To appear in SPIE Proceeding #2176 “Practical Holography VIII” (SPIE, Bellingham, WA, February 1994) paper #10 (in press).

Holographic stereograms as discrete imaging systems

Michael W. Halle
MIT Media Laboratory
20 Ames St.
Cambridge, MA 02139
USA

ABSTRACT

Unlike holograms of real objects, holographic stereograms consist of information recorded from a relatively small number of discrete viewpoints. As discrete imaging systems, holographic stereograms are susceptible to aliasing artifacts caused by insufficient or improper sampling. A characterization of sampling-related image artifacts in holographic stereograms is presented. Constraints on image extent and resolution imposed by sampling are outlined. Methods of reducing or eliminating aliasing artifacts in both photographically-recorded and computer-generated holographic stereogram images are described. Results of this analysis can be generalized to describe other autostereoscopic displays.

1. INTRODUCTION

Holographic stereography is the most widely used holographic technique for producing three-dimensional imagery from two-dimensional views. Holographic stereograms are the result of the merger of two approaches to three-dimensional imagery: display holography, with its roots in the work of Leith and Upatnieks [7], and “traditional” autostereoscopy used in much older lenticular and parallax barrier display technologies. Holography provides a medium with almost unlimited ability to record the appearance of light, while stereoscopy yields a way to reduce the huge amount of information needed to produce a visually acceptable three-dimensional image.

Because of their heritage, holographic stereograms are most often analyzed using stereographic terms such as “taking camera”, “view zone” and “view distance”: the display presents information captured by a camera (either an actual photographic device or a synthetic analog) to a view zone some distance from the display. This approach to analysis is closely related to the process of making the display. In general, this “view-based” paradigm of both traditional and holographic stereoscopy offers a straightforward way to understand the appearance and distortions of a three-dimensional autostereoscopic image.

However, some image properties and readily observable image artifacts are more difficult to explain solely using a view-based analysis. For example, some existing holographic stereograms, particularly very deep ones, exhibit stripes or banding in certain areas of the image. The cause of this artifact and possible methods for its prevention and cure are not immediately evident from the characteristics of the views that compose the image. In this paper, alternative paradigms for examining stereograms, such as wavefront analysis and sampling theory, are combined with the more traditional view-based approach in order to provide a more comprehensive explanation of stereogram characteristics and image artifacts. For simplicity, however, this analysis is limited to purely geometrical optics. Diffractive effects are considered in another, more general work [5], and appear significant mostly for small view aperture sizes.

2. A SIMPLE HOLOGRAPHIC STEREOGRAM MODEL

In this section, we define a simple holographic stereogram model as a basis for further discussion. Although this stereogram represents a particular class of display (specifically, a two-step horizontal parallax only (HPO) holographic stereogram), the results drawn from an analysis of this type are applicable to a wide range of holo-stereographic displays and geometries. The simple two-step model follows the work of deBitetto [3], with modifications by Benton [1], and is shown in Figure 1. The master hologram (H1) consists of a series of narrow vertical slit holograms, each of width W_{slit}, which are sequentially exposed to reference and object beams. N_{slits} is the number of slit holograms. The total width of the H1 is then $W_{H1} = N_{slits} \cdot W_{slit}$.

The “object” to which the slits are exposed is a ground-glass diffusion screen located a perpendicular distance D_{view} away. Onto this diffusion screen, a series of perspective views is back-projected. These perspective views are captured by a camera, moving along a straight track, with the optical axis of the camera’s lens perpendicular to the direction of camera travel. The camera’s lens stage translates horizontally to recenter one particular depth plane of the scene (known as the *image plane*). The geometric relationship between the image plane and the capture camera corresponds one to one to that between the diffusion

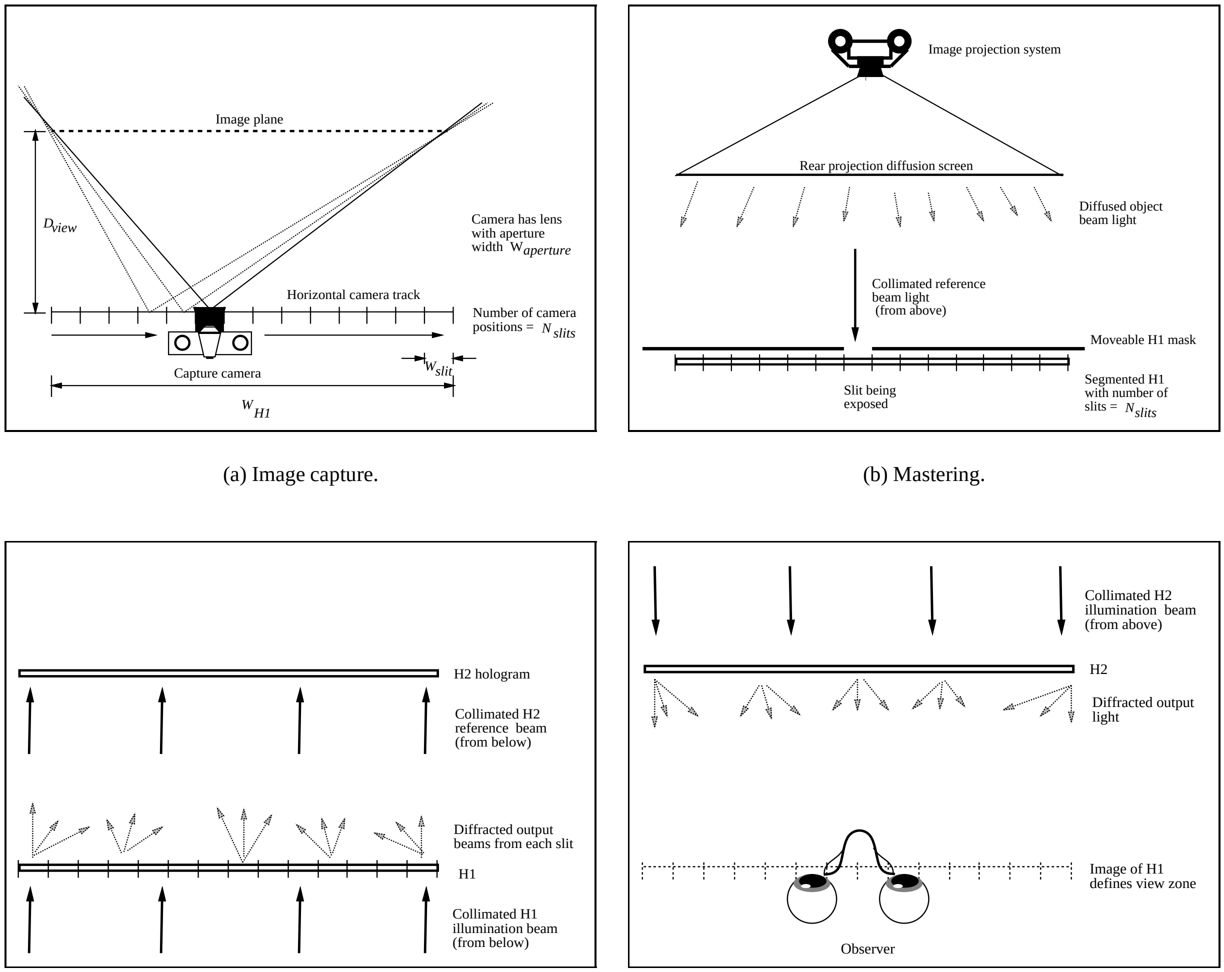

(a) Image capture.

(b) Mastering.

(c) Transferring.

(d) Viewing.

Figure 1: A simple holographic stereogram geometry.

screen and the holographic slit apertures. The lens of the camera has an aperture of width $W_{aperture}$. The one-to-one scale between the image and the display implies that the camera moves a distance W_{slit} between views.

The holographic recording process consists of unmasking each slit one at a time and projecting the corresponding perspective view onto the ground glass. The reference beam for each slit is provided by an collimated, off-axis source. The H1 that results from this process contains information about all the different perspective views of the object that were captured by the horizontally scanning camera. To simplify viewing, this master hologram can be transferred without distortion, astigmatism, or magnification using a phase-conjugate collimated illumination source.

The transfer hologram (H2) is located at a distance D_{view} away from the H1, positioned at the plane of the aerial images of the diffusion screen projected by the slits of the master. The H2 is referenced with a beam phase-conjugate to that of the final illumination source, so that the transfer hologram will in turn project an unabberated and unmagnified image of the slits of the H1 to a position D_{view} away from the H2. This distance is the *view distance* from which the viewer observes the hologram. The size of the *view zone*, the region of space where image information from the H2 can be be observed, is precisely the size of the H1 (W_{H1}).

3. SINGLE POINT IMAGING

Insightful analysis of the simple holographic stereogram model can best be achieved by reducing to a single point the scene captured by the taking camera. A single point object sends out light in a spherical wavefront with the point at its center. The curvature of the wavefront is greater near the point and becomes more shallow further away.

When a hologram is made of a point object, each part of the holographic medium forms an interference pattern capable of diffracting light from an illumination beam into a wavefront of the same shape as that of the original point. This wavefront, being spherical, has identical curvature horizontally and vertically. The exact curvature of the wavefront depends on the distance from the point to the hologram. In this way, the hologram records depth information about the point.

A human observer analyzes wavefront shape in three ways. First, the lens of the eye modifies the incoming wavefront and focuses it onto the retina. This focusing is called accommodation. Second, the observer's eyes sample the wavefront from two slightly different positions. Light from the point appears to be coming from two different directions, represented by different positions in the two visual fields. This image disparity is the basis of stereopsis. Finally, a moving observer samples the wavefront at many different positions: the direction of light from the object and thus the object's position in the visual field changes as a result of the observer's motion. Motion parallax is the term used to describe this means of wavefront analysis. These techniques for determining wavefront shape are primarily used by humans to determine the three-dimensional location of an object.

If a single image of the point is captured by a camera and projected onto a diffusion screen, a spherical wavefront is emitted with a center located at the surface of the screen where the image of the point appears. Information about the point's depth is lost: the image of the point is identical independent of the separation between point and camera. Regardless of the curvature of the wavefront of the original point, the wavefront of the point's image will be centered somewhere on the plane of the diffusion screen. The shape of this wavefront, however, remains spherical.

If a hologram is made of the diffusion screen, it too must record and play back only spherical wavefronts diverging from the diffusion screen plane. This hologram could be made either conventionally by exposing the entire surface of the holographic plate at once, or by exposing it sequentially, section by section. As long as the exposure is made carefully to assure stability and uniform exposure, a hologram produced by either method should faithfully reproduce the shape of the point image's wavefront.

A spatially multiplexed holographic stereogram is one example of a sequentially exposed hologram. Exposing each slit of the stereogram to the same object defeats the purpose of slit exposure, of course. Stereograms are exposed instead to a sequence of perspective views, one view per slit. The simplest such sequence is a series of identical images of a single object point. In this case, the shape of the resulting reconstructed wavefront is the same as for a true hologram of the diffusion screen. However, because the effective path length of the object beam changes slightly as it passes through the projection mechanism during the exposure of different views, the phases of the different segments of the wavefront reconstructed by the slits are relatively random and thus mutually incoherent.

A holographic stereogram made with identical views for each slit can only record and display points located at the image plane. In contrast, the perspective views taken of a point located at a plane other than the image plane will not be identical; rather, the image of the point will change position from one perspective view to the next. The position of the point is assumed to be fixed during camera capture, so any motion of the image of the point from view to view is due solely to parallax and not to actual motion of the point itself. When the stereogram is viewed, the observer is presented with a perception of the point composed of all of the discrete perspective images of that point. The observer's position selects which point image is visible.

Because only a discrete number of views is recorded into the slits of the holographic stereogram, the display only approximates the appearance of the original point. The original point appears to move continuously as a function of viewer movement, while the stereogram's image of the point jumps in position from one perspective view to the next. Figure 2 shows the relationship between the motion of the observer, of the actual point, and of the displayed image of the point. Only from an observer location at the middle of each slit does the position of the stereogram's image of the point match that of the point itself. For all other observer locations, the stereogram's approximation is not exact.

The stereogram's approximation can also be analyzed using a wavefront-based approach. The wavefront of the original point is approximated using small pieces of wavefronts emitted by point images on the image plane. Each slit aperture of the stereogram contributes one piece to the output wavefront. This piecewise wavefront approximation has the gross curvature of the wavefront of the original point. However, curvature for small portions of the wavefront is that of a point on the image plane.

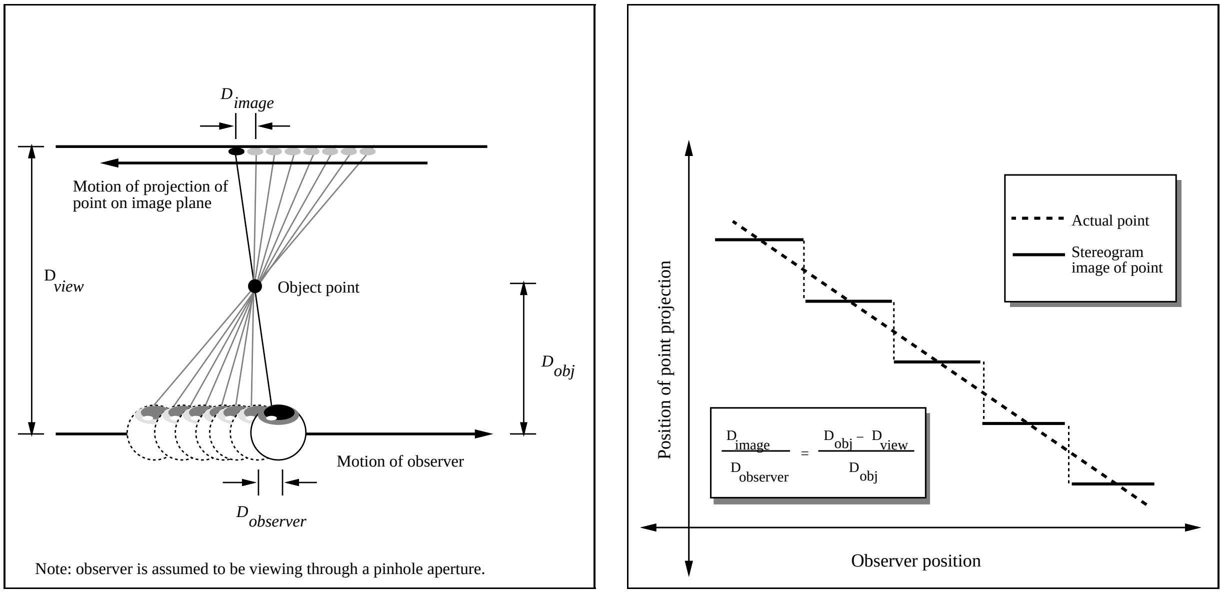

Figure 2: The relationship between observer motion, the apparent motion of an object point, and the apparent motion of a stereogram image of the point.

Points closer than the image plane will be composed of pieces with too shallow a curvature; points further away will have wavefront approximations that bulge from section to section. Figure 3 shows the shape of the two wavefront approximation. Furthermore, the wavefront approximation is not smooth, but instead has discontinuities where the edges of two adjacent wavefront sections meet. These discontinuities correspond to viewing locations where the observer sees the image of the point jump from one position to another.

This stereogram wavefront approximation only applies in the horizontal direction. Vertically, the image of a point is viewed through the same slit independent of the observer location. The wavefront is emitted by only one slit; the shape of the wavefront is unaffected by the stereogram slit structure. Therefore, in the vertical direction, the point will appear to be located at the image plane, independent of its original position.

4. VISUAL ARTIFACTS IN HOLOGRAPHIC STEREOGRAMS

Traditionally, stereogram holographers have used simple criteria to determine W_{slit}, the width of the stereogram's slit apertures. Too wide a slit produces an image that jumps noticeably from one frame to the next. Too narrow a slit complicates stereogram production by increasing the number of perspective views and holographic exposures. One commonly used compromise is to set W_{slit} equal to the width of the pupil of the human eye (approximately 3mm). As the observer moves in front of such a stereogram, the eye is always entering a new slit, giving the appearance of a continuously changing image. If the observer's eye straddles the boundary of two slits, the resulting presented image is a linear combination or blending of the perspective views recorded in those slits. The weighting of the two images depends on how much of each slit falls into the observer's eye.

Figure 4 compares the visual appearance of an actual point to a stereogram image of the same point as an observer slowly moves from one slit to the next. Unfortunately, for an arbitrarily small point lying relatively far from the image plane, the stereogram approximation differs significantly from the appearance of the actual point. Specifically, a gap exists in between the images of the point in the two views. Rather than consisting of a single point, the image is broken up into a line several discontinuous points. The gaps between the point images consist of extraneous image data. The display looks as if a horizontally-smeared image of the point is moving behind a picket fence-like pattern located on the image plane. Because this pattern does not lie at the same depth as the original point, a depth conflict may occur, diminishing the intended three-dimensional effect of viewing the stereogram.

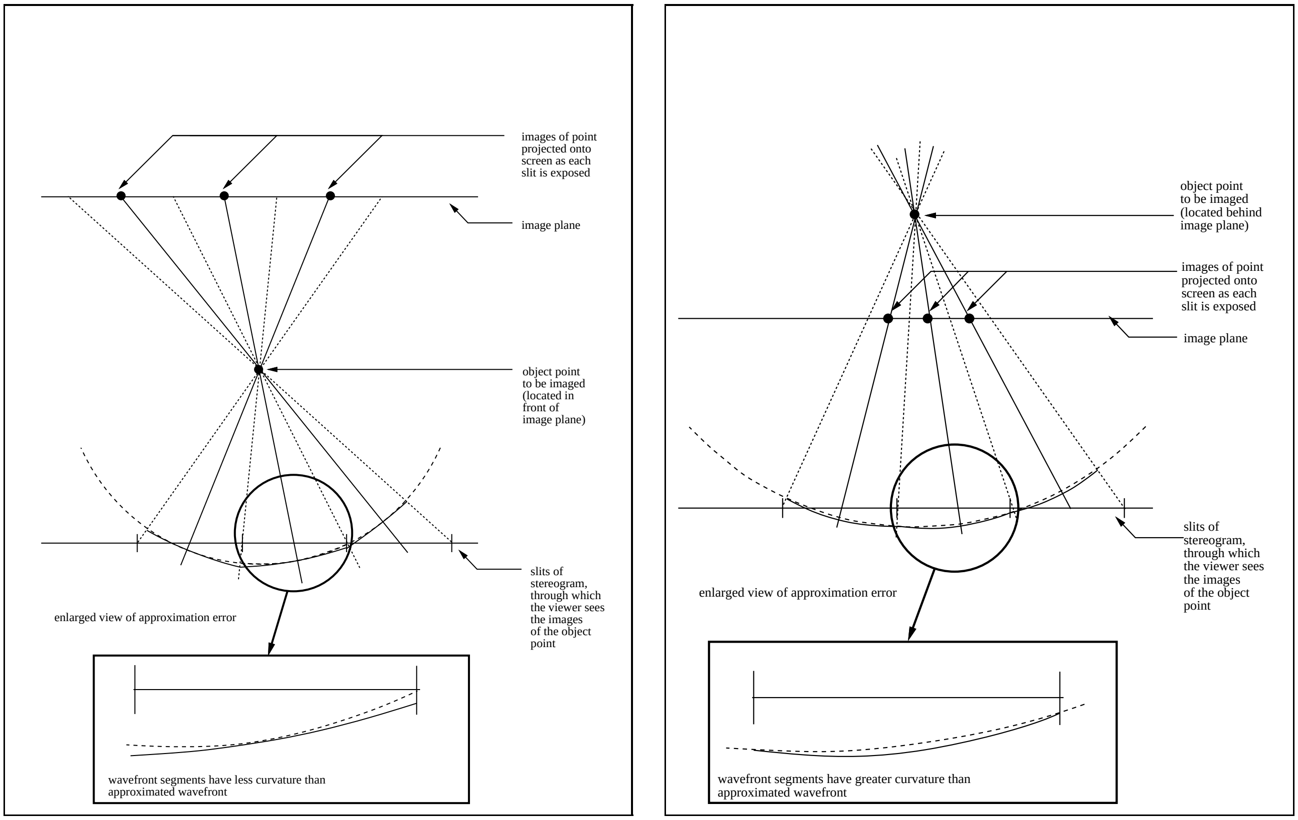

(a) Object in front of image plane.

(b) Object behind image plane.

Figure 3: Stereogram wavefront approximations for object points not located on the image plane.

The stereogram's discrete nature is responsible for this artifact. Sampling theory provides some insight into the problem. Discrete systems are prone to signal errors when their sampling rate drops below twice the maximum frequency of the continuous signal they record. This minimum frequency of sampling is known as the Nyquist rate. Aliasing artifacts, the result of insufficient sampling, are low frequency, periodic errors in the output signal.

In the stereogram, the parallax of a point is sampled into discrete viewpoints. The low frequency gaps that break up the image of the point are aliasing artifacts that occur when the number of perspective views is too small. Proper sampling for a given point depends on the point's size (its spatial frequency) and the distance is moves from view to view (its spatio-temporal frequency). This second parameter, in turn, depends on the point's distance from the image plane. Because aliasing in stereograms are produces by the improper combination of image data from adjacent views, it is called *inter-perspective* or *inter-view* aliasing.

5. EXAMPLES OF INTER-PERSPECTIVE ALIASING

Inter-perspective aliasing artifacts are not uncommon in holographic stereograms. This section presents three examples of stereograms with noticeable aliasing artifacts. The first, *Wedding Portrait of Bill Parker and Julie Walker*, is shown is Figure 5. The perspective views for this image were captured using a Mitchell movie camera traveling on a linear track approximately 3 meters long. The camera did not have a translating lens stage; the holographic printer geometry was adapted to use non-recentered perspective views at the loss of spatial resolution. Approximately 100 perspective views are incorporated into the stereogram. The closeup of the image shown in Figure 5(b) shows that while the main subjects are well-portrayed, some background details are repeated horizontally several times. The objects in the background are located significantly behind the image plane of the stereogram.

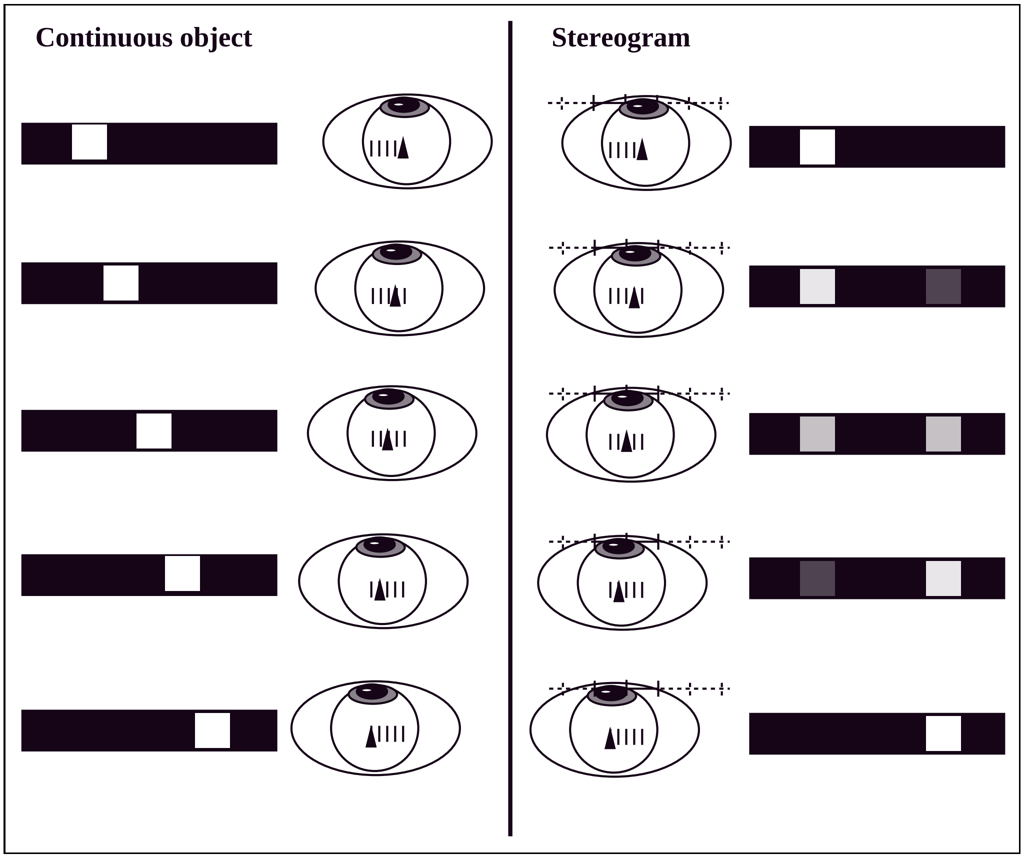

Figure 4: The visual appearance of a holographic stereogram of a point object compared to the appearance of the point itself.

(a) Entire image.

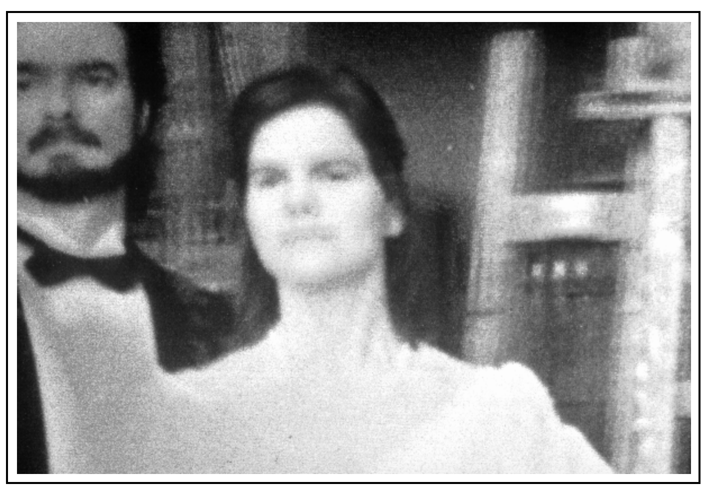

(b) Detail, 105mm lens, f/22.

Figure 5: *Wedding Portrait of Bill Parker and Julie Walker*, Michael Klug, MIT Media Laboratory. Holographic stereogram generated from photographically-recorded images, approximately 100 views.

The next two examples of aliasing artifacts come from computer-generated holographic stereograms. Computer graphics permits greater freedom of three-dimensional image composition and capture compared to photograph-based methods of stereography, but also permits points and lines of high resolution to appear at any depth plane without the filtering effects of lens blur and other optical degradations. As a consequence, deep computer generated holographic stereograms almost always have some sort of aliasing artifact. William J. Molteni's *Star Wars* holographic stereogram (Figure 6) is one of the earliest examples of computer generated images to use distortion compensation to minimize perspective distortion. This distortion compensation permitted the creation of a very deep image. The low resolution of the perspective images used to make the stereogram prevents aliasing in all parts of the image except those closest to the viewer. Figure 6(b) shows an enlargement of the edge of the frontmost circle. (The circle is elliptical because the photograph was taken with a camera positioned closer to the hologram than the intended view distance.) Small gaps exist in the most-vertical lines of the circle, but disappear as the line becomes more horizontal.

Finally, Figure 7 shows a photograph of *Lunar Lander* by Michael Teitel. This image is a computer generated image of a space shuttle landing onto a runway at a fictional moon base. This stereogram, one of the earliest to display an animated, shaded image, consists of approximately 100 perspective views each with a resolution of about 512 pixels square. This resolution is sufficiently high to produce noticeable sampling artifacts. Figure 7(b) shows a closeup of the stars in the background of the image. The stars are good examples of single point objects. The gaps in between the images of each star are clearly visible. For this example, and for the closeups of the other two stereograms, the lens of camera used to take these pictures was closed down to approximate the size of the pupil of the human eye. In Figure 7(c), the lens aperture was opened to permit the light from more slits to be recorded. The line of points representing each star's image widens because the images from more perspective views are visible to the camera. However, the presence and size of the gap between the points in a line is unaltered by changes in the viewing aperture. The end of the runway also exhibits aliasing.

6. CHARACTERISTICS OF INTER-PERSPECTIVE ALIASING

The exact appearance of inter-perspective aliasing artifacts is dependent on the perspective sampling rate. As the number of perspective views increases, more perspectives become simultaneously visible to the observer, filling in the gaps between the existing views. Figure 8(a) shows the appearance of the same point as in Figure 4, but as imaged by a stereogram with twice as many slits. At some level of sampling, the images of the point from all the contributing views will abut each other with no gaps. This sampling level represents the Nyquist rate for the given point and depth.

Alternatively, aliasing artifacts can be avoided by increasing the horizontal extent of the point so that the image of the point from one view abuts the image from the next. This process lowers the spatial frequency of the point's detail to a level that can be recorded without artifacts by the holographic stereogram. In the signal processing world, the process of frequency limiting is known as *filtering*, an operation which results in a signal said to be *bandlimited*. Figure 8(b) depicts the result of viewing a bandlimited point image.

From a wavefront analysis, bandlimiting creates not just one, but a series or range of wavefronts for each image point. These wavefront ranges have the property that the ranges for a single point in two perspective views have continuously varying curvature. In other words, the wavefront from the left side of one point image has the same curvature as the right side of the point image in an adjacent view. The bandlimiting process produces an incoherent sum of wavefronts that has no discontinuities in curvature. Figure 9 shows the result.

An important characteristic of all aliasing artifacts is that while they can be prevented, they cannot, in general, be removed from a signal once they are present. This characteristic is also true of inter-perspective aliasing. It would seem, for instance, that the collection of point images that should join together to represent a single point could be smeared together by using an extended-length horizontal light source to illuminate the hologram. Sadly, the gap between the point images is located on the image plane; as such, it is subject to neither source-size nor chromatic blur. Since post-processing removal of aliasing artifacts is not possible, steps must be taken in image creation to maintain sufficient sampling rates.

7. QUANTITATIVE BANDLIMITING

In order to perform proper bandlimiting for alias-free holographic stereograms, the resolution for detail located at a given distance from the image plane must be determined. For one specific depth plane, the image plane, the stereogram approximation is exact; therefore, detail of arbitrary resolution can be displayed at the image plane without concern for inter-perspective aliasing. The images of points located at all other depths must have a horizontal extent equal to the distance that the point

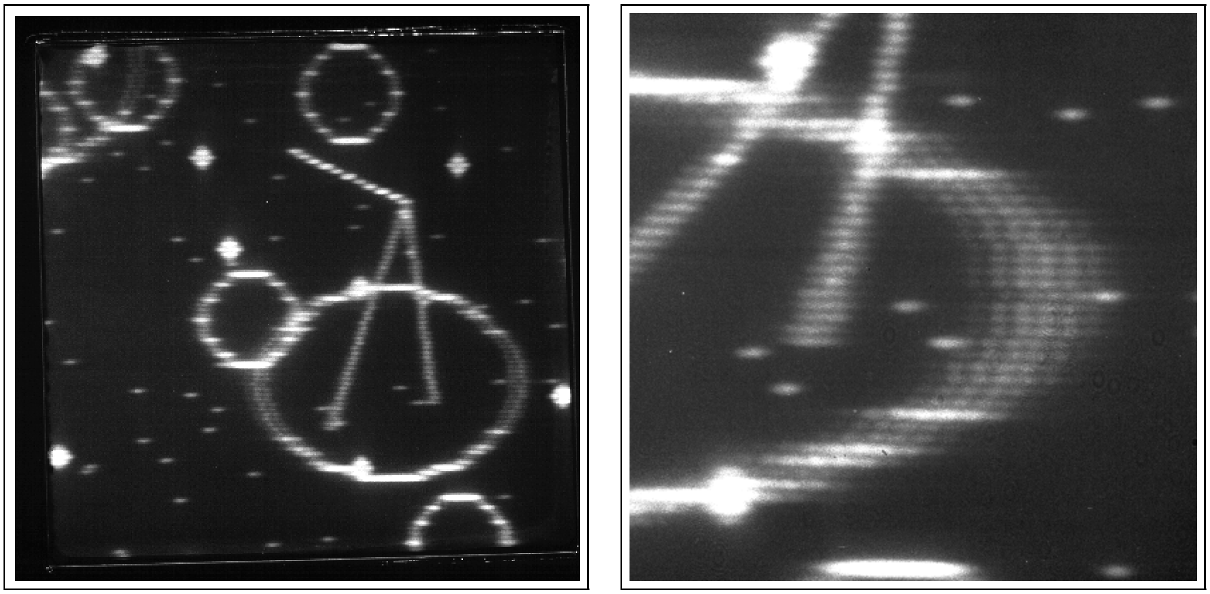

(a) Entire Image.

(b) Detail, 105mm lens, f/22.

Figure 6: *Star Wars*, William J. Molteni, Polaroid Corporation. Computer generated holographic stereogram, approximately 100 views.

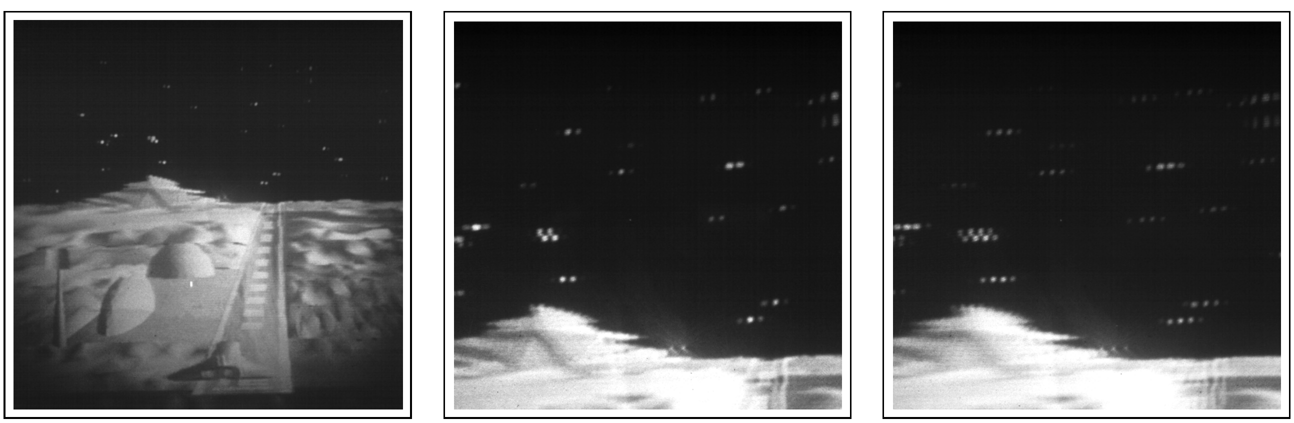

(a) Entire Image.

(b) Detail, 105mm lens, f/22.

(c) 105mm lens, f/4.

Figure 7: *Lunar Lander*, Michael Teitel, MIT Media Laboratory. Computer generated holographic stereogram, approximately 100 views.

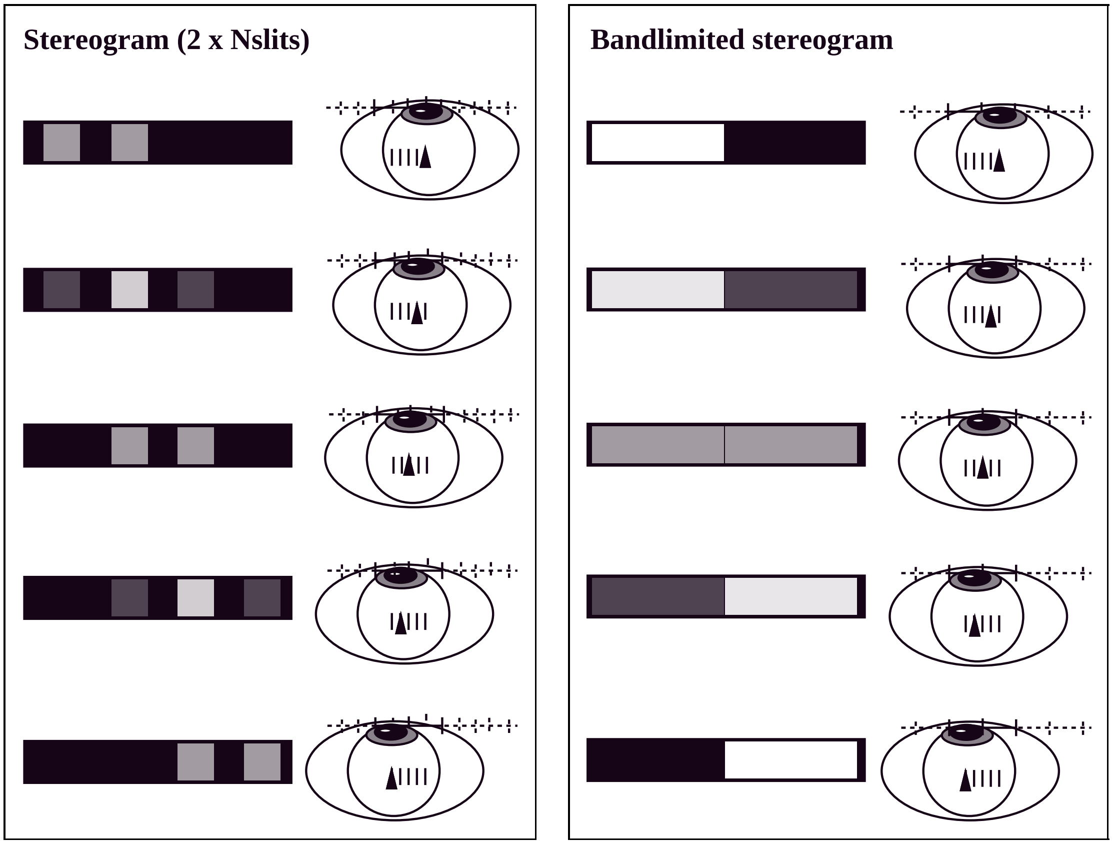

(a) (b)

Figure 8: Means of reducing or eliminating aliasing artifacts.

appears to hop across the image plane from one perspective view to the next. This horizontal size assures that no gap will exist between the two point images.

One way to perform this bandlimiting process is to filter the image plane projection of each image point based on the point's depth. This operation is called two-dimensional filtering because the filter is working on each of the two-dimensional projections of the scene. Figure 10(a) shows the width of the projection of the image point onto the image plane. Note that a point located at the slit plane of the hologram must have an infinitely wide projection to be bandlimited. This result is due to the effect of camera perspective: a point located at the aperture of the camera will cover the entire horizontal field of view.

Another approach to bandlimiting consists of selectively filtering the depth planes of the three-dimensional scene before projection, thus eliminating the effect of camera perspective on the filtering process. This technique is called three-dimensional filtering because it is performed directly on the three-dimensional scene volume instead of on the various two-dimensional projections. (In either case, however, filtering is only performed in the direction of sampled parallax. In an HPO stereogram, for instance, the filtering operation is one-dimensional in the horizontal direction.) Once again, the image plane can display arbitrarily high levels of image detail. The hologram slit plane, in contrast, is sampled only N_{slits} times, so no more than N_{slits} individual points can be be displayed horizontally at the slit plane. A three-dimensional point at the slit plane, therefore, can have a horizontal extent no smaller than W_{slit}. For all other depths, the maximum resolution can be found by a linear extrapolation of these two boundary conditions, as shown in Figure 10(b). This three-dimensional resolution limit is consistent with the two-dimensional result determined above: the projection of a bandlimited three-dimensional point onto the image plane produces a bandlimited two-dimensional point, and vice versa.

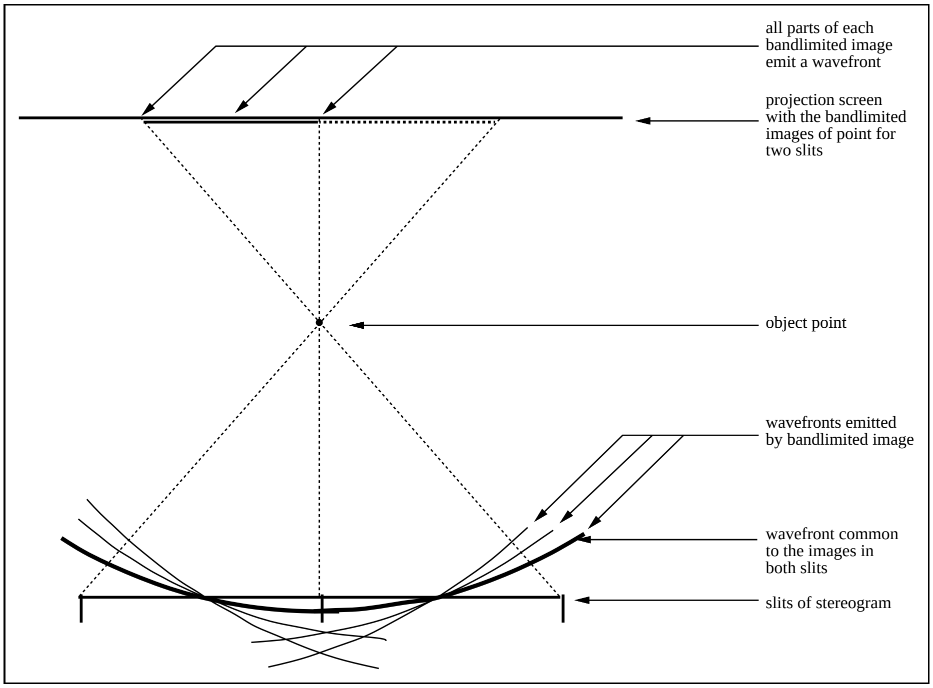

Figure 9: Bandlimiting smoothes the shape of the wavefront emitted from many slits.

8. BANDLIMITING IN PRACTICE

Several techniques exist to bandlimit holographic stereogram images. The easiest way is to increase the band limit by increasing the number of perspective views and holographic slits. The perspective information about the scene is sampled at a higher rate, increasing the likelihood that objects in adjacent point images will overlap. Increasing the number of perspective views may be a costly and difficult solution to the aliasing problem, however. Most stereogram printer setups use a slit mask of a fixed size, so changes in slit width are difficult. Also, an increase in the number of perspective images and holographic exposures may increase the expense and difficulty of producing the hologram. Last, the resolution of the finest details in the background of the scene may be both difficult to determine and non-essential to the effectiveness of the image as a whole. Increasing the resolution of the entire stereogram just to eliminate artifacts in a small, unimportant region of the image is wasteful and inconvenient.

The next bandlimiting technique, perspective averaging, increases the number of perspective views while keeping fixed the number of holographic exposures. Stereogram bandlimiting can be viewed as a way of preventing any perspective information from being missed by the taking camera. If, instead of recording one perspective from the middle of each slit, the camera recorded many perspectives from different horizontal positions within the slit, perspective information is less likely to be omitted. These images can be averaged together to find the incoherent sum (or blur) of the slit's different perspectives. The averaging should be done so that image plane detail averages with itself through all perspective views. Objects on the image plane will thus remain sharp, while objects further from the image plane will be progressively more blurred. The resulting averaged images, when exposed, produce a bandlimited stereogram display.

Recording extra perspective views and mathematically averaging them is difficult when making holograms from photographic sequences. This averaging process can, however, be approximated by the taking camera itself. First, if the camera is moving along a linear track, the camera's shutter can be left open as the camera moves through the part of the track that corresponds to a slit. The blur recorded on the film will be the sum of all the perspective data in the slit region and is bandlimited. A translating lens stage is required to keep image plane details sharp.

An even simpler approach to bandlimiting using a photographic camera is to use the lens to provide the necessary blur. If the horizontal size of the lens aperture matches the size of the slit, the lens itself captures the perspectives for the entire slit

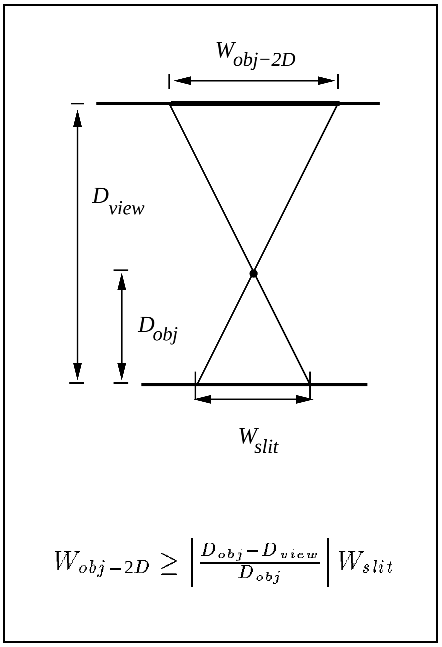

(a) Two dimensional filtering.

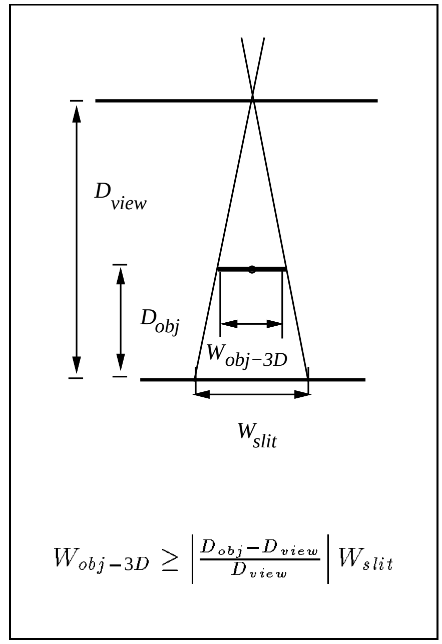

(b) Three-dimensional filtering.

Figure 10: Image band limits for points of a given depth displayed in a stereogram with slit width W_{slit}.

and combines them onto the film. The lens should be focused on the image plane to maximize resolution at that plane. In practice, most lenses have circular apertures, so vertical resolution will be sacrificed to provide horizontal bandlimiting. A circular aperture also weights the perspectives near the middle of the slit more heavily than those on the slit edges, so the chosen aperture size may need to be larger than the slit to achieve suitable results. Practical limits on lens size prevent this optical type of bandlimiting from being used in all cases. However, this technique casts the bandlimiting procedure as a way of matching the depth of field of the capture camera and the final stereogram.

Bandlimited computer generated stereograms can be made using any computer graphics techniques that produce physically accurate image blur. The technique of computing and averaging many slightly different perspective views described previously is perhaps the most general and most frequently used. Perspective views can also be filtered using depth-selective image processing techniques. Alternatively, the image volume can be pre-filtered before image generation. This pre-filtering is most applicable to volumetric images on a regular grid.

9. CONCLUSIONS

Proper sampling is an important but almost universally ignored aspect of stereographic imaging. Inherent limits exist on an image's spatial resolution, dictated by image depth and slit size. When these limits are exceeded, distracting aliasing artifacts break up the image and diminish the image's dimensionality. In many cases, simple techniques can be used to bandlimit the image, eliminating aliasing artifacts. In others, the scene itself can be manipulated so resolution constraints are not violated.

The analysis and methods described in this paper are applicable to not just two-step, HPO stereograms. The resolution limits and sampling criteria presented here can be applied as well to full parallax and one-step stereograms, ultragrams, and non-holographic multiplexed technologies such as parallax barrier and lenticular displays. Bandlimiting is especially important in displays and display types with a small number of perspective views. As interest in stereoscopic and autostereoscopic displays increases, a close analysis of resolution constraints offers a glimpse at the fundamental limits of different display technologies, and ways to live within those constraints.

10. ACKNOWLEDGEMENTS

This work was done as part of several different projects, sponsored by the Design Staff of the General Motors Corporation, Honda R&D Company, International Business Machines, and Nippon Electric Company.

11. REFERENCES

1. Stephen A. Benton, “Photographic holography”, in *Optics in Entertainment*, SPIE, 1983.
2. Stephen A. Benton, “Survey of holographic stereograms”, in *Processing and Display of Three-Dimensional Data*, SPIE, 1983.
3. D. J. DeBitetto, “Holographic panoramic stereograms synthesized from white light recordings”, *Applied Optics*, 8–8:1740–1741, August 1969.
4. I. Glaser and A. A. Friesem, “Imaging properties of holographic stereograms”, in *Three-Dimensional Imaging*, volume 120, SPIE, 1977.
5. Pierre St. Hilaire, “Optimum sampling of holographic stereograms”, *Applied Optics*, in press.
6. M. C. King, A. M. Noll, and D. H. Berry, “A new approach to computer-generated holography”, *Applied Optics*, 9–2:471–475, February 1970.
7. E. N. Leith and J. Upatnieks, “Photography by laser”, *Scientific American*, 212:24–35, June 1965.
8. J. T. McCrickerd, “Comparison of stereograms: Pinhole, fly’s eye, and holographic types”, *Journal of the Optical Society of America*, 62:64–70, January 1972.