-pixel. A hyper-pixel is a device which has the following two properties:

1. It may emit modulated light over a hemisphere of orientation. (It is like a high resolution video projector with a 180 degree fish eye lens.)
2. It may sense light over a hemisphere of orientation. (It is like a high resolution video camera with a 180 degree fish eye lens.)

By covering a surface in such hyper-pixels, a display of arbitrary optical properties may be created. (This is a generalisation of the fly's eye integral photograph in Okoshi '76.)

Consider a planar slab of material which has a 2-D array of hyper-pixels on both sides. Such a device is shown in Figure. 6.

For every pixel on the left surface, the incoming light may be detected in all directions. This information is transferred to those hyper-pixels on the right surface which correspond to the location that the unencumbered rays would have reached when they crossed the righthand surface. These intensities are then displayed by the righthand hyper-pixels in the appropriate directions. (The converse is also true for light traveling in the opposite direction.) Figure 6 shows light emitted from a single hyper-pixel and the light entering the hyper-pixels from which it received illumination information. This renders the intermediate volume effectively invisible even though it contains the bus structure and whatever else will fit in the space. Creating such an effect will be dubbed "computational transparency" as opposed to the physical transparency of uniform dielectric materials.




Fig. 6: Planar Two-Sided Planar Hyper-Display.

To overcome the problem of reflections at the surface is quite straightforward. For a given hyper-pixel, the light it should emit because of computational transparency is provided to it by the hyper-pixels on the other side of the slab. The amount of light falling onto the hyper-pixel itself is known from its own sensors. We also know the reflective properties of the hyper-pixel's lens. We can compute the amount of light reflected by the lens in a particular direction from the light falling onto it. We may subtract this amount of reflected light from the amount the hyper-pixel is supposed to emit because of computational transparency. Provided that the transmitted light is bright enough, the surface reflections may be eliminated by this subtraction. (If the slab has total darkness on one side, the reflections on the other side can't be eliminated. Even in the future it is impossible for incoherent light sources to shine negative light.)

There is no reason to build just a planar device. A sphere covered in hyper-pixels would render the contained volume invisible. Probably the most challenging region of the display to make invisible would be the outline of the sphere, where the hyper-pixels would be nearly edge on, and the optics of dielectric media make the reflections the brightest. Also, unfortunately, the bus structure of such a

device would need to connect every hyper-pixel to every other hyper-pixel. This connectivity is shown in Figure. 7. The physical interconnections might actually just be a very fast linear bus.

We need to consider the thorny issue of resolution. A hyper-display would only be transparent down to the physical spacing of the hyper-pixels (the spatial resolution). Below that size there would be either aliasing artifacts or a loss of clarity. The angular resolution of the hyper-pixels would determine the depth of field of the display. Objects in the far distance behind the display would be out of focus.




Fig. 7: Bus Structure for Spherical Hyper-Display

4.2 Hyper-realistic Remote Presence

So far we have considered the case of a transparent volume. With sufficient bandwidth it is possible to use a hyper-display to create convincing remote presence. Consider a convex spherical hyper-display in one location and a concave spherical hyper-display in a second location. On the "convext" sphere, the hyper-pixels faces outwards. On the "concave" sphere the hyper-pixels face inwards. This arrangement is illustrated in Figure 8.

Light shone onto the convext display would illuminate the interior of the concave display. Objects inside the concave display would appear to be at the corresponding location inside the convex display. Shadows cast onto the convex display would shadow objects within the concave display. A person inside the concave display would see everything in the room containing the convex display, including his own reflection in mirrors etc. It would be the ultimate video phone. (By the time it is possible to build such a device, bandwidth may be cheap enough to afford the phone bill.)




Fig 8. Hyper-realistic Remote Presence

4.3 A General Graphics Model

A completely general graphics model would simulate the interior of a concave hyper-display and create all of the corresponding optical properties. It would take the incoming light values from the surface hyper-pixels and combine them with interior illumination to compute the light values to be emitted. For a fixed interior scene, this computation would be a linear mapping of the incident intensities. The transformation from incident light to emitted light can be called the "transluminance matrix". Instead of a display being used to display an image in terms of pixel values, it would display a transluminance matrix. This matrix would be applied to the incident illumination to compute the emitted light once every refresh cycle, even if the model did not change.

Needless to say, all of the computation for a completely general hyper-display would be daunting. The null case for empty space is beyond the current state of the art in terms of bus architectures, computation, and display technology. The computation of a synthetic transluminance matrix is also extremely costly, being equivalent to computing global illumination for hundreds of images. It can also be argued that such a display would be very wasteful computationally since all views of a scene are computed simultaneously.

4.4 An Interesting Special Case

As stated at the beginning of this paper, a useful property of a good holy grail is that early steps towards the goal should yield encouraging results. An incremental steps towards a hyper-display is shown in Fig. 9. It shows a 2-D display with a micro-camera focused on the user. A second camera points backwards away from the rear of the screen. Using the data from the cameras, and special algorithms, it is possible to apply an approximate transluminance matrix for the specular reflections and refractions in a scene, in near real time. An early prototype of such a system, using a single camera, has already been built in our lab.




Fig. 9. Flat Panel Display Plus Two Cameras.

5.0 Acknowledgements

This paper was written in the Advanced Technology Group of Apple Computer, Inc.. It reflects many stimulating discussions, especially with Lance Williams, who suggested the use of video projectors for casting shadows onto real objects, as well as introducing me to the problem of invisibility.

6.0 References

Bajura, Michael, Henry Fuchs, Ryutarou Ohbuchi, "Merging Virtual Objects with the Real World: Seeing Ultrasound Imagery within the Patient", SIGGRAPH '92, Chicago, July 26-31, 1992, pp 203-210.

Batelt, H. and N. Streibl, "Three dimensional display based on thick holographic phase components", Optical Engineering, Nov. and Dec. 1985, Vol. 24, No. 6.

Blinn, James F., "Simulation of Wrinkled Surfaces," Computer Graphics, Vol. 12, No. 3, August 1978, pp 286-292.

Bšrner, R., "Autostereoscopic 3D-imaging by front and rear projection and on flat panels", Display, Vol. 14, No. 1, 1993.

Buzak, Thomas S., "A Field-Sequential Discrete-Depth-Plane Three Dimensional Display", 1985 SID International Symposium Digest of Technical Papers, May, 1985, ISSN 0097-966X.

Deering, Michael, "High Resolution Virtual Reality", SIGGRAPH '92, Chicago, July 26-31, 1992, pp 195-202.

Duff, Tom, "Compositing 3-D Rendered Images", Computer Graphics, Vol 19., No. 3, 1985 pp 41-44.

Eichenlaub, Jesse B., "Developments in Autostereoscopic Technology at Dimension Technologies, Inc.", SPIE Col.

1915 Stereoscopic Displays and Applications IV (1993) / 177.

Harris, L. D., J. J. Camp, E. L. Ritman, and R. A. Robb, "Three-Dimensional Display and Analysis of Tomographic Volume Images Utilizing a Varifocal Mirror", IEEE Transactions on Medical Imaging Col. MI-5, No. 2, June 1986.

Jansson, D. G., and J. T. Goodhue, "Medical applications of a new 3-D display system", Proc. National Computer Graphics Association, Baltimore, Maryland (1981).

Kameyama, Ken-ichi, Koichi Ohtomi and Yukio Fukui, "Interactive volume scanning 3-D display with an optical relay system and multidimensional input devices", SPIE Vol. 1915 Stereoscopic Displays and Applications IV (1993).

Lucente, Mark, Pierre St.-Hilaire, Stephen A. Benton, Derrick L. Arias, John A. Watlington, "New Approaches to Holographic Video", SPIE Vol. 1732, Holographics International '92 / 377.

Okoshi, Takanori, "Three-Dimensional Imaging Techniques", Academic Press, 1976, ISBN 0-12-525250-1.

Pausch, Randy, Thomas Crea, and Matthew Conway, "A Literature Survey for Virtual Environments: Military Flight Simulator Visual Systems and Simulator Sickness", Presence, Vol. 1, No. 3, 1992.

Porter, Thomas, and Tom Duff, "Compositing Digital Images", Computer Graphics, Vol. 18, No. 3, 1984, pp 253-260.

Reeves, William T., "Approximate and Probabilistic Algorithms for Shading and Rendering Structured Particle Systems", Computer Graphics, Vol. 19, No. 3, 1985.

Reeves, William T., David H. Salesin, Robert L. Cook, "Rendering Antialiased Shadows with Depth Maps", Computer Graphics, Vol. 21, No. 4, July 1987, pp 283-292.

Shires, Mark, "3D real-time Video Display Using a Cylindrical HOE", SPIE Vol. 2043 Holographic Imaging and Materials (1993).

Shires, Mark, "Real-time Flat Panel Solid-state Holographic Stereogram", SPIE Proceedings, Fifth International Symposium on Display Holography, July 18-22, 1994.

Shiwa, Shnichi, Nobuji Tersutani, Kenjii Akiyama, Susumu Ichinose and Tadahiko Komatsu, "Development of Direct-View 3D Display for Video-phones Using 15 inch LCD

and Lenticular Sheet.", IEICE Trans. Inf. and Syst. Vol. E77-D, No. 9, Sept. 1994.

St. Hilaire, Pierre, Stephen A. Benton, and Mark Lucente, "Synthetic aperture holography: a novel approach to three-dimensional displays", Vol. 9, No. 11 / November 1992 / Journal of the Optical Society of America.

Solatan, Parviz, John Trias, Waldo Robinson, Weldon Dahlke, "Laser Based 3D Volumetric Display System", SPIE Vol. 1664 High-Resolution Displays and Project Systems, (1992) / 177.

Sutherland, Ivan E., "Head-mounted three dimensional display", AFIPS Conference Proc. Vol. 33, Pt. 1, Fall Joint Computer Conference, San Francisco, California, Dec 9-11 1968 pp 757-764.