This article appeared in *Computer Graphics*, ACM SIGGRAPH, 31(2), May 1997, pp 58–62. © 1997 Michael Halle, ARR.

Autostereoscopic displays and computer graphics

Michael Halle

Surgical Planning Laboratory
Department of Radiology
Brigham and Women's Hospital
Boston, Massachusetts, USA
mhalle@bwh.harvard.edu

ABSTRACT

Autostereoscopic displays present a three-dimensional image to a viewer without the need for glasses or other encumbering viewing aids. Three classes of autostereoscopic displays are described: re-imaging displays, volumetric displays, and parallax displays. Re-imaging displays reproject an existing three-dimensional object to a new location or depth. Volumetric displays illuminate points in a spatial volume. Parallax displays emit directionally-varying image information into the viewing zone. Parallax displays are the most common autostereoscopic displays and are most compatible with computer graphics. Different display technologies of the three types are described. Computer graphics techniques useful for three-dimensional image generation are outlined.

1 INTRODUCTION

After many years of relative obscurity, three-dimensional displays have recently become both increasingly popular and practical in the computer graphics community. This interest can be attributed to many factors. In our daily lives we are surrounded by synthetic computer graphic images in print and on television, and can now even generate similar images on personal computers in our home. We also have holograms on our credit cards and lenticular displays on our cereal boxes. And has it really been so many years since we first saw Princess Leia projected into thin air in the Star Wars motion picture? In fact, the general public has been excited about three-dimensional images since the days when stereoscopes graced every mantelpiece at the turn of the century, through the 3D movie craze of the early 1950's, the wonder of holography in the 1960's, and the new frontier of virtual reality today. With each new technology or movie, the excitement seems to grow.

Developments in the computer graphics industry have also done their part to make spatial images more practical and accessible. In the business of computer graphics, the computational power now exists for desktop workstations to generate stereoscopic image pairs for interactive display. At the high end of the computational power spectrum, the same advances that permit intricate object databases to be interactively manipulated and animated also permit large amounts of image data to be rendered for high quality three-dimensional displays. Finally, there seems to be a general realization in the research and scientific community that the two-dimensional projections of three-dimensional scenes traditionally referred to as "three-dimensional computer graphics" are insufficient for inspection, navigation, and comprehension of some types of multivariate data. For these databases, the oft-neglected human depth cues of stereopsis, motion parallax, and to a lesser extent ocular accommodation are essential for image understanding.

The broad field of virtual reality has driven the computer and optics industries to produce better stereoscopic helmet- or boom-mounted displays, as well as the associated software and hardware to render scenes at rates and qualities needed to produce the illusion of reality. However, most journeys into virtual reality are currently solitary and encumbered ones: user often wear helmets or other devices that present the three-dimensional world to them, and *only* to them. Presenting a three-dimensional image to a casual passerby, a group of collaborators, or an audience requires a different technology: autostereoscopic displays.

Surgical Planning Laboratory, Department of Radiology, Brigham and Women's Hospital, 75 Francis St., Boston, MA, 02115, mhalle@bwh.harvard.edu.

The information contained in this paper was collected in part while the author was at the MIT Media Laboratory.

AUTOSTEREOSCOPIC DISPLAYS

Autostereoscopic displays present a spatial image to a viewer without the use of glasses, goggles, or other viewing aids[1]. Autostereoscopic displays are appealing because they offer the best approximation to the optical characteristics of a real object. As a result, though, there is much misunderstanding and misinformation by those who would oversell the capabilities of a particular technology. This paper will try to outline the strengths and practical limitations of the different technologies by classifying them into broad categories.

Our current understanding of physics does not include a practical way of forcing photons to change direction in the absence of an optical medium. Thus, a fundamental and general statement can be made about all spatial displays, whatever its particular technology. This paper will refer to this requirement as the *projection constraint*:

A display medium or element must always lie along a line of sight between the viewer and all parts of a spatial image.

1. In the typically vague terminology of the three-dimensional display field, *autostereoscopic* has two different common meanings: the broad one used here, and a narrower one restricted to the technology that this paper calls parallax displays. The second usage is semantically more correct; unfortunately, no existing term encompasses the more important broader classification.

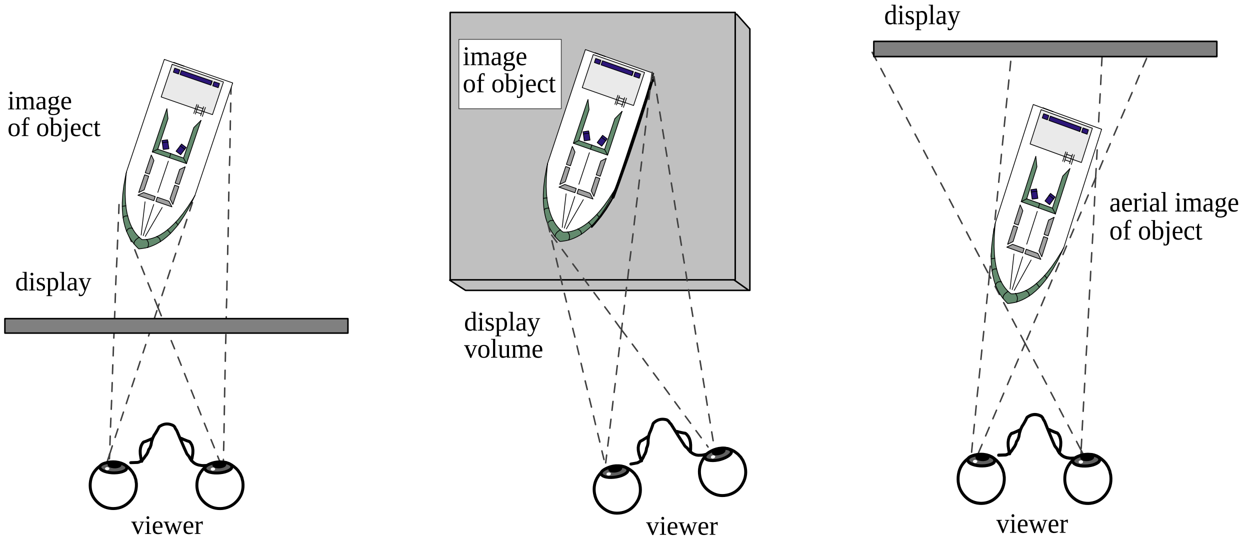

Fig. 1: All parts of an object must lie along a line of sight between the viewer and some part of the display.

Photons must originate in, or be redirected by, some material. The material can be behind, in front of, or within the space of the image, but it *must* be present. All claims to the contrary violate what we understand about the world. Figure 1 shows the possible relationships between the image and the display. A corollary to this constraint is the observation that air, water, or smoke are, in general, very poor display media. Images appearing "in mid-air", called *aerial images*, will invariably have originated not in the air but from some other medium. Technologies lavished with claims of mid-air projection should always be ly have originated not in the air but from some other medium. Technologies lavished with claims of mid-air projection should always be scrutinized with regard to the fundamental laws of physics.

A specific and practical result of the projection constraint is that no matter where a spatial image appears with respect to its display, the image will be clipped by the display's physical boundaries. If for instance, an image appears in front of its display, a sufficient translation of the viewer will cause part or all of the object to intersect and "fall off" the edge of the display. This condition, known as a *window violation*, is particularly disturbing for aerial images.Figure 2 illustrates a window violation.

Physically realizable autostereoscopic displays can be classified into three broad categories: re-imaging displays, volumetric displays, and parallax displays. Re-imaging displays capture and re-radiate the light from a three-dimensional object, perhaps to a new location in space. Volumetric displays span a volume of space, allowing individual parts of the space to be illuminated. Finally, parallax displays are surfaces that radiate light of directionally-varying intensity. Displays of each type have been used in commercial display systems, and each has inherent strengths and weaknesses.

Re-imaging displays

Re-imaging displays are technically the simplest of autostereoscopic displays. Re-imaging displays do not by themselves produce a three-dimensional image. Rather, they affect the appearance of another three-dimensional image in some way. The most trivial re-imaging displays is a plain piece of glass. The back surface of the glass intercepts rays of light traveling in different directions with different intensities. The energy of the light is propagated to the front surface of the glass, where light is re-radiated in the same direction and with the same intensity as when it was captured.

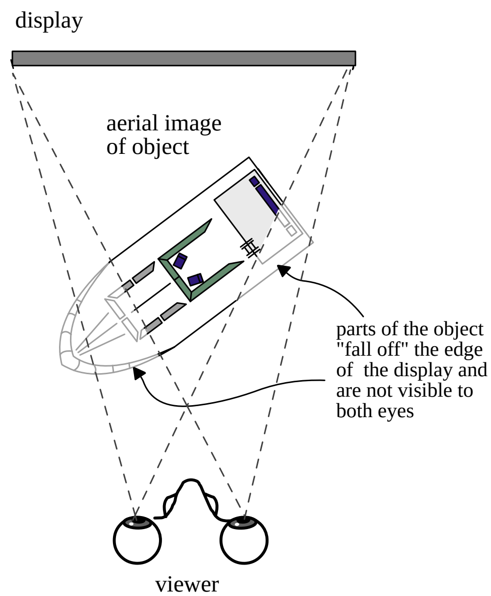

Fig. 2: A window violation.

Although the piece of glass is a very simple device, it does illustrate that even a passive optical element can display an arbitrarily complex light field that maintains the three-dimensionality of a scene.

A mirror is only slightly more complex than plain glass, but it is capable of altering the direction of all the rays of light entering it without changing either the intensity or the color of the light itself. A semi-silvered mirror can superimpose the light from two three-dimensional scenes and re-radiate the result. Mirrors, semi-silvered or not, are perhaps the most effective (and cost-effective) three-dimensional displays used in theme parks today.

More complex re-imaging displays are based on lenses and mirrors with optical power that can translate the position of an object in depth or distort it into a different three-dimensional shape. The three-dimensional display by Dimensional Media Associates uses a mirror system to relay a two- or three-dimensional object out in front of the display surface [8]. If the two-dimensional object is a CRT screen, a flat image of the screen will be appear to float in front of the device. The location of the viewer must be restricted to minimize window violations.

Another type of optical system was used by SEGA in an unsuccessful arcade video game called “Time Traveler” (marketed under the misleading term “hologram”) to relay and distort the appearance of a flat CRT into a curved surface. The successful commercial billing of re-imaging displays as “holographic” systems is a clear statement about how vivid the images they produce can be. On the other hand, the use of these devices in technical display applications is severely limited by their inability to display general three-dimensional information. Optical re-imaging is often incorporated into more general autostereoscopic displays.

Both of the other two classes of autostereoscopic displays described here, volumetric displays and parallax displays, can produce more general synthetic three-dimensional images. The major difference between the volumetric displays and parallax displays lies in the way they address the three-dimensional image volume. A volumetric display addresses individual points in the volume explicitly: input to this type of display can be a voxelized data volume or a display list of three-dimensional primitives. This data is then drawn in the three-dimensional space. In contrast, a parallax display device images the direction and intensity of light at many different locations on the display surface. Unlike the explicit three-dimensional input for the volumetric display, the parallax display’s input consists of two-dimensional projections such as photographic or synthetic images. Each two-dimensional image contains no explicit depth information. Instead, depth is implicitly encoded as positional disparity between different projections. The next two sections look more closely at volumetric and parallax displays.

Volumetric displays

Volumetric displays work by filling or sweeping out a volume of space. In the inexact terminology of three-dimensional imaging, volumetric displays are also called volume displays, slice stacking displays, and space filling displays. Several technologies of volumetric displays exist. One of the most elegant is the varifocal mirror display [21], shown in Figure 3. A varifocal mirror display uses a mirrored membrane of varying optical power to sweep an image of a CRT through different depth planes of a volume. By synchronizing the CRT display with the mirror’s oscillation, any point within the volume can be displayed. The greatest difficulty with varifocal mirror displays is building a high quality varifocal optic that can be oscillated at high frequencies.

Another group of displays use a spinning element to physically sweep out a volume. The element can be a simple rectangle spinning in a cylinder, or it can be a more complicated shape such as a helix. The spinning element can either have light sources such as light emitting diodes attached to it, or it can be scanned by an external focused light source. An early example of the concept was patented by Ketchpel in 1964 [9]. A example of a contemporary system is being developed at the United States Naval Command, Control and Ocean Surveillance Center [10]. Space-filling displays of this type will always face the challenges of mechanical scanning.

A final type of volumetric display fills the volume with a display medium that can be excited externally to the point where it emits light. This external stimulus can, for instance, come from lasers of different wavelengths that are scanned through the imaging medium. Finding display materials with the appropriate nonlinear optical properties has proved to be a great research challenge. Recent progress has been made at Stanford by Downing[3], incorporating principles described by Lewis [11]. The ideal display material, not yet discovered, must combine the qualities of optical

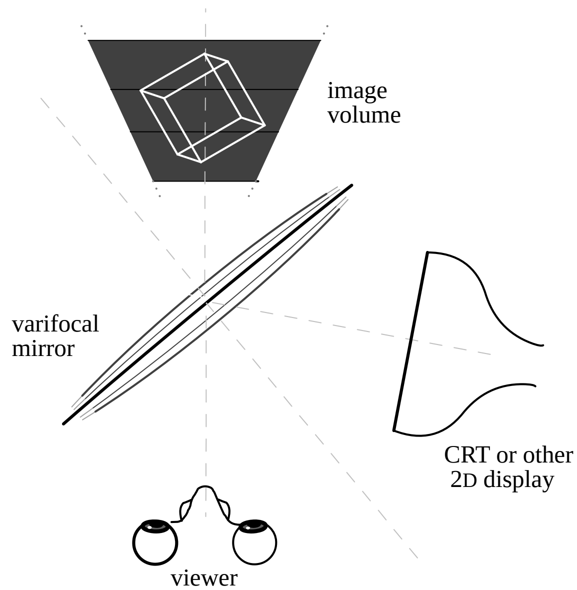

Fig. 3: A varifocal mirror display.

efficiency, low cost, and light weight to order to find widespread use.

Independent of the optical and mechanical technology, volumetric displays share several properties. First, the image they present is visible from a wide range of viewpoints, even permitting a viewer or group of viewers to walk all the way around the display. Second, the displayed image emits a continuous, uniform, spherical wavefront centered at each displayed point. The human eye can selectively focus on this wavefront, providing the sense of ocular accommodation. On the other hand, because the wavefront is uniform and omnidirectional, view-independent shading of objects is not possible. Even more critically, current displays do not exhibit arbitrary occlusion of one part of the image volume by another. For photorealistic scenes, occlusion is almost always the most important depth cue, much more important than ocular accommodation and often stronger even than stereopsis. Volumetric displays thus have most common use in non-photorealistic applications such as wireframe images and icon-based displays. Volumetric devices are appropriate for this type of application because they can vector-scan only the regions of space that the object spans, eliminating the display bandwidth that would otherwise be required to rasterize the entire image volume.

A similar technology to volumetric displays, stereolithography, has achieved widespread use in the CAD/CAM industry because the result of the scanning process is not the emission of light, but a formation of a solid computer-generated object made from a polymer material. Also related to volumetric displays is VOXEL Corporation’s VOXBOX holographic display system for radiology [5]. The VOXBOX uses a hologram to display stacks of tomographic images as an image volume. It shares many imaging properties of volumetric displays.

Parallax displays

Parallax displays consist of a surface covered with display elements that can emit light of varying intensity in different directions. The plain piece of glass “display” is a good way of thinking about parallax displays: the front surface of the glass is a continuum of sites that send out a hemisphere of light varying in intensity. A single output site radiates only information captured from one viewpoint, the corresponding site on the back surface of the glass. While this site's information contains no explicit notion of depth, the light emitted from several sites considered as a whole presents a three-dimensional image. Depending on a viewer's exact location, light traveling in different directions appears to originate from different parts of the glass. This visual information, intercepted by the viewer's two eyes, is processed to form a three-dimensional mental model of the scene.

The plain glass display analogy also demonstrates that parallax display devices can correctly show arbitrary occlusion of one part of an object by another. Occlusion is essential for the display and comprehension of photorealistic synthetic scenes. On the other hand, many types of parallax displays use information reduction techniques that approximate the shape of the wavefront of light-emitting points. These approximations diminish or eliminate any ocular accommodation depth cues.

Holograms

Holographic displays [1][17] are in many ways very close to the "piece of glass" display model. Holograms store wavefront information about an object as microscopic interference fringes during the holographic exposure process. When the developed hologram is illuminated, its interference fringe pattern acts as a complex diffractive lens that reconstructs the object light's direction and intensity. A holographic display reconstructs light so exactly that it shares many of the properties of a volumetric display, including providing ocular accommodation cues, without having to physically span the imaging volume. Unfortunately, the data bandwidth of a high quality display hologram is far beyond any current or envisioned image synthesis technology: a typical display has a spatial frequency exceeding 1500 line pairs per millimeter. In order for a hologram to be synthesized, the information that is contains must be reduced to manageable quantities.

Unfortunately, the alluring and eerie realism of holographic displays has lead to an increasingly common misuse of the term "hologram" to describe any display that is vaguely three-dimensional, and even some that are not. To be clear, display holograms are image-bearing diffractive optical devices. Other displays may be three-dimensional, but they are not holograms.

Parallax Barrier Displays

Several parallax display technologies were developed long before holography. The earliest were the parallax stereogram [6] and the parallax panoramagram [7]. Both display types depend on a device called a parallax barrier, an opaque material slotted with a series of regularly spaced vertical slits. A piece of film or other imaging medium is offset some distance behind the parallax barrier. Each slit in the barrier acts as a window onto a stripe of the section of film that lies behind it. Exactly which stripe of film is visible depends on the horizontal angle from which the slit is viewed. A parallax stereogram displays a stereoscopic image pair by interleaving columns of the two images on the film, one column of each image per slit. An appropriately-positioned viewer will see the right view of the pair through the slits with the right eye, the left view with the left. The parallax barrier blocks the opposite image from view. A stereoscopic image is thus produced.

The parallax panoramagram extends this concept by introducing thinner columns of more views behind each slit. Art^{n}'s PSCHologram is an example of this technology [16]. The more views that are present, the more naturally the image will appear to change as the viewer moves from side to side. Figure 4 shows a parallax panoramagram. Parallax panoramagrams are limited because the barrier, while necessary to block the unwanted views, also blocks light from getting to the viewer. Panoramagrams usually require banks of bright, diffuse lights located behind the film. Displays of this type rely on the fact the spatial and directional information is spatially multiplexed onto the film, which leads to

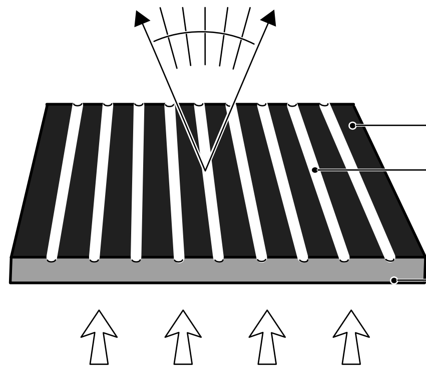

Fig. 4: A parallax panoramagram

several other problems. First, a viewer positioned far enough to the left or right of the display will be able to look through one slit to see the image data associated with the slit's neighbor. As a result, the image appears to repeat its perspectives as the viewer moves. If the viewer sees a correct view with one eye and repeated view with the other, the depth of the object can even appear to flip inside out (called a psuedoscopic image). Second, the resolution of the film limits the maximum number of views that can be displayed. The spacing of the slits determines the maximum spatial resolution of the display.

The parallax panoramagram is three-dimensional only in the horizontal direction; vertically, the image of the display behaves just as if it were a flat photograph. As a viewer moves closer or further from the display, vertical edges of the image will appear to shift naturally with respect to each other, just as they would in a real object. Horizontal edges, though, remain fixed relative to each other. This kind of display is said to be horizontal parallax only, or HPO. HPO displays are a useful engineering trade-off because they greatly reduce the information content of a three-dimensional image while still displaying stereoscopic and motion parallax information. For demanding applications, the view limitations and possible distortions of HPO would preclude its use in favor of a full parallax display.

Lenticular sheet displays

The word "lenticule" is a synonym for "lens", but the term "lenticular" has come to refer to a type of three-dimensional display that using an array long, narrow lenses instead of slits to display three-dimensional information. More correctly, this display type should be referred to as a lenticular panoramagram. Figure 5 depicts a lenticular panoramagram.

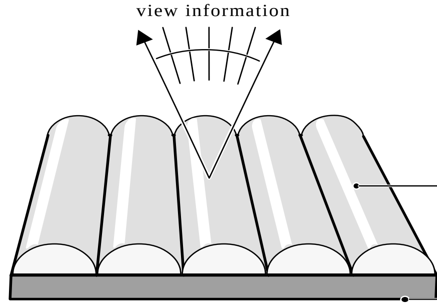

Fig. 5: A lenticular panoramagram.

This display type is functionally very similar to the parallax panoramagram. Each lens focuses on the image information located behind it and directs the light in different directions. If we think of one slit of a parallax panoramagram as similar to a camera with a pinhole as an aperture, one lenslet of a lenticular panoramagram is analogous to a camera lens. Cameras with lenses are more common than pinhole cameras because they collect more light from the scene. Similarly, a lenticular panoramagram is brighter and optically more efficient than the corresponding parallax panoramagram. The entire surface of the lenticular sheet radiates light; there are no dark stripes such as those produced by a parallax barrier.

Continuing the camera analogy, camera lenses come in a wide variety of focal lengths and materials and can be adjusted for focus, while the only adjustable variables of a pinhole camera are the pinhole diameter and the spacing between the pinhole and the film. Lenticular sheets are molded from plastic in a process that sets the width of the lenslets, the distance between the image and the lens, and each lenses' optical power. The quality of the sheet-making process also determines the optical aberrations that will be manifested in the final image. The optical power of the lens controls the angle of view through which the final image can be seen. Lenticulars are almost always made so that the film plane is located one focal length behind the lenses: the image data emerges collimated from each lenslet. Making high quality yet affordable lenticular sheets is one of the major difficulties of creating lenticular sheet displays.

Lenticular panoramagrams can also be used with a CRT or other two-dimensional display device to produce a dynamic three-dimensional image. The spatial resolution of the two-dimensional display must be high enough in the horizontal direction to provide both spatial and directional information. Optical alignment of the underlying display with the lens sheet is essential to producing distortion-free three-dimensional image. As always, the information bandwidth of the display increases as more directional information is added.

Like parallax panoramagrams, lenticular panoramagrams display only horizontal parallax. Another display type, the integral photograph or integram, uses spherical lenses instead of cylindrical

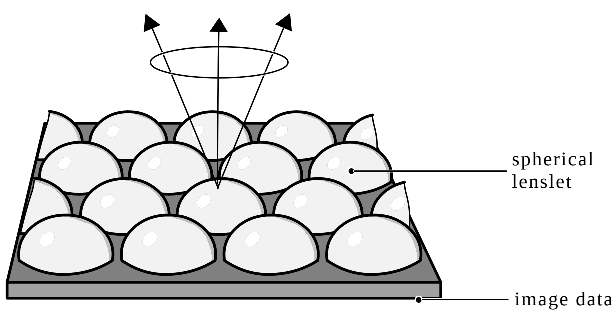

Fig. 6: An integral photograph or integram.

ones to present horizontally and vertically varying directional information, thus producing a full parallax image. Figure 6 shows the integram's spherical lens array. Integrams are less common than their cylindrical lensed counterparts mostly because even more of their spatial resolution is sacrificed to directional information.

Holographic stereograms

The optical systems used in parallax and lenticular panoramagrams are tightly constrained by their how much information can be stored on their imaging medium. They are also prone to image repeating caused by crosstalk between adjacent display elements. The holographic stereogram overcomes these problems by combining the information storage capacity of the hologram with the information-reducing image discretization of the panoramagram [2]. (As Okoshi notes, the holographic stereogram should rightly have been called a holographic panoramagram [15].) Instead of macroscopically encoding directional information as interleaved stripes of image, a holographic stereogram optically records the same information microscopically as fringe patterns.

The details of the holographic stereogram's recording process are beyond the scope of this text. The basic idea of the process is to make a series of holographic windows or aperture through which two-dimensional projections of image data can be seen. Each aperture is a distinct optical element; this property eliminates the problem of element crosstalk and image repeating. "One-step" holographic stereograms are recorded directly onto the final display media using spatially discrete slit apertures. One-step holograms are useful for rapidly produced single-copy images. "Two-step" holographic stereograms use a spatially discretized master hologram to record the image information. The master hologram can be transferred to make many final holographic images. The two-step process is most frequently used when multiple copies of the hologram are desired.

Holographic stereograms can be recorded in a number of media and are suitable for stamping and publication. The holographic stereogram's major weakness is the complexity of its exposure apparatus and the difficulties of lighting and color that result from the hologram's diffractive properties.

Electro-Holography

The fact that display holograms are even possible is directly due to the fact that the physical process of interference that forms the holographic fringe pattern is a parallel optical computation. Electro-holography, the technology of forming the same type of fringe pattern electrically, does not have the benefit of this natural computer. Instead, the fringe patterns for electronically-generated holograms must be conventionally computed and then output to a physical device capable of diffracting light and producing the three-dimensional image. Electro-holography depends heavily on information reduction techniques such as the elimination of vertical parallax and holographic stereogram-like discretization of spatial and directional information. Recent progress on developing both the display technology and the computational algorithms used to compute the fringe patterns has been encouraging [18] [12], but the technology is still far from producing high quality three-dimensional images using affordable hardware.

GUIDELINES FOR PARALLAX DISPLAYS

At the current time, parallax displays are the three-dimensional displays most commonly used with computer-generated images. Parallax displays are popular for three primary reasons: they can be published and mass-produced, they can be made in a wide range of sizes, and they can produce photorealistic images. For these reasons, these basic guidelines for generating effective spatial images will concentrate on parallax displays. Here are some important things to consider when undertaking the process of creating three-dimensional displays.

Produce effective imagery.

Stereopsis is only one of the cues we use to evaluate an object's dimensionality. Without stimulation and agreement of all depth cues, spatial images can appear uninteresting or even visually confusing. Strong occlusion, shading, and perspective depth cues are even more important in three-dimensional images than in 2D renderings and photographs.

Know the display technology. Respect its limitations.

Some day, software packages will take care of the details of matching rendering to a display's characteristics, allowing the designer to concentrate on the content of the image. Unfortunately, that day is not yet here. Each display type has its own properties, and each manufacturer has their own particular process. These constraints simply must be respected in order to produce undistorted, high quality images.

If you are involved in your first three-dimensional imaging project, it is best to work either directly with the manufacturer or with a designer who has experience creating such images. They will able to help you plan your imagery to avoid common problems such as window violations and conflicting depth cues. They will most often also be able to provide correct parameters for a particular display.

The process of rendering for most parallax displays consists of moving a computer graphics camera along a track in front of the virtual scene, capturing images at regular intervals. The varying camera viewpoints provide the spatial and directional information that is encoded into the display when it is made. Here are some basic camera and image parameters required for generating parallax displays:

- Size and position of the camera track or grid.
- The camera parameters or transformation matrix for the computer graphics camera view.
- The number of perspective images that must be rendered.
- Pixel resolution and image format of the image data.

Horizontal parallax only images have a specific viewing distance where the varying horizontal and fixed vertical perspectives of the object match. Viewers not located at the correct distance will see images with cylindrical lens distortion. To minimize this problem, be sure to match the computer graphic camera's view distance to that specified for the display.

In addition, all parallax displays impose inherent limits on the spatial resolution of the image volume that depends on image depth [4]. For example, the horizontal resolution of a parallax panoramagram is limited by both the barrier slit width and the width of the vertical image slices on the film. Similarly, the horizontal spatial resolution of a lenticular panoramagram cannot exceed the width of a lenslet at the display surface. Lenticular sheets are also limited by lens blur and diffraction. Exceeding these limits produces images that break up into horizontal, jaggy pieces. Many manufacturers may be only anecdotally familiar with these limitations.

Model the imaging process; match the graphics to the display.

If you are trying to get the best possible spatial image, you must go beyond these basic guidelines and actually model the imaging process of image synthesis, recording, and display. The computer graphics image data must match where the light from the final display actually goes. One of the most difficult part of this process is not the computer graphics algorithms, but actually measuring and understanding the optical characteristics, geometry and distortions of the display device. Correct matching of graphics and display can produce significant gains in image realism and lucidity, but simplifying this process remains an area of continuing research and development.

CONCLUSION

A range of autostereoscopic displays exist today that can be used in applications ranging from advertising tools to air traffic control consoles and publishable images for scientific visualization. As the demand for such displays increases, the 3D and computer graphics fields face the challenge of demystifying three-dimensional technology and simplifying the image generation process. Sensible image design, selection of an appropriate display device, and adherence to its limitations can yield realistic, understandable, and uniquely effective three-dimensional images

REFERENCES

[1] Benton, Stephen A. Display Holography: an SPIE Critical Review of Technology. Proc. SPIE, Holography, A86-32351 14-35, (1985), pp. 8-13.

[2] Benton, Stephen A. Survey of Holographic Stereograms. Proc. SPIE, Processing and Display of Three-Dimensional Data, 367, (1982), pp. 15-19.

[3] Downing, Elizabeth et. al. A Three-Color, Solid-State, Three-Dimensional Display. Science 273, 5279 (August 30, 1996), pp. 1185-1189

[4] Halle, Michael W. Holographic Stereograms as Discrete Imaging Systems. Proc. SPIE, Practical Holography VIII, 2176, (May 1994), pp. 73-84.

[5] Hart S.J. and M.N. Dalton, Display Holography for Medical Tomography, Proc. SPIE, Practical Holography IV, 1212 (May 1990), pp. 116-135.

[6] Ives, F. E. U.S. Patent No. 725,567, (1903).

[7] Ives, H. E., A Camera for Making Parallax Panoramagrams, J. Opt. Soc. Amer., 17, (Dec. 1928), pp. 435-439.

[8] Kawamoto, Wayne 3-D Images That Float in Air. Byte (Oct. 1995).

[9] Ketchpel R. D. U.S. Patent No. 3,140,415, (July 7, 1964).

[10] Lasher, Mark et. al. Laser-Projected 3D Volumetric Displays. Proc. SPIE, Projection Displays II, 2650 (1996), pp.285-295.

[11] Lewis, J. D. et al. A True Three-Dimensional Display. IEEE Transactions on Electron Devices, ED-18, 9, (Sept. 1971), pg. 724-732.

[12] Lucente, Mark and Tinsley A. Galyean. Rendering Interactive Holographic Images. Proceedings of SIGGRAPH '95 (Los Angeles, CA, Aug. 6-11, 1995). In Computer Graphics Pro ceedings, Annual Conference Series, 1995, ACM SIGGRAPH, New York, 1995, pp 387-394.

[13] McAllister, David F. Stereo Computer Graphics and Other True 3D Technologies. Princeton Univ. Press, 1993.

[14] McKenna, Michael and David Zeltzer. Three Dimensional Visual Display Systems for Virtual Environments. Presence, 1, 4 (Fall 1992), pp 421-458.

[15] Okoshi, Takanori. Three-Dimensional Imaging Techniques. Academic Press, New York, 1976.

[16] Sandin, Daniel J., et. al. Computer-Generated Barrier-Strip Autostereography. Proc SPIE, Non-Holographic True 3D Display Technologies, 1083, (Jan. 1989), pp. 65-75.

[17] Saxby, Graham. Practical Holography, 2nd edition. Prentice Hall, Dec. 1994.

[18] St.-Hilaire, Pierre et. al. Scaling Up the MIT Holographic Video System. Proc. SPIE, Fifth International Symposium on Display Holography, 2333, (Feb 1995), pp. 374-380.

[19] Starks, Michael R., Stereoscopic video and the quest for virtual reality: an annotated bibliography of selected topics. Proceedings Stereoscopic Displays and Applications II, 1457. SPIE, pp. 327-342, Aug. 1991.

[20] Starks, Michael R., Stereoscopic video and the quest for virtual reality: an annotated bibliography of selected topics. Proceedings Stereoscopic Displays and Applications III, 1669. SPIE, pp. 216-227, June 1992.

[21] Traub, A.C. Stereoscopic Display Using Varifocal Mirror Oscillations. Applied Optics, 6,6 (June 1967), pp. 1085-1087.