The Generalized Holographic Stereogram

by

Michael W. Halle

S.B., Computer Science
Massachusetts Institute of Technology
Cambridge, Massachusetts
1988

Submitted to the Media Arts and Sciences Section,
School of Architecture and Planning
in partial fulfillment of the requirements
for the degree of

Master of Science in Visual Studies

at the

Massachusetts Institute of Technology

February 1991

Signature of Author ______________________________

Michael W. Halle
Media Arts and Sciences Section
January 18, 1991

Certified by ______________________________

Stephen A. Benton
Professor of Media Technology
Thesis Supervisor

Accepted by ______________________________

Stephen A. Benton
Chairman
Departmental Committee on Graduate Students

Document last updated February 24, 1993.

The Generalized Holographic Stereogram

by

Michael W. Halle

Submitted to the Media Arts and Sciences Section,
School of Architecture and Planning, on January 18, 1991
in partial fulfillment of the requirements of the degree of
Master of Science in Visual Studies at the
Massachusetts Institute of Technology

Abstract

A general model for holographic stereograms in the context of discrete and continuous optical systems is presented. Building from a simple, highly constrained model, restrictions on the viewer's horizontal position and depth, the size and location of the plane of the stereogram's slits, the spatial resolution of the points in the imaged object, and the optical properties of the projection screen are relaxed one by one. Emphasis is placed on accurate modeling of the stereogram and on the correspondence between the photographic capture, holographic recording, and final viewing geometries. Bandlimiting and anamorphic distortion techniques useful for producing artifact-free images are presented. Discussion centers on horizontal parallax only, flat format, computer generated stereograms, but general conclusions applicable to other stereogram types are drawn. Specific examples of stereograms created using these techniques are shown and discussed.

Thesis Supervisor: Stephen A. Benton
Title: Professor of Media Technology

This research was supported in part by the Defense Advanced Research Projects Agency under Rome Air Development Center (RADC) Contract F30602-89-C-0022, and by the Design Staff of the General Motors Corporation.

Dedication

This thesis is dedicated to my parents. To my mother, who tirelessly strived to challenge her children, and to instill in them the desire to challenge themselves. And to my father, who never stopped trying to give us the opportunities that he never had. May this work and all that I do stand as a proud but humble tribute to their ideals and selflessness.

Acknowledgements

And now, a word of thanks.

Thanks to my cohorts in the Spatial Imaging Group, who were always around to let me banter about some sort of malformed idea, and who looked over drafts of this work. Especially noteworthy: John Underkoffler implemented much of the “weird rendering things” necessary to apply this paper’s concepts to the images presented. Michael Klug was the “other half”, doing the lab work without which all this stuff would be academic. He also provided several of the photographs used in this text. Both put up with a lot of false starts and provided extended comic relief through late nights and Thanksgivings.

Thanks to Stephen Benton, my advisor, for making available the fine facilities of the Spatial Imaging Group, and for his review of this lengthy document.

And thanks to my housemate Steve Pieper, whose timely acquisition of a Personal Iris created an unparalleled thesis production environment without which this document would not have been possible.

Contents

Chapter 1

Introduction

Ever since the first stereo photographs were produced in the 1800's, humanity has been fascinated with the idea of a medium capable of accurately recording and displaying a scene in three dimensions. Many schemes for three-dimensional display have been created over the years. Mechanical schemes, including parallax barrier displays and lenticulars to name just two, are inherently limited in the quality of their display, but have brought simple auto-stereoscopic "three-d" to the public. Their intricate construction played to the imagination of "garage inventors", researchers who through their ingenuity made small, empirically determined optimizations to their displays without truly understanding the consequences.

Laser-illuminated display holography, developed in 1964 by Leith and Upatnieks[10], was the first truly high quality three-dimensional display medium. The independent work of Denisyuk and Benton[9] that led to white light viewable holograms brought the new display medium out of the laser lab and into the practical, useful world. The optimistic popular press of the 1960's proclaimed the advent of holographic street signs and snapshots.

But the hologram is burdened by the fact that it is not only a display but a recording medium. Holographic recording must be done in monochromatic, coherent light, and requires that the objects being imaged remain stable to within a fraction of a wavelength of light. These requirements have hindered holography from gaining widespread use. In addition, the amount of optical information stored in a hologram makes the computation of holographic patterns very difficult. Until recently, the creation of synthetic display holograms by computer has not been practical.

Holographic stereography weds the structure and basic optical properties of mechanical systems with the huge information storage potential and image fidelity of holography. A holographic stereogram records a relatively large number of viewpoints of an object and uses a hologram to record those viewpoints and present them a viewer. The information content of the stereogram is greatly reduced from that of a true hologram because only a finite number of different views of the scene are stored. The number of views captured can be chosen based on human perception rather than on the storage capacity of the medium. The capturing of the viewpoints for the stereogram is detached from the recording process; image capture is photographic and optically incoherent, so that images of natural scenes with natural lighting can be displayed in a stereogram. The input views for traditional stereograms are taken with ordinary photographic cameras and can be synthesized using commonplace computer graphic techniques. Using recently developed true color holographic techniques, extremely high quality, accurate, and natural-looking display holograms can be produced.

The history of holographic stereograms shares something of the "garage inventor" level of understanding with its mechanical-3D brothers. The spatially multiplexed stereogram of today closely resembles the single-step stereograms of DeBitetto[3] in 1968 and the two-step transfers of stereograms made by King, Noll and Berry[8] in 1970. Although some modifications have been made to the stereogram's holographic exposure geometry, and some attempts have been made to understand and compensate for the distortions that occur in stereograms [4][13][14][6], the behavior of the stereogram's optical system as a whole has seldom been analyzed.

The purpose of this work is to begin that analysis, in order to see if the quality of stereograms can be improved by better understanding how they work. More specifically, a primary goal of this thesis is to explore the constraints that stereography places on the necessary photographic capture, holographic exposure, and final viewing geometries. For example, conventional display stereograms use an optical transfer step to separate the viewer from the plane of the hologram. But one-step stereograms are appealing for holographic printer applications. If a viewer steps away from a conventional one-step stereogram, the resulting image changes in aspect ratio and otherwise distorts. If these distortions can be understood, perhaps they can be eliminated, allowing practical one-step stereograms to be produced. In two-step stereograms, the size of the zone within which a viewer can see an image is limited to the size of the "master" holographic plate. The viewer must also be positioned at a distance corresponding to the separation between the master and the transfer plates during the holographic transfer step. For stereograms with long view distances, this constraint requires the holographic recording apparatus to be large, bulky, and expensive. A relaxing of the connection between holographic exposure and viewing would allow the simplified production of stereograms with large view zones at arbitrary view distances.

Wide applicability and practicality of method are foremost in importance to the approach of this research. The topics discussed here are of fundamental importance to all types of holographic stereograms, and the methods presented are designed to efficiently produce actual images. The amount of extra work required to compensate for image aberrations is about the same as needed for the other steps in the stereographic process. Several stereograms developed during the course of this work are presented as evidence of the correctness and practicality of these new techniques, as well as to provide specific examples within a historical context.

Again for practicality's sake, this text will concentrate on stereograms made with slit apertures and presenting horizontal parallax only. While full-parallax stereograms most closely mimic the natural world, HPO stereograms provide the viewer with most of the three-dimensional information about the scene with a greatly reduced number of camera viewpoints and holographic exposures. The principles presented, however, apply equally to both HPO and full-parallax stereograms. For similar reasons, flat format stereograms are discussed to the exclusion of cylindrical formats, and diffusing projection screens are used over cylindrical lenses and other large optical elements. The general ideas transcend the precise choices of format or design.

In part for practical reasons, this text places its emphasis on issues related to computer-generated stereogram images. The concepts of bandlimiting and distortion compensation implemented here could be implemented either with physical optics or the computer equivalents of those optics. Unlike physical optical elements, computer models are configurable, aberration-free, and relatively inexpensive. Computer image processing permits large amounts of image information to be manipulated in a precise, predictable and adaptable way. Changes in holographic geometry, for example, can be accomodated by altering software parameters instead of fabricating new lenses. Indead, the wide range of holographic formats and sizes presented here almost necessitates the use of computer image processing.

A prevailing theme throughout this thesis is the importance of correspondence between the three stages of stereogram creation: photographic capture, holographic recording, and final viewing. Attention to this correspondence is of paramount importance in order to minimize image distortions. Similarly, to correct for distortions, a stereogram that violates this correspondence in some way can be altered so that correspondence is achieved or approximated. Each step must be accurately modelled to understand how the effect of each stage matches that of the other two. To simplify this process, a minimal, highly constrained stereogram model will first be presented. Once that the behavior of that model is understood, constraints of viewer position, slit size, object resolution, slit position, and various optical restrictions will be removed one by one. Finally, a general model for holographic stereograms will be presented.

Chapter 2

Stereogram Basics

The holographic stereogram is a means of approximating a continuous optical phenomenon in a discrete form. In display holo-stereography, the continuous three-dimensional information of an object's appearance can be approximated by a relatively small number of two-dimensional images of that object. While these images can be taken with a photographic camera or synthesized using a computer, both capture processes can be modeled as if a physical camera was used to acquire them. The photographic capture, the holographic recording, and the final viewing geometries all determine how accurately a particular holographic stereogram approximates a continuous scene. This chapter presents the basic principles of holo-stereography for visual display and lays the groundwork for later chapters on the analysis of sampling and distortions effects.

The simple stereogram model

The type of stereogram to be used as a first example is similar to the one described by DeBitetto, with modifications to the holographic exposure setup made by Benton. It consists of a single holographic plate comprised of a series of thin vertical slit holograms exposed one next to the other across the plate's horizontal extent. Each slit is individually exposed to an image projected onto a rear-projection screen some distance away from the plate. Once the hologram is developed, each slit forms an aperture through which the image of the projection screen at the time of that slit's exposure can be seen. The images projected onto the screen are usually views of an object captured from many different viewpoints. A viewer looking at the stereogram will see two different projection views through two slit apertures, one through each eye. The brain interprets the differences between the two views as three-dimensional information. If the viewer moves side to side, different pairs of images are presented, and so the scene appears to gradually and accurately change from one viewpoint to the next to faithfully mimic the appearance of an actual three-dimensional scene.

In the beginning of this chapter, the assumption will be made that the slit apertures are about as wide as the pupil of the eye, and that when the viewer looks at the hologram each eye sees through one and only one aperture at a time. During viewing, then, the viewer's face must be right up against the surface of the plate, an awkward and inconvenient location, as shown in Figure 2.1. Later in this chapter, more convenient ways to view a holographic stereogram will be discussed.

The holographic exposure setups used to expose all of the stereograms discussed in this chapter follow the same basic layout, similar to the common two beam off-axis holographic setup used to make holograms of real objects. The layout is shown in Figure 2.2. Two beams of mutually coherent light are used to expose the holographic recording material. The reference beam diverges from a point source at the same location as the one that will eventually be used to illuminate the hologram. The final hologram is to be illuminated from above and behind; for practical reasons, the entire optical setup must be flipped "on its side" so that the illumination beam, and by direct consequence the reference beam, can travel parallel to the table surface and yet strike the plate at the correct angle. The reference

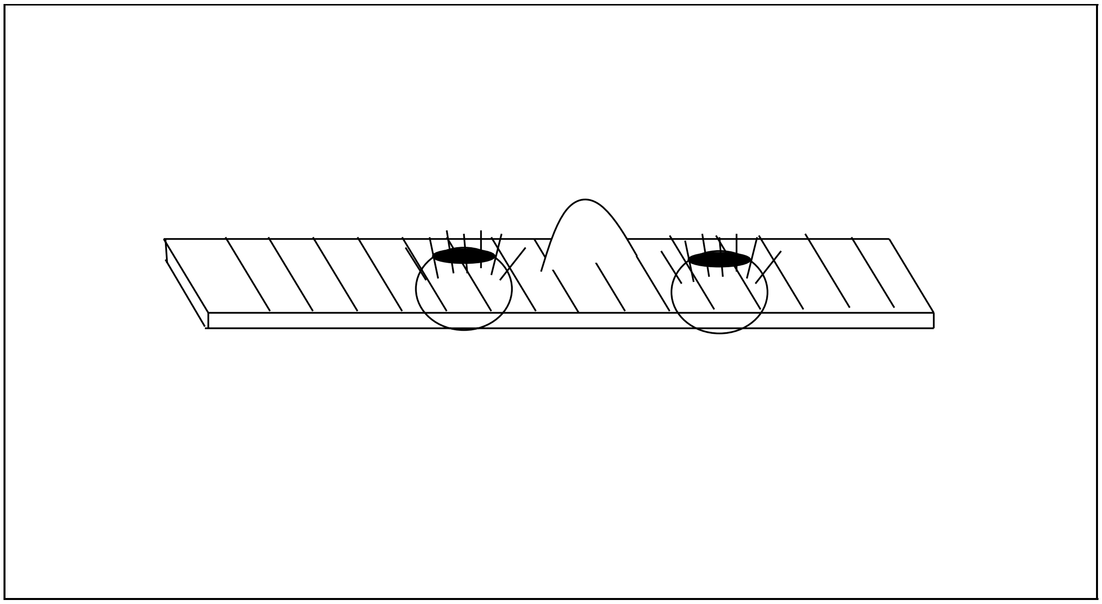

Figure 2.1: Viewing a simple stereogram requires that the viewer's eyes be positioned at the plane of the holographic plate.

beam is collimated so that each slit of the holographic plate will be referenced with a beam of the same direction, independent of the slit's precise lateral position on the plate.

The object beam is diverged and used to project an image of the camera's recording medium, typically photographic frames on cinéfilm, onto the projection screen. The *projection screen* serves as a two dimensional object that can be changed from view to view. The actual projected image, whose extent is defined by the *projection frame*, is the visible subregion of the projection screen in any particular view. The projection screen itself is a subregion of a plane of infinite extent called the *projection plane*.

The projection screen directly faces the holographic plate and slit mechanism. The details of this mechanism vary for different types of stereograms, but in all cases, the holographic plate is covered by a piece of optically opaque material with a slit-shaped hole that masks off all but a stripe for exposure. During exposure, this stripe is exposed both to the image on the projection screen and to the reference source. Either the plate or the slit

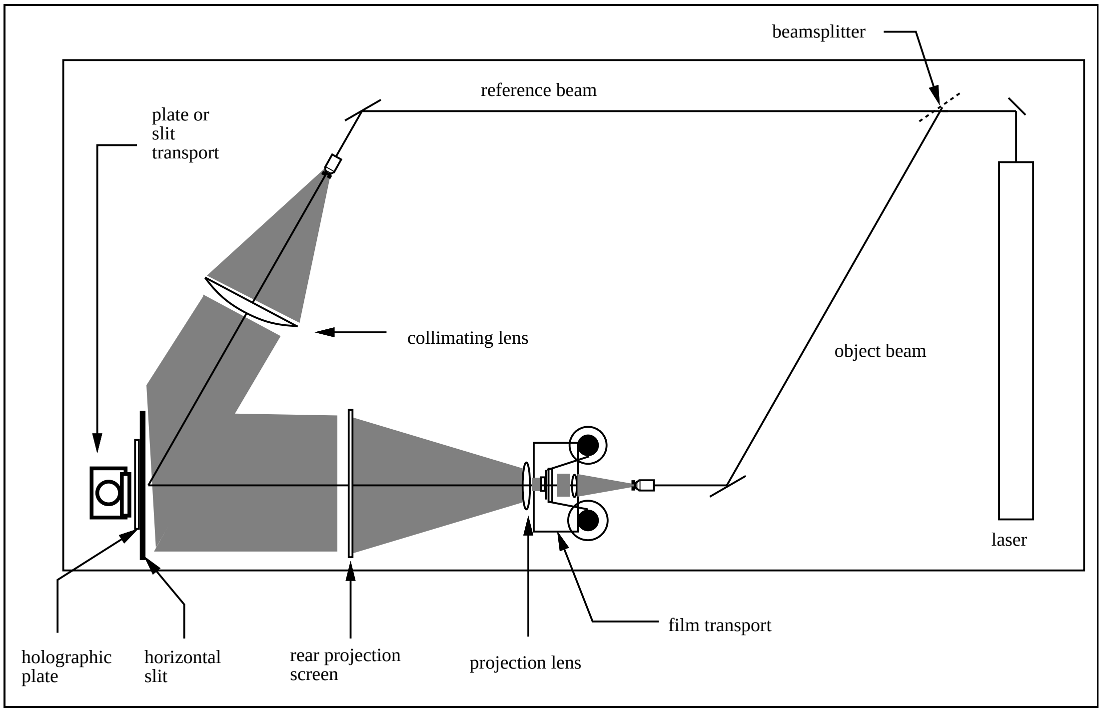

Figure 2.2: An above view of the general holographic table layout used for making the stereograms of the types described in this chapter.

is moved between one exposure and the next so that the next adjacent stripe of holographic material can be exposed.

In display holo-stereography, the pictures imaged onto the projection screen are projectional views of an object recorded with some sort of camera. For a point on an object to be visible to the stereogram's viewer from a particular viewpoint, the image of that point must fall in the projection frame for the slit that corresponds to that view location. Depending on the exact projection and exposure setup, the projection frame (or the projection screen itself) may appear stationary as the viewer's eye moves from view to view across the surface of the hologram; alternatively, the image of the projection frame may move across the viewer's field of view as if it were at a finite distance from the plate. Because the setup is flipped sideways, images projected onto the screen must similarly flipped. The top of the projected images must point towards the reference beam, which defines from where the eventual overhead illumination will come. "Top" on the projection screen, then, is toward the interior of the table, while "left" is into the table and "right" is toward the sky.

Simple camera stereogram

The first example of a stereogram geometry centers the projection frame in front of each slit being exposed during every exposure. When a stereogram made with this geometry is illuminated with the reference source, every slit forms an image of the projection screen that is centered in front of it. Figure 2.3 shows several slits and the location of the projection screens that each slit forms. Such a stereogram can be produced with a holographic recording apparatus that fixes the location of the slit aperture with respect to the projection frame. The following exposing apparatus satisfies this geometrical constraint. The plate is fixed to a movable plateholder, and positioned behind a fixed horizontal slit aperture. The projection screen is centered in front of each slit. Figure 2.4 shows this setup.

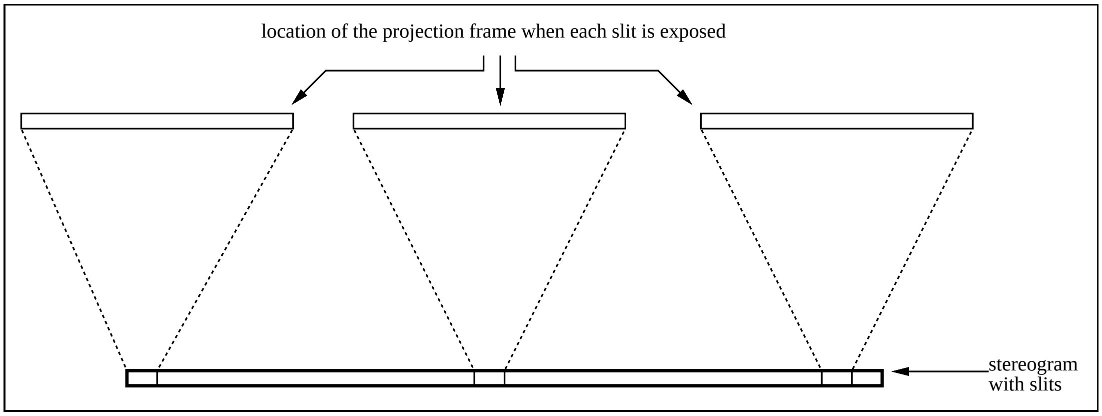

Figure 2.3: Three slits of a holographic stereogram using footage taken with a simple camera. The projection screen is always centered in front of the slit being exposed during exposure and viewing.

The stereogram is exposed in the following way. The plateholder is positioned at the upper end of its travel so that the eventual leftmost slit is behind the slit aperture. The image intended to be seen through the leftmost slit is projected onto the projection frame and a reference beam simultaneously exposes the slit of hologram. The intensity and direction of light from every part of the projection frame is recorded as a latent image in the slit-shaped area of holographic material. The movable plateholder is then repositioned so that the next unexposed segment of plate is behind the slit aperture, a new image is projected onto the projection screen, and another exposure is made. The process continues until the entire width of the hologram has been exposed. After the plate is developed, the resulting diffraction pattern on the emulsion (either a phase or an amplitude grating depending on the type of processing used) will, when viewed with a collimated monochromatic source at the reference beam angle, appear as a sequence of slit-like windows, each presenting a different image that appeared on the projection screen.

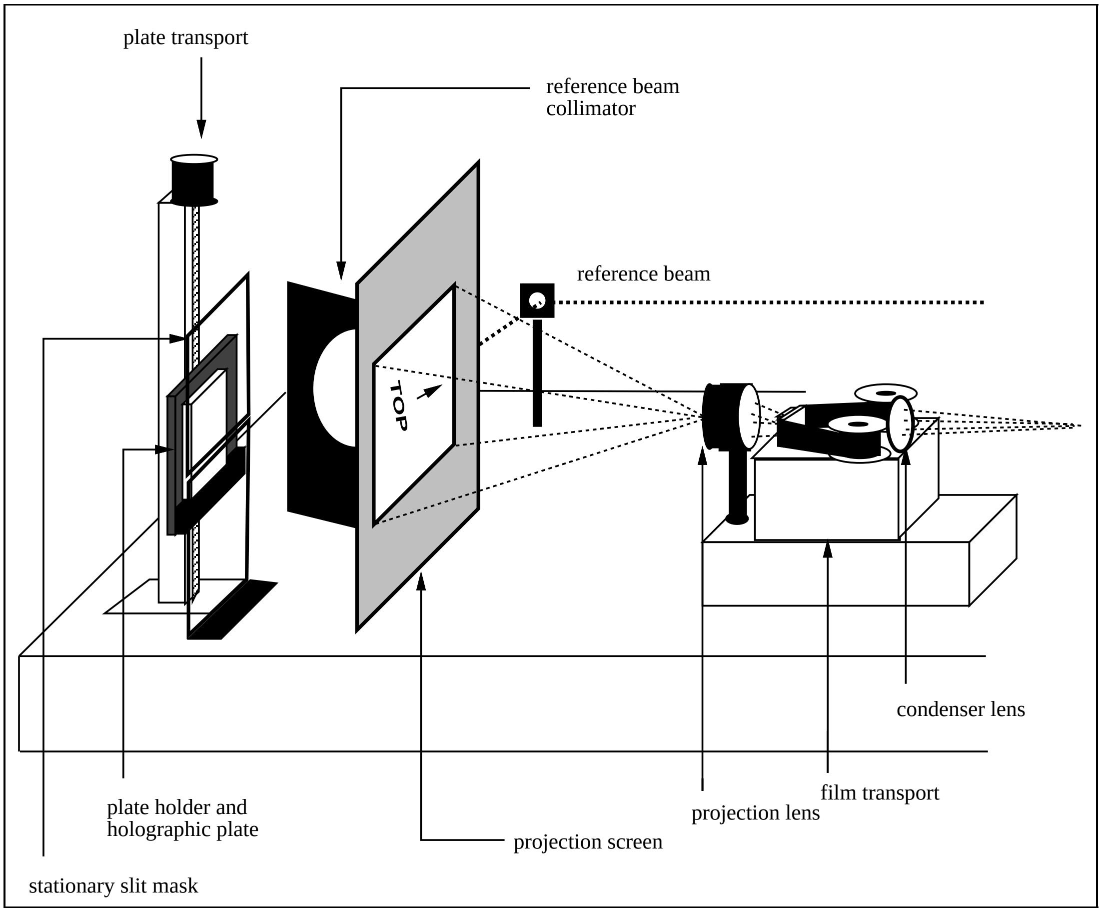

Figure 2.4: Holographic table layout for simple camera geometry.

How a stereogram displays depth

A closer look at the details of this stereographic exposure provides some insight into how the projectional views of an object should be captured so that the final stereogram produces an accurate three-dimensional image of that object. Imagine that the image projected on the screen consisted of a single bright point centered in a field of black. Each slit of the stereogram is exposed to this test pattern with the bright point centered in front of it. The image neither changes nor moves with respect to the slit from one exposure to the next. When the final hologram is viewed, the viewer's two eyes fall behind two different slits of the hologram. The point appears to be directly in front of both of the viewer's eyes. The viewer interprets the two images stereoscopically as if a single point were located at infinity. This binocular depth cue is very strong; horizontal image parallax provides most of the viewer's depth sense. However, two other inaccurate depth cues provide conflicting information to the stereogram viewer. First, the viewer's eyes must still focus on the projection plane to focus on the point, so focus cues indicate that the point lies on the projection plane. Second, because the slits are recorded horizontally, the plate records only a single vertical perspective. The same perspective is presented independent of the viewer's vertical position. If the viewer moves vertically, the point will move as if it were located at the vertical projection plane. The precise significance of focus and vertical parallax cues in stereograms has not been fully studied, but it appears to be minor compared to binocular cues.

Stereogram camera geometry

Using the observation that a stationary point appears to be at infinity as a landmark, the correct camera geometry needed to accurately capture a three-dimensional scene can be inferred. To appear at infinity, then, an object point must remain at the same position in every camera view. This constraint implies that the camera should face the same direction, straight ahead, as each frame is captured. The camera moves along a track whose position and length corresponds to the final stereogram plate. The camera takes pictures of a scene from viewpoints that corresponds to the the locations of the stereogram's slits. The plate is planar, so the camera track must be straight, not curved. The camera must be able to image the area corresponding to the projection frame onto its film; thus, the frame defines the cross section of the viewing pyramid with its apex is located at the camera position, as shown in Figure 2.5. Because the projection frame bounds the camera's image, the size of the projection frame and its distance from the slit determine the angle of view of the image and thus the maximum (and optimal) focal length of the camera's lens.

The film plane of the stereogram capture camera is always parallel to the plane of the scene that corresponds to the projection plane (the *capture projection plane*) in order to image it without geometric distortions onto the focal plane of the lens. If the film plane were tipped with respect to the projection plane when the image was captured, that image that appears on the film would be scaled vertically by an amount that varied from the left side of the image to the right, turning the bounds of the projection frame into a shape that resembles a sideways keystone. Keystone distortion is shown in Figure 2.6.

The correspondence between the photographic capture and holographic exposure geometries for the simple camera stereogram is shown in Figure 2.7. If this correspondence is maintained, the images of all object points, not just points located far from the camera, will appear to be at the same depth as in the original scene. So a complicated three-dimensional scene composed of many points will appear undistorted if correspondence between the two geometries if maintained. To uniformly scale an object in all dimensions, the holographic recording geometry can be a scale model of the photographic capture geometry. For proper scaling between the two geometries, the angles A and B in Figure 2.7 should match. If this restriction is violated, one or more dimensions of the object will appear too great or too small in extent (for instance, the object's depth may be exaggerated or reduced).

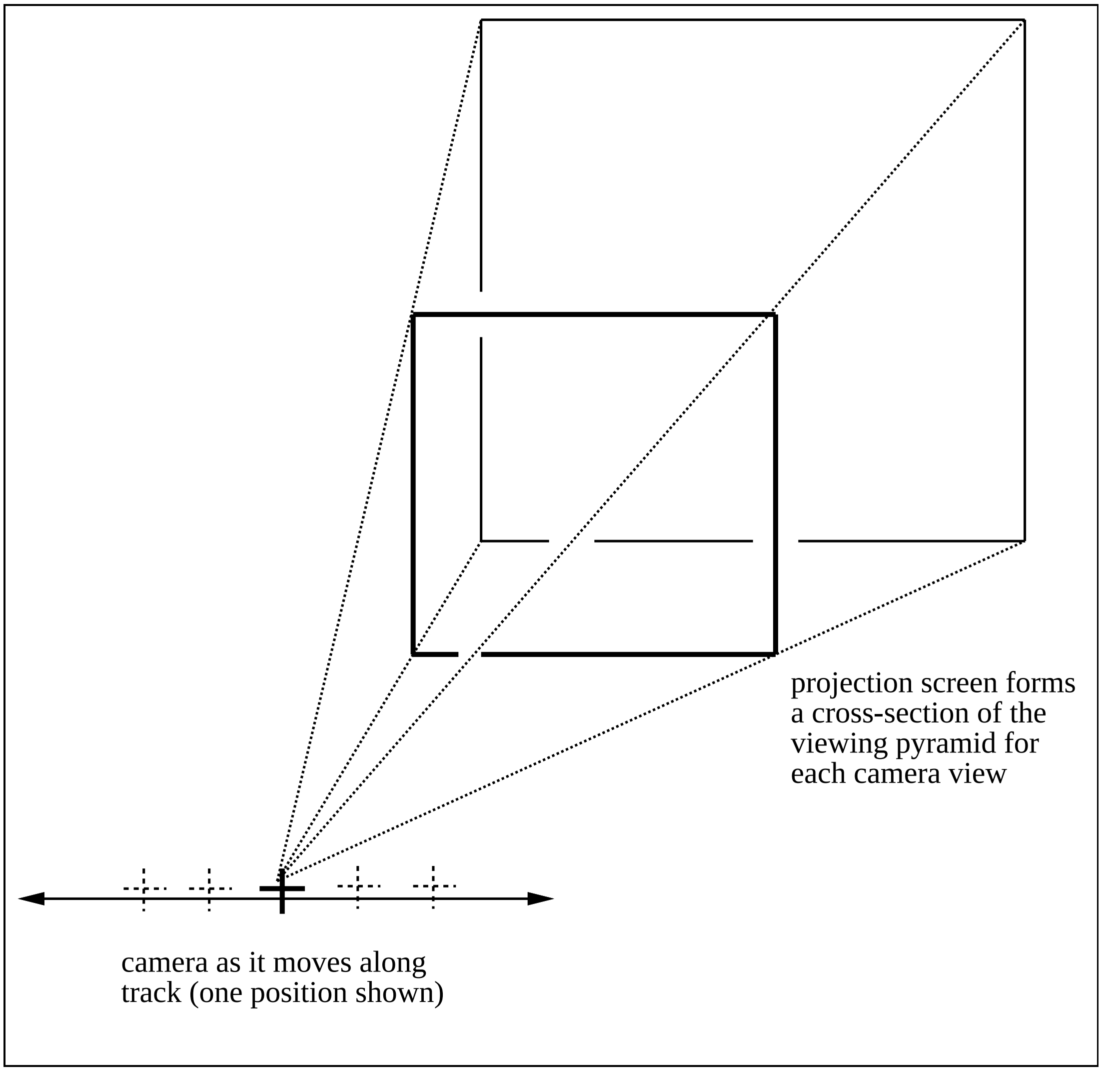

Figure 2.5: As the camera moves along the stereogram track, it images the space defined as a pyramid with apex at the current camera position and a cross-section defined by the projection plane. The camera is always pointing directly ahead.

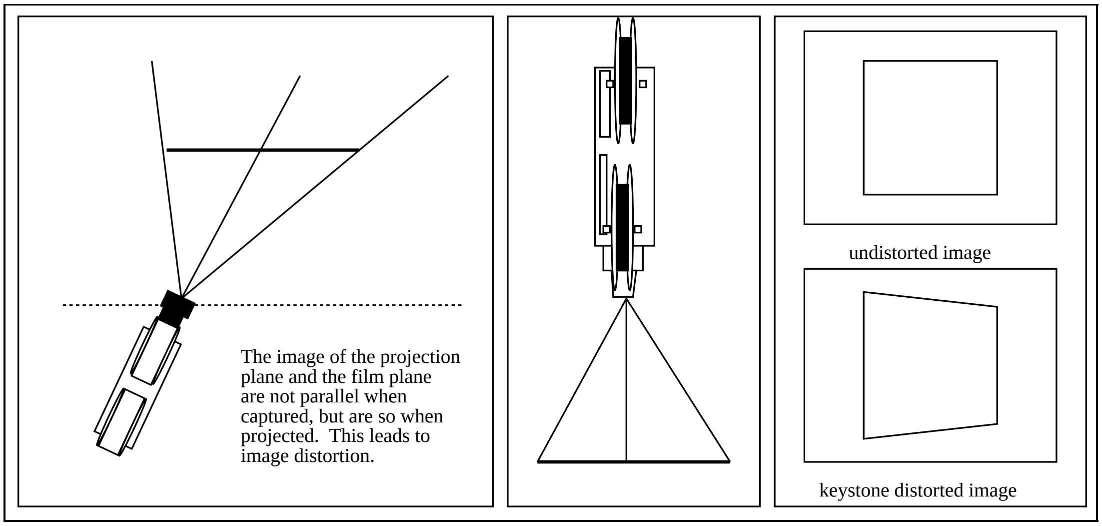

Figure 2.6: Keystone distortion occurs when the relationship between the film plane and the projection plane is not maintained from photographic capture to holographic recording.

Recentering camera stereograms

The above stereogram exposure geometry is well suited for objects far from the camera because the image of the object wanders little from frame to frame, always remaining in the camera's field of view and thus always visible to the stereogram viewer. However, distant objects are seldom the center of interest in three-dimensional images because the different perspectives captured over the view zone have little disparity and, as a result, convey little sense of depth. Objects at more interesting locations, closer to the camera, wander across the frame from one camera view to the next and tend to be vignetted in the camera's image at either or both extremes of the camera's travel. The solution to the problem is to alter the capture camera to always frame the object of interest as it records the photographic sequence. Effectively, this change centers the object plane in every camera frame so that it remains stationary on the film from view to view. Object points in front of or behind the stationary plane will translate horizontally from view to view, but at a slower rate than they would in a simple camera stereogram.

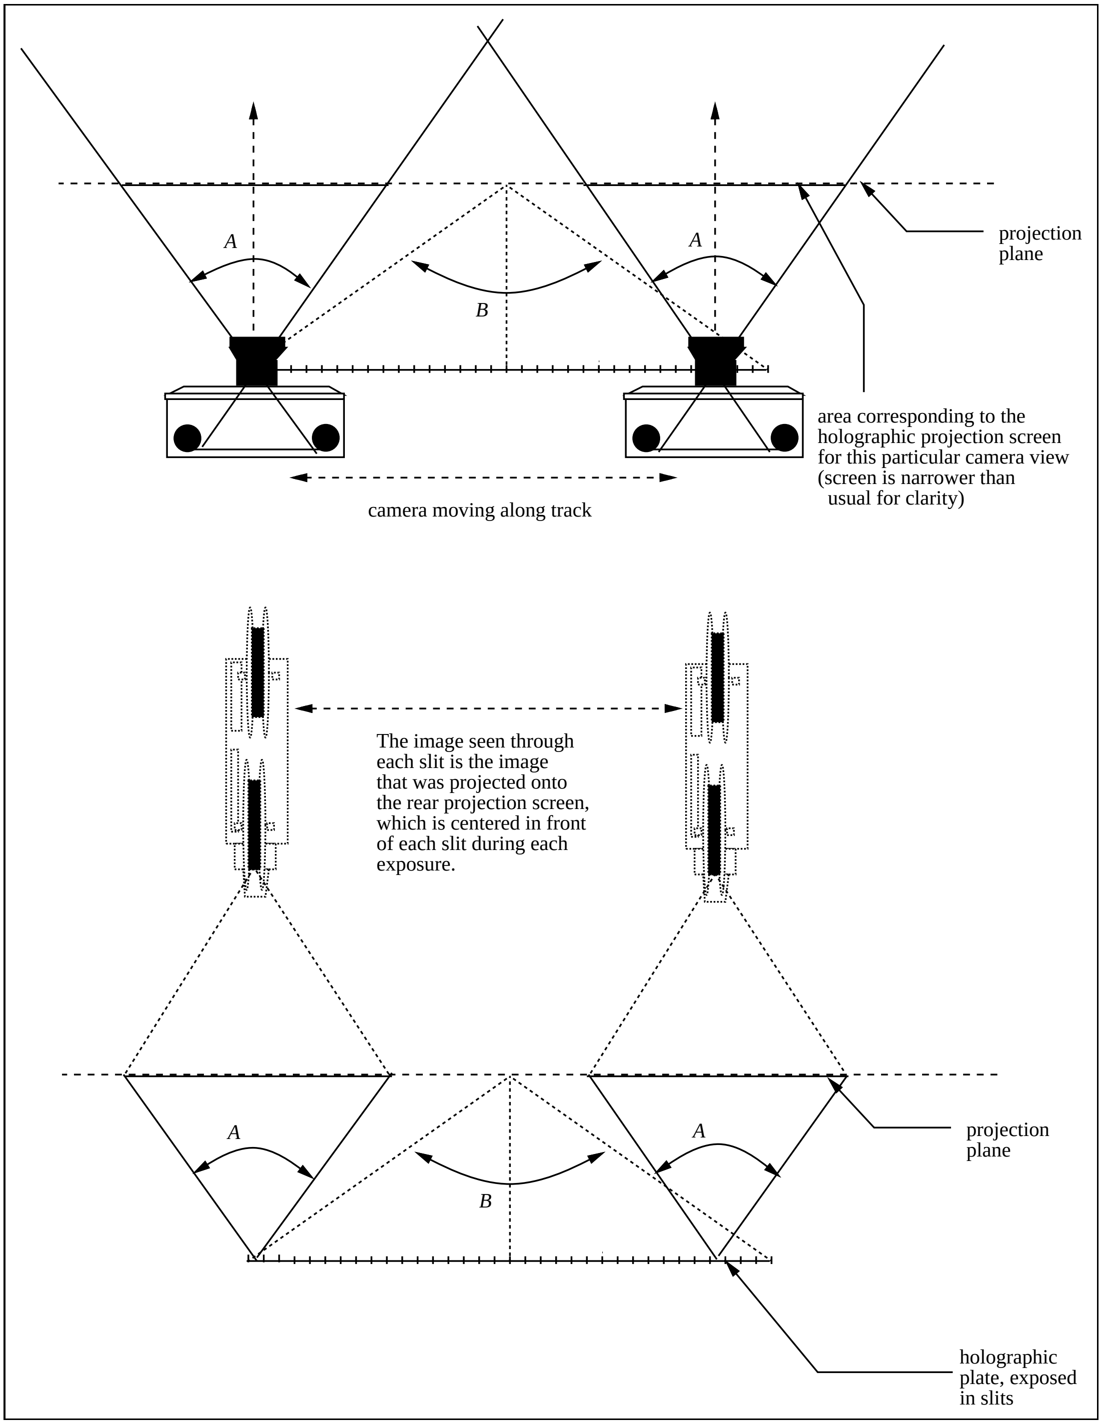

Figure 2.7: The relationship between holo-stereographic camera and recording geometries.

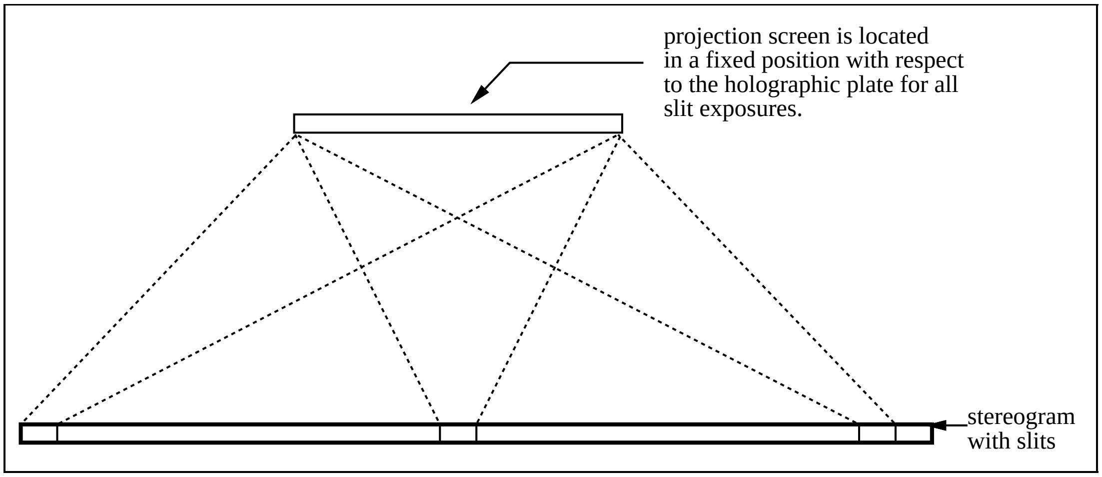

Figure 2.8: Three slits of a holographic stereogram exposed using a recentered camera geometry.

Altering the camera geometry requires changes in the holographic exposure geometry needed to produce undistorted images. The projection screen is no longer centered in front of the slit aperture during all exposures. Instead, the holographic plate holder is stationary and the slit in front of it moves from exposure to exposure. Thus, the projection frame is fixed in space relative to the plate for all exposures, rather than being centered in front of each slit during each exposure. In this geometry, called the “recentered camera” geometry, only one projection frame position exists for all slits, as shown in Figure 2.8. Holographically, such a stereogram can be realized using a table layout like that shown in Figure 2.9.

Because the projection frame and the plate are fixed with respect to each other for every slit exposure, the projection frame appears to be at its true, physical location in space. If the bright point on a black field is projected onto the screen and all the slits are exposed to that image, the point will appear not at infinity as before, but at the projection plane distance instead just as if it were a hole cut into a black card. In effect, as the viewer looks at the final stereogram, the projection frame no longer seems to follow the viewer but instead appears stationary in space. If an image of the object plane of the original scene

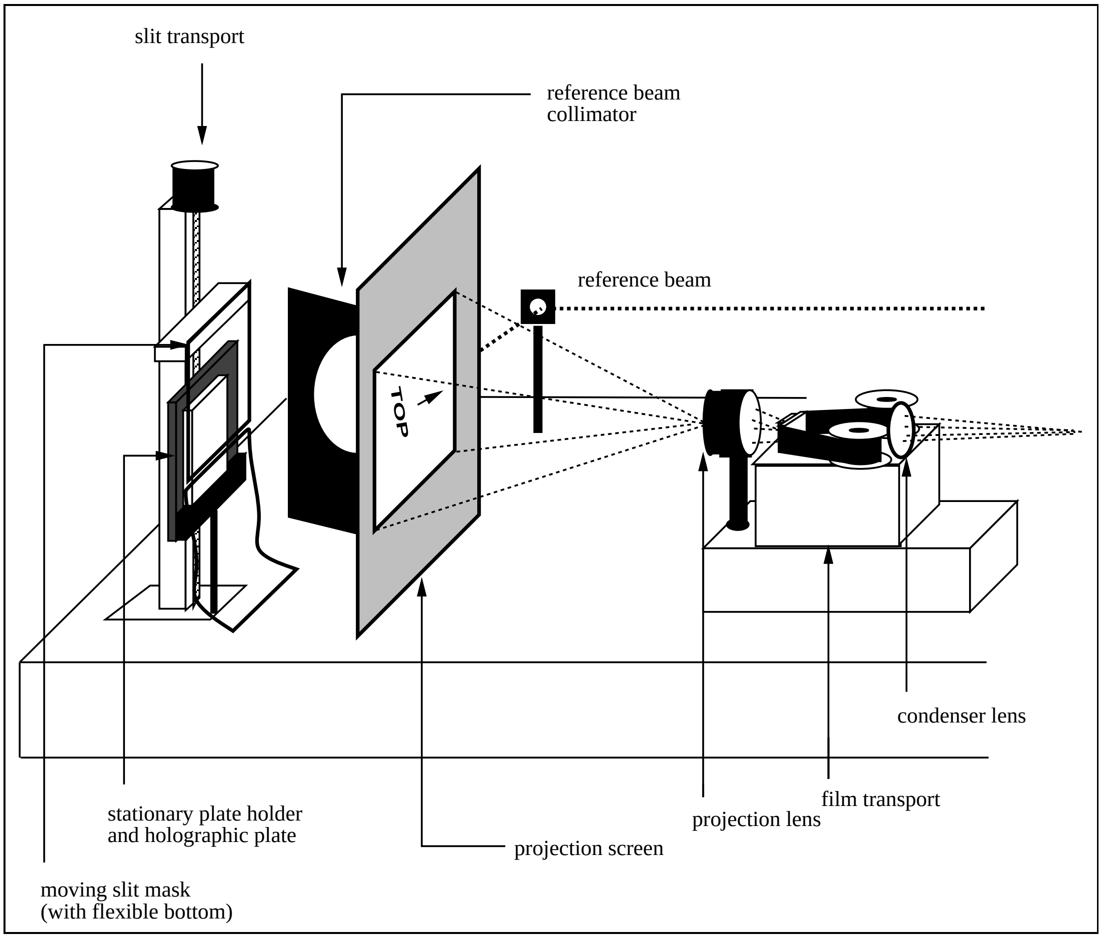

Figure 2.9: Holographic table layout for a recentered stereogram.

remains stationary on the projection screen, then, the object plane of the original scene and the projection plane of the final hologram will lie at the same depth.

Recentering cameras

The type of camera necessary to take pictures for this type of stereogram is based on, but more complicated than, the simple camera used to capture the infinitely distant object. This new type of camera is called a *recentering camera*. Recall that in the simple camera image capture, the image of a nearby object point translated across the camera's film plane as the camera moved down its track taking pictures. In a recentering camera, the lens and the film back of the camera can move independently from each other, so the film plane can be translated at the same rate as the image of the object of interest. The film and image move together through all frames, so just as desired the image appears stationary in all the resulting images. A view camera with a "shifting" or "shearing" lens provides this type of recentering. A picture of a recentering camera is shown in Figure 2.10. The lens of the camera must be wide enough to always capture the full horizontal extent of the object plane without vignetting the image at extreme camera positions.

Once again, a correspondence must exist between the camera capture and the holographic exposure geometries. In the recentering camera system, the necessary translation of the camera's lens adds another constraint that must be maintained. A point in the middle of the object plane must always be imaged into the middle of the film plane, and must always be projected onto the middle of the projection frame. The angle subtended by the object frame as seen from the camera must equal to the angle subtended by the projection frame as seen from the slit. If, for example, the focal length of the lens of the taking camera is changed, the amount of lens translation required and the size of the holographic projection frame would also have to be adjusted. The relationship between the camera and holographic geometries is illustrated in Figure 2.11. As in the simple camera case, angles

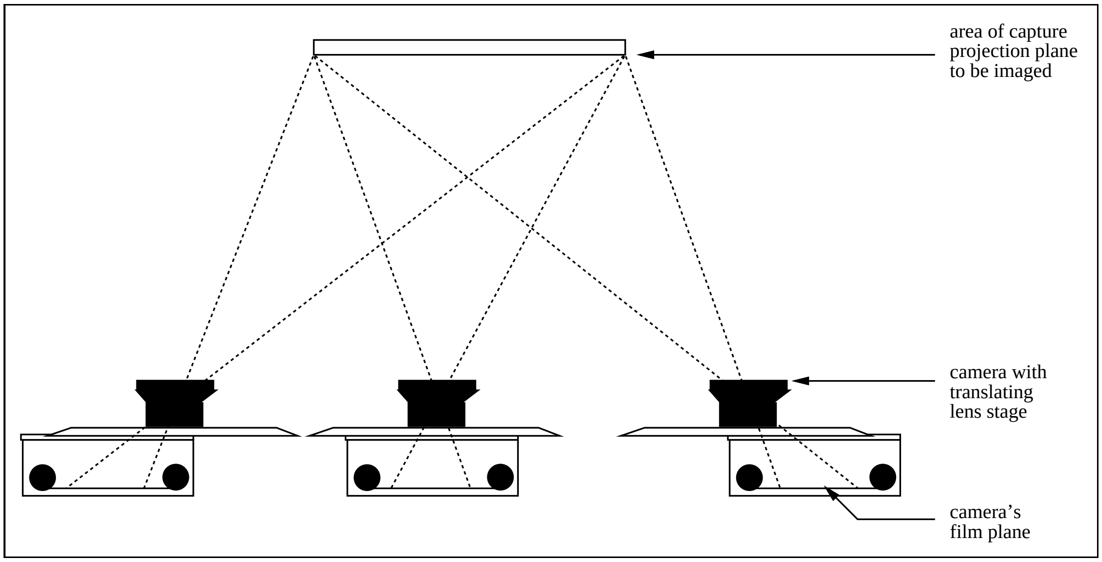

Figure 2.10: A recentering camera. As the camera moves down the track, the film back translates with respect to the lens to keep the image of an object at a finite distance stationary on the camera's film plane.

A and B must be equal in the capture and holographic recording geometries.

Wrong ways to recenter

The advantage of a camera that recenters the object of interest from one frame to the next has been known to stereographers for some time. However, many different ways of moving the camera can produce recentered views; most of them, however, produce significant distortions in the final stereoscopic image. Two common distortion-inducing methods are shown in Figure 2.12. For instance, a rig with a camera moving on a circular track, always facing the centrally located object of interest, is especially simple to construct. This setup is geometrically equivalent to one in which a stationary camera takes pictures of a subject rotating on a turntable. However, the necessary correspondence between the position of the camera from which a particular image is taken and the position in space from which the

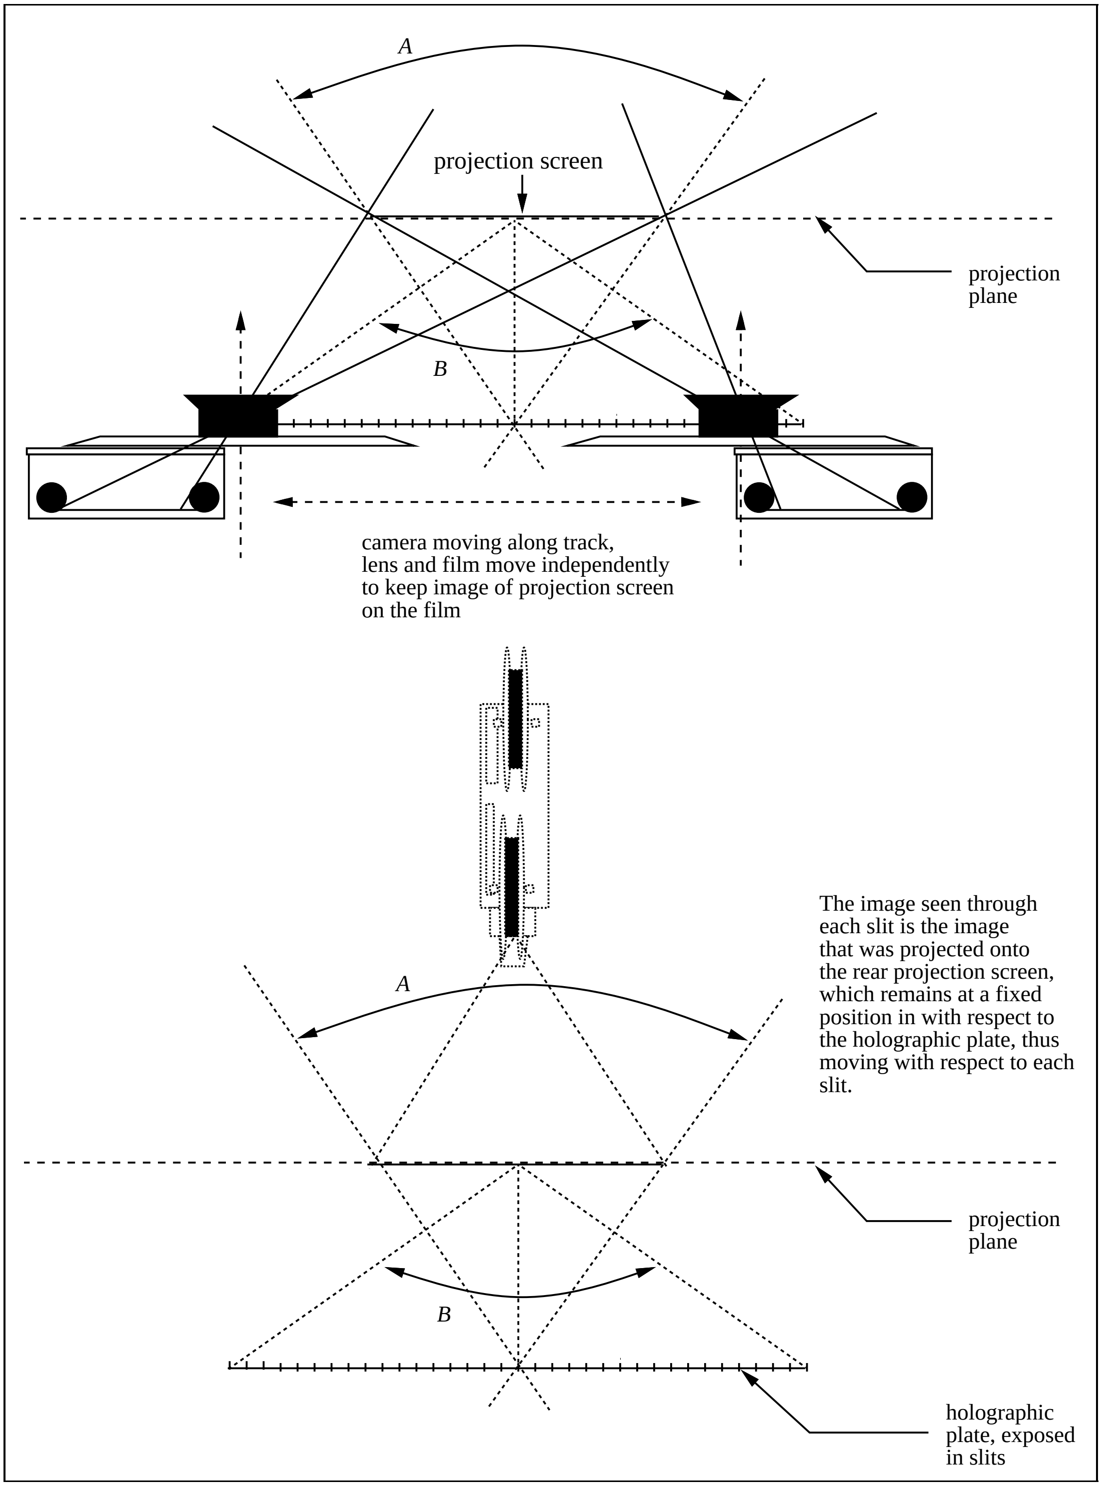

camera moving along track, lens and film move independently to keep image of projection screen on the film

The image seen through each slit is the image that was projected onto the rear projection screen, which remains at a fixed position in with respect to the holographic plate, thus moving with respect to each slit.

Figure 2.11: The relationship between capture and recording geometries for a recentering camera stereogram.

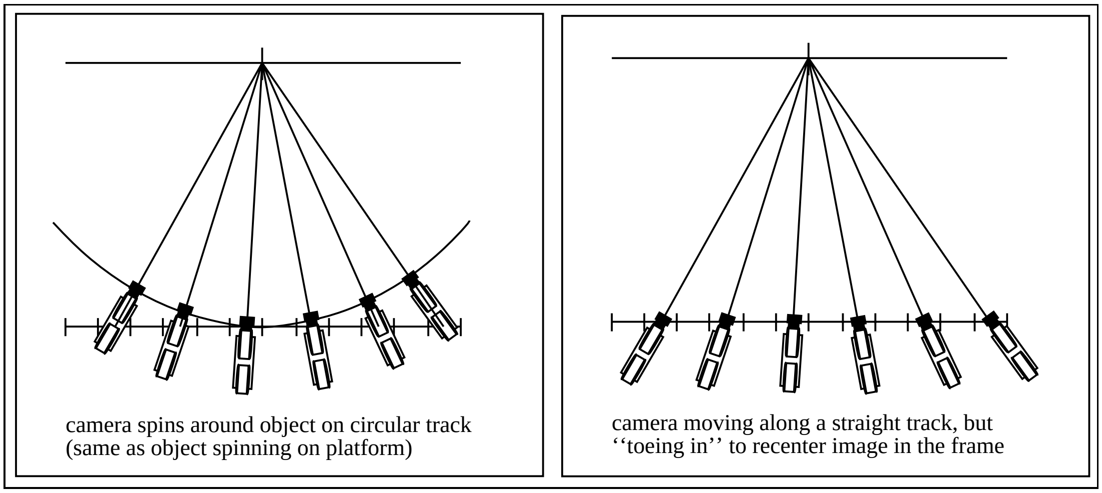

Figure 2.12: Two common ways of capturing images for a stereogram that introduce distortions in the final hologram. Both methods violate the required correspondence between capture, recording, and viewing geometries.

stereogram viewer sees that image is now broken. It makes no physical sense for a viewer moving along one path (the viewing zone defined by the stereogram plate) to be shown images that were captured by a camera moving along another path (the camera moving in an arc). Because of the plate is flat, the camera *must* move along a straight, not curved, track.

Another way to recenter the images is to move the camera along a straight track, but to swivel or “toe in” the camera to face the center of the object frame before every exposure. This method suffers from the keystoning distortion described earlier because the film plane is no longer parallel to the object plane during capture, while the film plane is indeed parallel to the projection plane during holographic recording. However, the camera is at least in the same physical location when the image is captured as the viewer as the viewer will be when the image is viewed, so that the perspective distortion is less severe. Unfortunately, the plane of focus of the camera does not correspond to the plane of the projection screen, so this stereogram geometry is prone to image blur if the taking lens is not sufficiently stopped down.

To summarize the previous discussion, there are two common methods of producing a distortion-free holographic stereogram from a sequence of images: the first in which the projection frame is located directly in front of the slit during each exposure and the plate translates with respect to it (the "simple camera" geometry), and the second in which the screen is centered in front of the plate throughout all the exposures and the slit moves from one exposure to the next (the "recentering camera" geometry). The first method has the advantage that the camera needed to acquire the projectional images is the easier to build, but the input frames tend to vignette objects that are close to the camera. The second method requires a more complicated camera, but moves the plane of the image where no vignetting occurs from infinity to the object plane. The camera complexity of this method is less of an issue if a computer graphics camera rather than a physical camera is used.

Advantages of recentering-camera stereograms

The projection frame in a recentered camera stereogram forms a "window" of information in space, fixed with respect to the stereogram and located at the depth of the projection plane. The usefulness of this fixed window becomes important when the slit hologram is optically transferred in a second holographic step in order to simplify viewing. To maintain the capture-recording-viewing correspondence in any stereogram, the viewer's eyes must be located at the plane of the slit hologram. When the stereogram is a physical object, the viewer's face must be immediately next to a piece of glass or film. However, a holographic transfer image can be made so as to project a real image of the slit master hologram out into space, allowing the viewer to be conveniently positioned in the image of the slits, as shown in Figure 2.13.

A transfer of a stereogram is directly analogous to a transfer of a continuous master

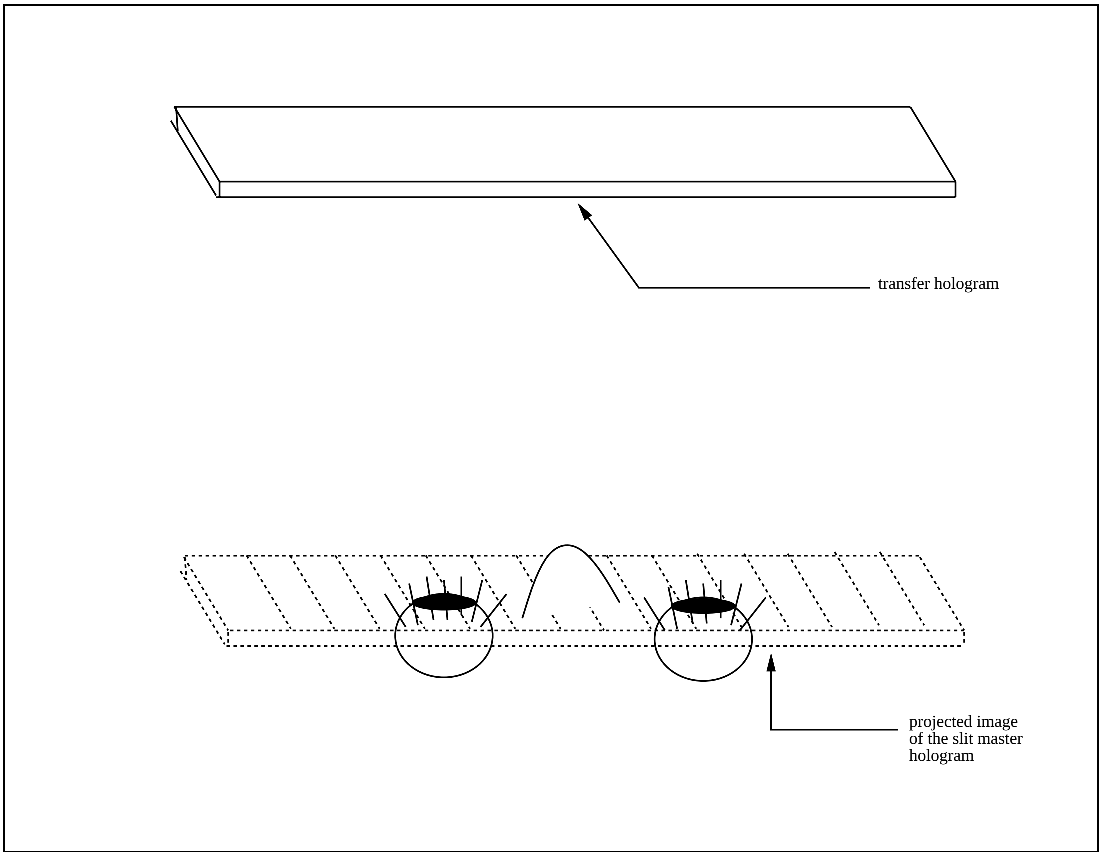

Figure 2.13: If a transfer hologram of a slit master stereogram is made, the viewer can stand at the plane of the slits, or even pass through it, without suffering facial lacerations.

hologram of a real object. The transfer of the slit master is done by illuminating all the slits of the master from the same angle from which the reference beam arrived but in the opposite direction, forming a perfect conjugate illumination source. Each slit then projects a real image of the projection screen out to the projection plane. The transfer hologram is also placed on this plane, and records the light that strikes it from the master. The transfer hologram can then be illuminated with a beam that is conjugate to its reference source, projecting an image of the master slits out into space at the projection plane-master separation distance. An important motivation for transfer holograms, besides the convenience of viewing, is that white light illumination of the hologram is possible.

In the simple camera stereogram, the images of the projection frames that the slits of the master project to the transfer during mastering are shifted with respect to each other because each frame image is centered directly in front of its slit. Thus, the frames cannot completely overlap each other. In the case of the recentered camera stereogram, however, all images of the projection frames precisely overlap on the projection plane.

When the transfer hologram is made, the position of the transfer plate on the projection plane determines what window of that plane will be visible to the viewer. In the recentering-camera stereogram, this window is clearly defined by the projection frame: all information from all slits overlaps there, with no data wasted off the frame's edge. In the simple-camera case, some information from every slit (except the center one) will miss the transfer's frame and as a result will never be visible to the viewer. This fact, which can be thought of as a double vignetting of two windows (one defined by the transfer plane and the other located at infinity), eliminates any advantages that a simple camera stereogram has for imaging distant objects. All extra information captured with the simple camera is cut off by the transfer's edges. More noticeable, though, are problems due to the fact that most of the slits do not fully cover the transfer frame with their image of the projection screen.

The area of the transfer frame that overlaps the projection frame projected by a single slit defines what part of the transfer hologram will bear image information when the viewer looks through the image of that slit when the transfer is viewed. The rest of the transfer, having no information about the projection frame for that slit, will appear black. In effect, as the viewer moves from left to right through the transfer hologram's view zone, an image window appearing at infinity slides across the transfer plate, also from left to right. Again, the information that is contained in the window located at infinity that does not fall within the transfer frame is never visible and thus was captured for naught. The limitations of a simple camera stereogram transfer are pictured in Figure 2.14. In contrast, in a recentering camera stereogram master, all parts of the transfer hologram are image bearing from all viewpoints. The "everything seen, nothing wasted" property of recentering camera holographic stereograms is especially important when image capture is difficult or troublesome, which is true, for example, if the images needs to be computed frame by frame. Computer time is better spent calculating visible parts of images than invisible parts; the recentering camera geometry defines precisely the areas of visibility when the transfer is viewed.

So, for transfer holography, a recentering camera has significant advantages over the simple camera. For one-step holography, the shortcomings of a simple camera stereogram are not as severe, and the simplicity factor may weigh more heavily when deciding between the two methods. The following chapters will rely, however, on the fact that the two camera types are actually quite similar, and that both types present distortion-free images to the viewer. The differences between the two types can be downplayed in order to show more important concepts.

Wavefront approximation in stereograms

So far, the discussion of stereograms has centered completely around viewer perception, with no regard to how light from the stereogram approximates light from the real object.

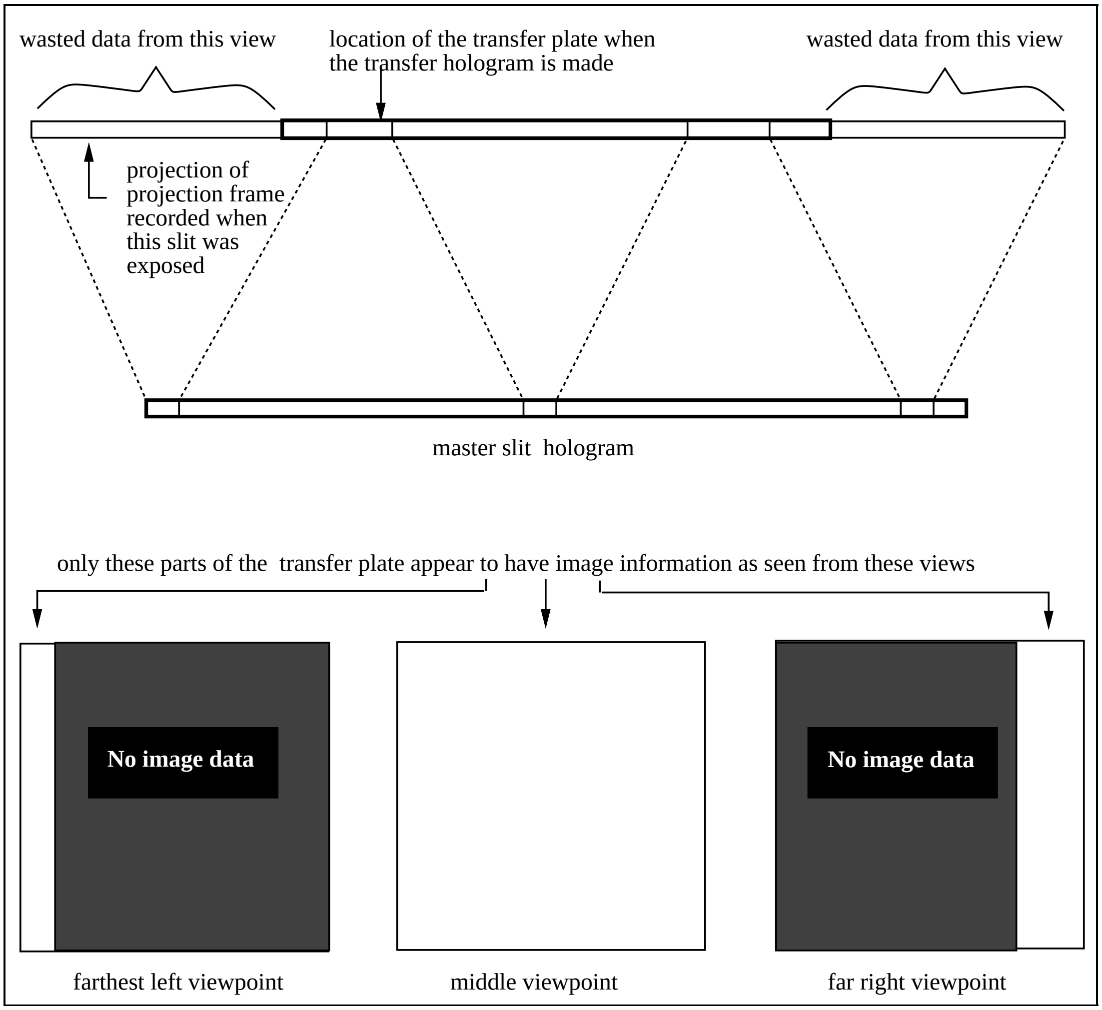

Figure 2.14: In a simple camera stereogram, the projection frames of many slits do not completely overlap the transfer frame when a transfer hologram is made. As a result, a significant amount of image data is never seen and thus wasted, while no image is visible in parts of the transfer from certain viewpoints because no data about those parts was recorded.

A stereogram mimics the light emitted by the original scene using a piecewise wavefront approximation. This approximation can be investigated using the simple stereogram model described above. Only one spacing of wavefronts need be considered because of the monochromatic nature of the illumination source. For simplicity, a recentering camera stereogram model will be used in the following explanation.

Wavefront of points on the projection screen

An isolated point in space radiates light in a spherical pattern. The radius of the curvature of the wavefront is a measure of how distant the source is from the observer. The response of the eye's focusing mechanism is the visual system's way of measuring a wavefront's curvature. Because a hologram records the direction of light striking it during exposure, the wavefront of a perfectly illuminated hologram of a single point is also spherical, with a radius of curvature identical to that of the original point.

If all the points in a scene are restricted to a single depth plane, all the points will produce wavefronts of identical curvature. Each slit of a stereogram is a hologram of just such a planar object, with all objects points for each slit being restricted to the projection plane. A recentered-camera stereogram in which each slit is exposed to an unchanging test pattern of a white point on dark background will, when illuminated, produce almost the same wavefront as a true hologram made of the same pattern. The only difference between the two wavefronts is a series of stepwise phase changes between segments of the stereogram's wavefront, a difference to which the eye is not very sensitive. Thus, the holographic and stereographic test patterns are almost identical.

Wavefronts of points at other depths

A stereogram would be rather uninteresting, though, if all objects in its scene were restricted to a single depth plane. Every slit of the hologram sees, as does the real hologram of the projection screen described above, a collection of wavefronts radiated by points on the screen. The radius of curvature of these wavefronts varies based only on the two-dimensional position of the radiating part of the screen. While each slit is only exposed to a two-dimensional screen, the resulting stereogram produces an approximation to the light emitted by objects in a three dimensional field. The wavefront approximation is composed of pieces of the wavefonts of the fixed, regular pattern pattern radiated by points on the screen. The example holographic stereogram has horizontal parallax only; as a result, the piecewise approximation is only complicated in the horizontal direction. Vertically, the wavefronts for all points in the three dimensional field appear to be, and in fact are, emitted from the projection plane, independent of the intended depths of the points.

Horizontally, the only depth location at which an object point may lie so that a stereogram emits a wavefront of the same shape (except for phase variations) as the original point is at the projection plane. For any slit spacing, each slit sees the same image of the screen at the same absolute position, so the slit structure of the hologram does not affect the shape of the reconstructed wavefront. However, if an object point lies at a depth different from the projection plane, the accuracy of the wavefront produced by the stereogram cannot be maintained.

For example, if an object point lies between the projection and slit planes, the projection of the object point onto the projection plane moves right to left in the frames used to expose the slits in left to right order. The projection's motion, in other words, moves across the projection plane in a direction opposite the viewer's motion. The rate of this motion is inversely proportional to the distance from the viewer to the object point, minus the rate at which the projection screen itself appears to move. Equation 2.1 describes this motion.

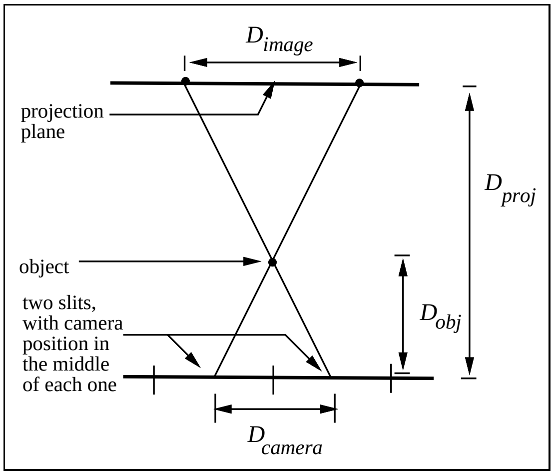

$$D_{image} = \frac{D_{proj} - D_{obj}}{D_{obj}} D_{camera} \tag{2.1}$$

The moving image of the object point on the projection screen provides a series of wavefront segments that, in total, approximates the point's true wavefront. An object point in front of the projection plane would produce a horizontal curvature smaller than a point located at the projection plane. The image points that make up the approximated wavefront are all located on the projection plane. So the stereogram's wavefront approximation is composed of pieces that are "flatter" than the true wavefront. An illustration of this approximation is shown in Figure 2.15.

Similarly, an object point located behind the projection plane will produce a greater radius of curvature than do the points on the projection plane. In the projected images for each slit, though, the object point's image moves in the same rather than the opposite direction from the viewer. When the stereogram image is reconstructed, the resulting wavefront will have "bumps", being composed of wavefront subsections from the projection screen that have a smaller radius of curvature than the actual point. Such a wavefront looks like the one shown in Figure 2.16.

Clearly, the further the object point is from the projection plane, the less similar the curvature of the wavefront segments is to the wavefront being approximated, and thus the worse the accuracy of the approximation becomes. However, independent of the quality of the approximation, there are no restrictions on the location of object points: points at

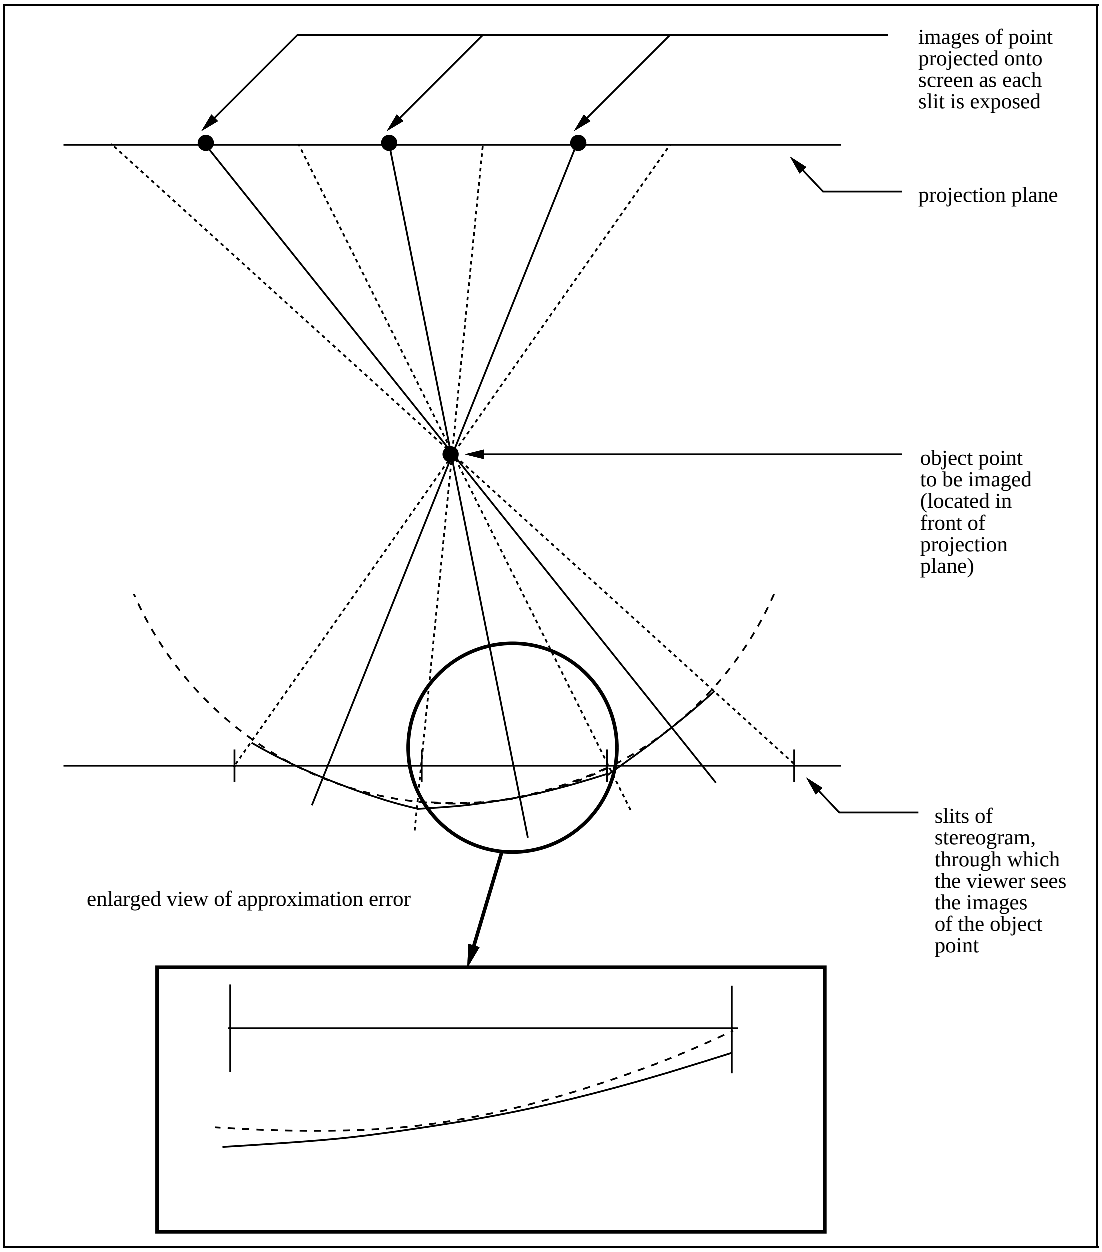

Figure 2.15: Object points in front of the projection plane are approximated by wavefront segments of lesser curvature than the actual wavefront.

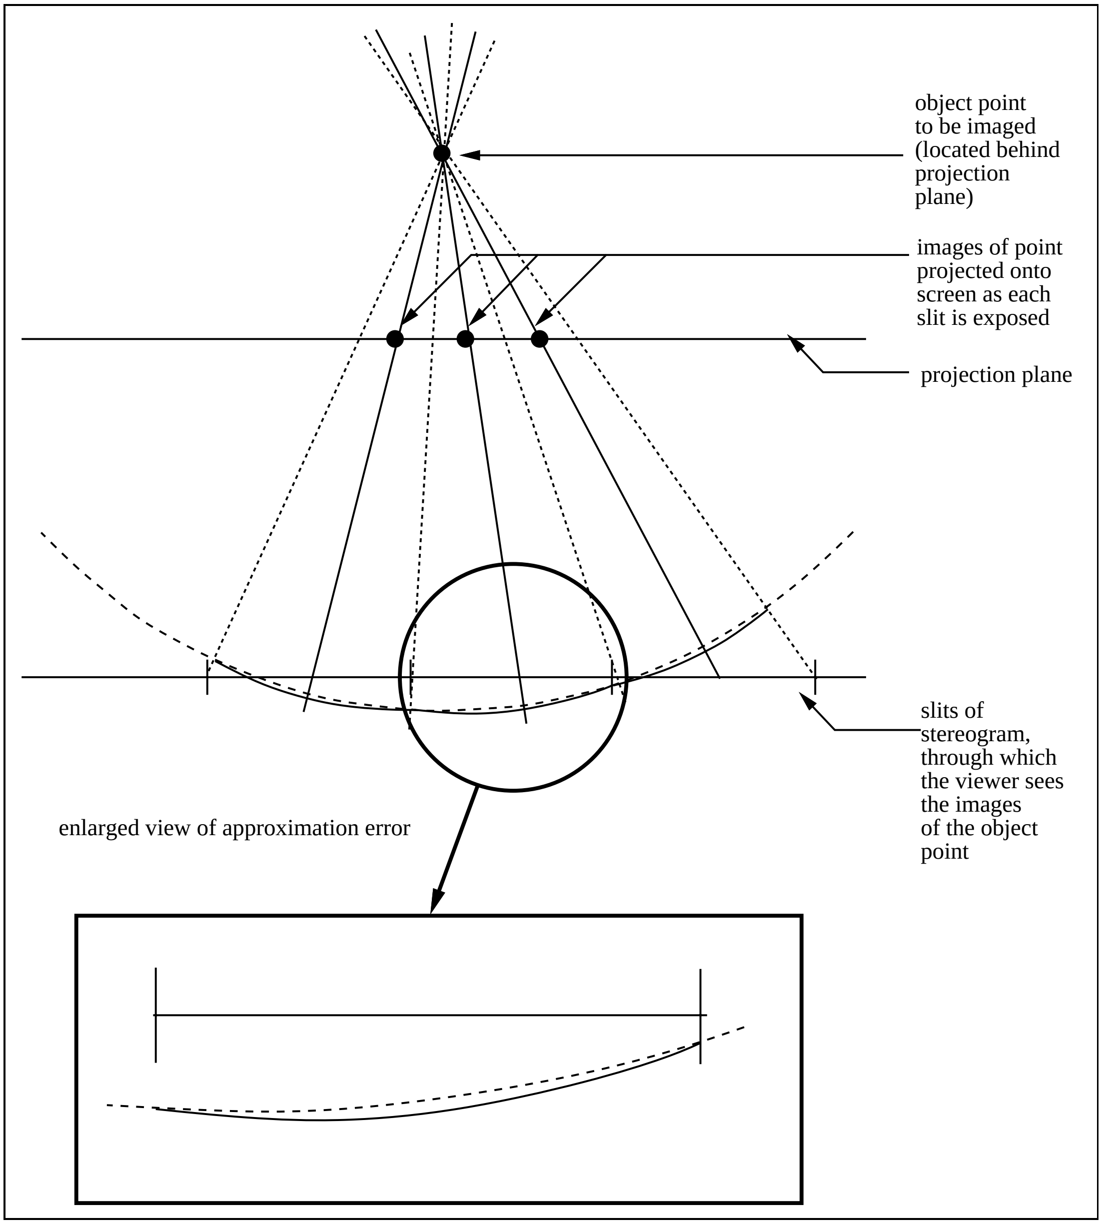

Figure 2.16: Object points behind of the projection plane are made up of wavefront segments of greater curvature than the actual wavefront.

any depth can be approximated to some degree. Combined with superposition arguments, this fact assures that objects of any spatial extent can be reconstructed using a stereogram, although image artifacts may be present.

Minimizing error in wavefront approximation

The errors of stereographic wavefront approximations and their direct correspondence to the distance from the projection plane establishes a fundamental rule of stereograms. To minimize the effect of wavefront errors, position the apparent horizontal focus of the projection screen at the plane of the image of greatest visual importance, or position the plane to straddle several important objects at different depths. While all the object points on one side of the plane will have segments of too great a radius of curvature, and points on the other side will have too little, the overall error over the set of object points will be minimized.

Fortunately, the degradation of wavefront accuracy as object points are positioned far from the projection plane is a gradual one; compromises in projection plane position can be made without greatly affecting the quality of the image. The image quality tradeoff between position and accuracy of wavefront is very similar to depth of field tradeoffs in a standard variable aperture camera lens: objects on either side of the plane of focus are somewhat blurred, but the entire object appears reasonably sharp. Depth of field is a term that traditionally refers to the rate at which the amount of image blur increases as a function of distance from the plane of focus and is directly controlled by the size of the lens' aperture. The next chapter will elaborate on this connection between the depth of field and stereograms.

Chapter 3

Bandlimiting

In the last chapter, two different approaches to the analysis of stereograms were presented: one based on purely visual considerations, the other on the accuracy of a stereogram's approximation to an object's true wavefront. The wavefront argument showed that a stereogram can duplicate the wavefront curvature only of objects at a depth corresponding to the projection plane; the wavefronts of points at all other depths are approximated using a piecewise circular wavefront. This chapter will investigate the visual cost of this approximation, its relation to the discrete nature of stereograms and other sampled systems, and explore ways of preventing approximation-caused image artifacts.

Continuous and discrete images

The simple stereogram model of the previous chapter was immune to approximation artifacts because the viewer was constrained to see through one and only one slit at a time. The "seams" between adjacent pieces of the wavefront never entered the eye, so they were never seen. This viewer constraint will now be relaxed, leaving the viewer free to look through

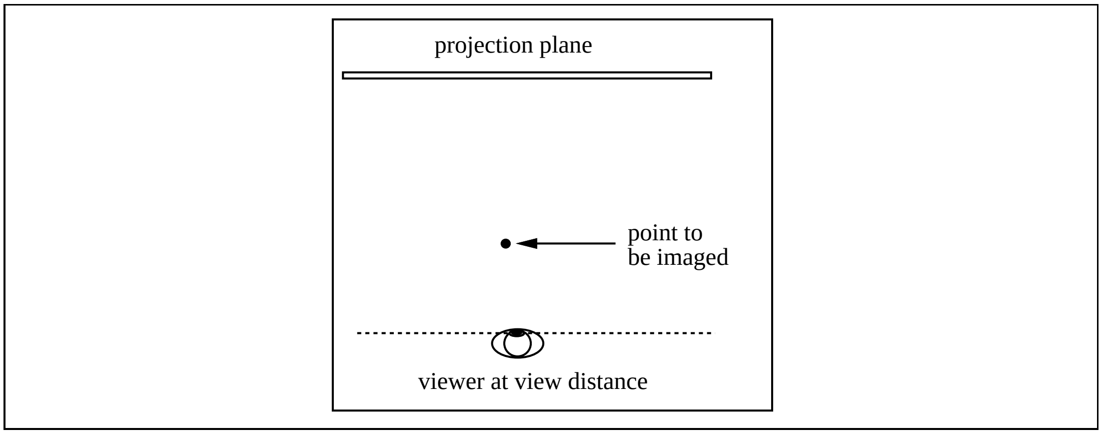

Figure 3.1: The stereogram test geometry. A single point in space lies between the viewer and the projection plane.

any part of the stereogram. No longer must the eye necessarily look through one and only one slit; as is often the case in “real-world” stereograms, the eye position may straddle the border between two adjacent slits. Imagine, for example, that a point in front of the image plane is visible in the projection screens used to expose two adjacent stereogram slits. The geometrical relationship between the viewer, the projection screen, and the object point is shown in Figure 3.1. When the viewer’s eye is completely within either slit, the viewer will see an image of the point as captured from the camera position corresponding to the middle of each slit. The difference between the real object with its continuous wavefront and the stereogram with its sampled wavefront becomes evident when the viewer whose eye is completely in one slit moves incrementally towards the other slit. Figure 3.2 is a comparison of the visual difference between the continuous and sampled system.

Were the viewer looking at a real object, such an incremental motion would cause the object to appear to move slightly with respect to the projection plane. On the other hand, in a stereogram, such a small shift in the apparent position of the object point is not possible because the object is imaged from only finitely many equally-spaced camera positions along the camera track. A point cannot appear to move any less than the hop its

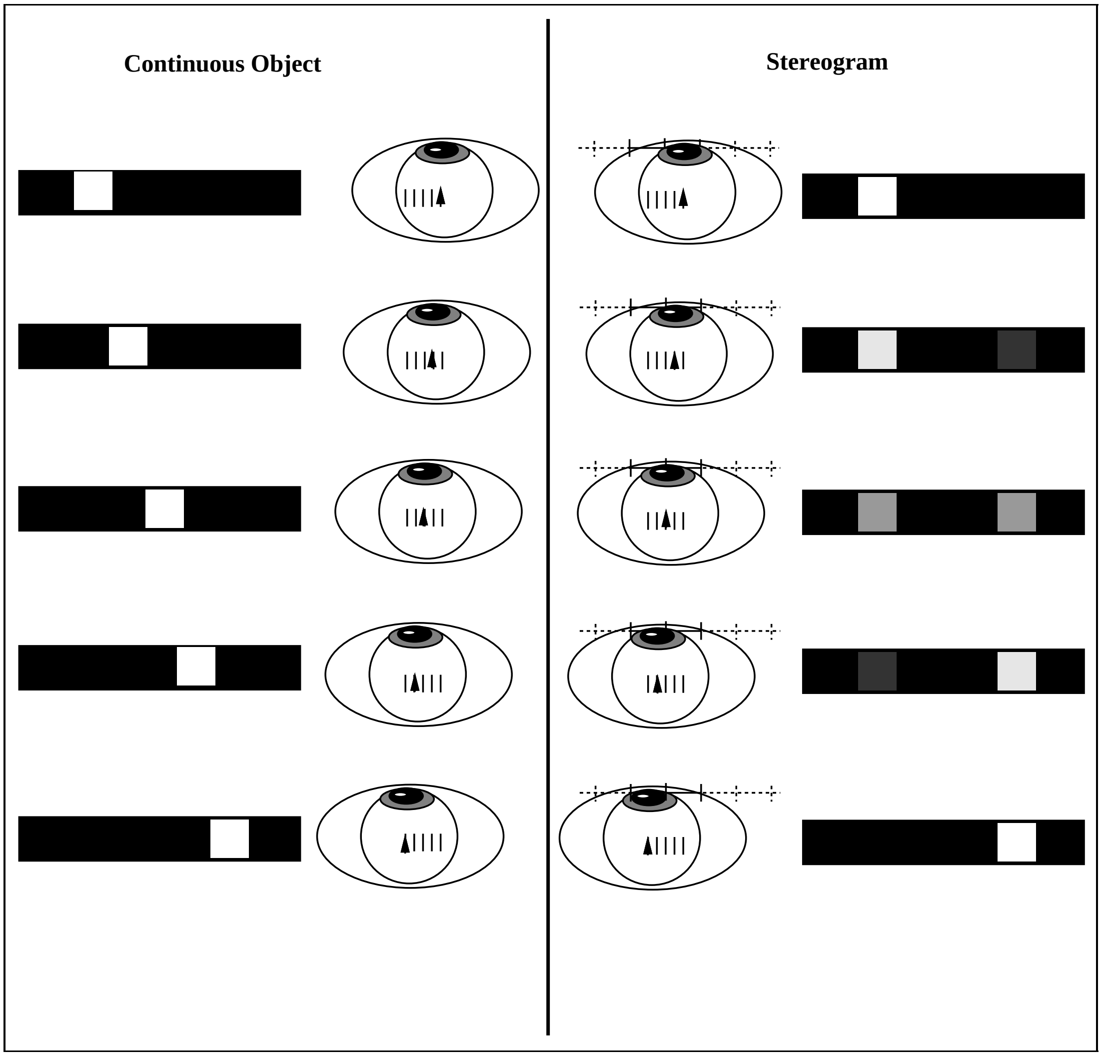

Figure 3.2: Comparison between viewing a continuous image and a stereogram of a point. In the continuous system, the image of image appears to move gradually as the viewer moves. In the stereogram, the image of the point can only appear at discrete locations and fades from one position to the next as the viewer moves.

image takes from one camera view to the next. Instead of a small change in position, a different effect occurs. A faint image of the object as seen from the second slit's camera position becomes superimposed with the image from the first slit. The separation of the old and the new image is proportional to the object point's distance from the projection plane; it is, in fact, the distance that the image of an object point hops from one slit view to the next. As the viewer moves, both old and new images of the object appear to move as if they are stuck to the projection plane (which, in fact, they are). The relative contributions of the first and second slit's images is dependent on how much of each slit falls within the confines of the pupil of the viewer's eye. For instance, a viewer exactly straddling the border between the two slits would see two half-intensity images of the object point, one in its old and one in its new position. The image continues to change as the viewer moves until the new image completely fades in and the old one fades away. This analysis ignores the effect of diffraction at the slit edges, a subject for future study.

The visual effect of crossing the slit boundary can be explained by examining the shape of the stereogram's wavefront approximation. A viewer's eye completely in one slit sees a wavefront emitted from the projection screen for every object point. The eye's focusing mechanism determines the depth of the object based on the shape of the wavefront, so monocular focus cues indicate that the object is at the projection plane. Viewer motion toward the second slit brings a tiny piece of the wavefront emitted by the second image of the object point into view. This second wavefront is centered around and thus seems to come from the object point's second image on the projection screen. No continuity is guaranteed between the slope of the first and second wavefronts. As the viewer continues to move into the second slit, more and more of the second wavefront and less and less of the first enters the pupil, until finally the discontinuity between wavefront sections falls outside of the eye, leaving visible only the second wavefront and thus the second point.

Perception of stereogram artifacts

Unfortunately, the jumpy motion that the object follows as the viewer moves across the stereogram may not be readily interpreted as continuous motion. When the eye's pupil straddles the slit boundary and the object point is seen as it appears from two different camera positions, extraneous information may separate the two images of the point. This information is composed of image detail on either side of the object point in the perspective views used to expose the two slits. This image detail may also be replicated, compounding the problem. If this unwanted extra information is present, the object point does not seem to be one continuous, solid object, but rather two almost identical objects separated in space.

In real stereograms of complicated scenes, this replication of object detail can be very noticeable. For instance, parts of the image far from the projection plane appear vertically striped, with horizontal detail on distant objects and backgrounds being the most visibly incorrect. The visual effect is similar to seeing a person walking behind a picket fence; the gaps between the images of the object move as if fixed to the projection screen and seem to occlude the object behind. A stereogram slit width on the order of the width of the pupil of the human eye, commonly used in practice, implies that the viewer's eyes will almost always straddle a slit boundary, so any unwanted image artifacts that occur are visible from almost any viewer position.

Stereograms and aliasing

The discrete nature of the stereogram is the underlying cause of this image artifact. Each slit, and the corresponding camera image, is a single sample of an object's appearance as seen from a particular location. Each image, then, is a record of the relative positions of all the points in an object. The stereogram image sequence can in turn be thought of as a sampled record of an object's apparent velocity through space as the viewer moves over a range of viewpoints. If this velocity is sampled at too low a rate, or the apparent position of some points on the object change too rapidly (their velocity is too high), the stereogram suffers from image artifacts. These artifacts fall into a general category, called aliasing artifacts, that plague all insufficiently-sampled discrete systems. Stereogram aliasing is otherwise known as inter-view aliasing because it occurs between adjacent camera views. The severity of aliasing is directly proportional to the stereogram's sampling rate and is controlled by the number of stereogram slits per unit distance.

To better see this relationship between aliasing and sampling rate, imagine that the viewer's pupil exactly straddles by a slit boundary, so that an object point off of the projection plane will appear as a double image. Now, double the number of slits of the stereogram, which is the same as doubling the sampling rate. The visual result is shown in Figure 3.3. If the viewer is still straddling two slits, the viewer will still see two camera views simultaneously, but the views come from cameras now one-half the distance apart as were the previous images. Because the distance the taking camera moves from view to view is directly proportional to how far an object point appears to move from view to view, the two images of the object point that the viewer sees are half as far apart as were the original images. As a result, the amount of unwanted intervening detail has been cut in half.

Eliminating aliasing artifacts

One way, then, to reduce the problem of stereogram aliasing is to sample at a high rate by capturing more views and exposing more slits. At some level of sampling, diffraction, imperfections of the holographic recording process or limitations of the human eye will obscure the presence of the artifact. This approach is straightforward, but impractical and inelegant. First, increasing the number of camera views and slits greatly reduces the simplicity of stereogram production. Narrow slits are difficult to manufacture and are prone

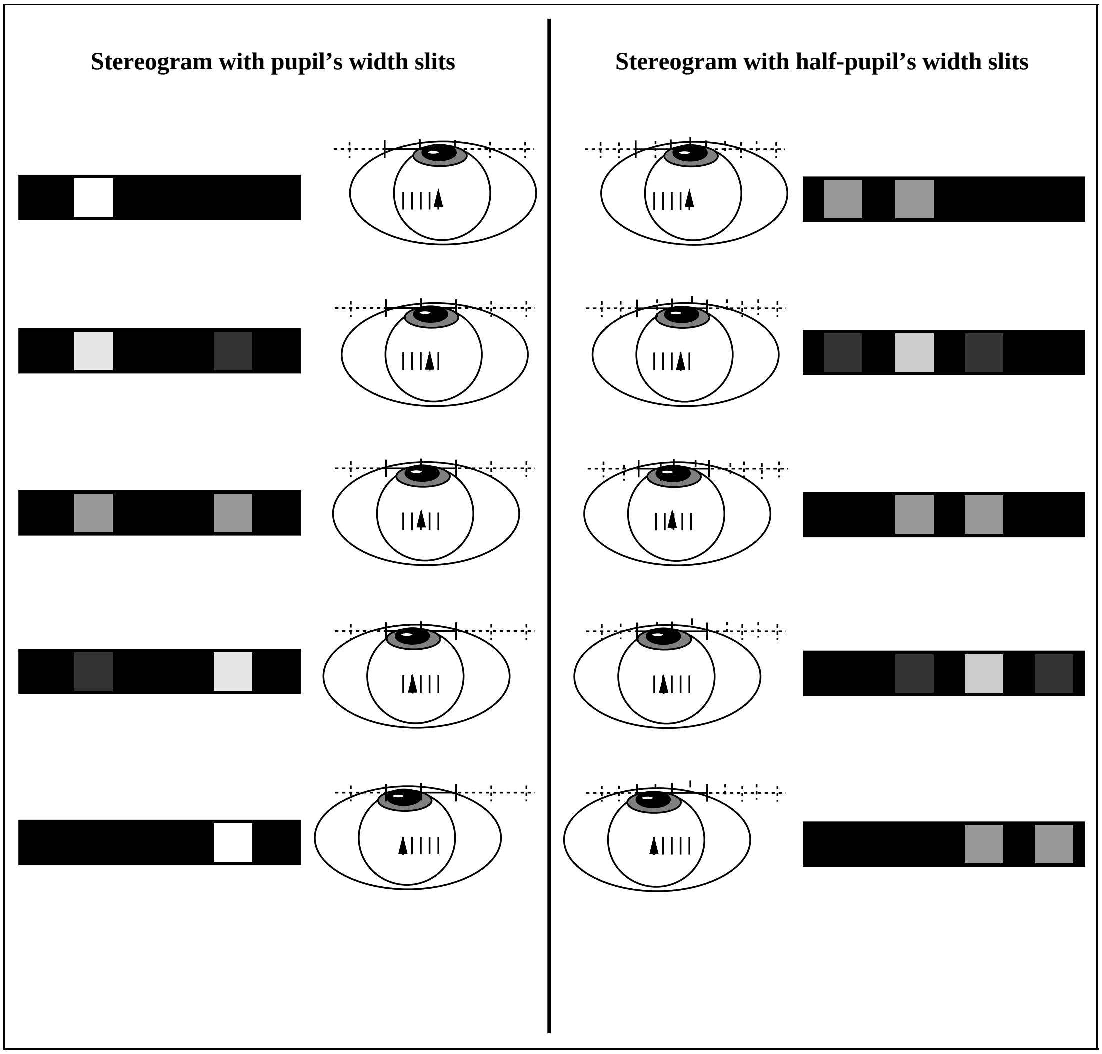

Figure 3.3: If the number of slits of a stereogram is doubled, the eye sees several views taken from cameras closer together. The gap between the images of the object point seen from one viewpoint is half as large.

to diffraction effects. Lastly, increased sampling does not eliminate the problem, it merely scales it down; for any finite sampling rate, a point sufficiently far from the projection plane will still exhibit aliasing artifacts.

Bandlimiting

In other sampled systems, such as digital audio and video, aliasing is eliminated by the process of bandlimiting, the removal of fluctuations in a signal of a frequency higher than that which can be accurately captured by the sampling signal or medium. The same type of approach can be used for stereograms. Up to this point in the analysis, the extent of the images of object points has not been well defined: points could be arbitrarily small. As a result, the frequency of occurrence of object points could be arbitrarily high. The process of bandlimiting in holographic stereograms consists of selectively reducing the spatial frequency of points in the images of the scene by increasing their horizontal extent. The high spatial frequencies are thus filtered out of the stereogram's input views. Proper filtering assures that the image of a point in one frame smoothly blends into the image of the point in the next frame, with no gap or extraneous detail ever falling between the two images. The following example illustrates the visual effect of bandlimiting, and follows Figure 3.4.

Return to the case in which the viewer's eye straddles the slit boundary of a stereogram of single point located off the projection plane. Recall that the distance that separates the two images of the the point from the two slits is the distance that the image moves from one projected view to the next. Now, imagine that the two point images expand horizontally in both directions until they just touch. Each point is now as wide as the distance that originally separated the two image points. Another way to calculate this width is to analyze its endpoints. One end of a point's image extent is where the point appears to be from the left border of the slit. Similarly, the other end of the point is where the image appears to be

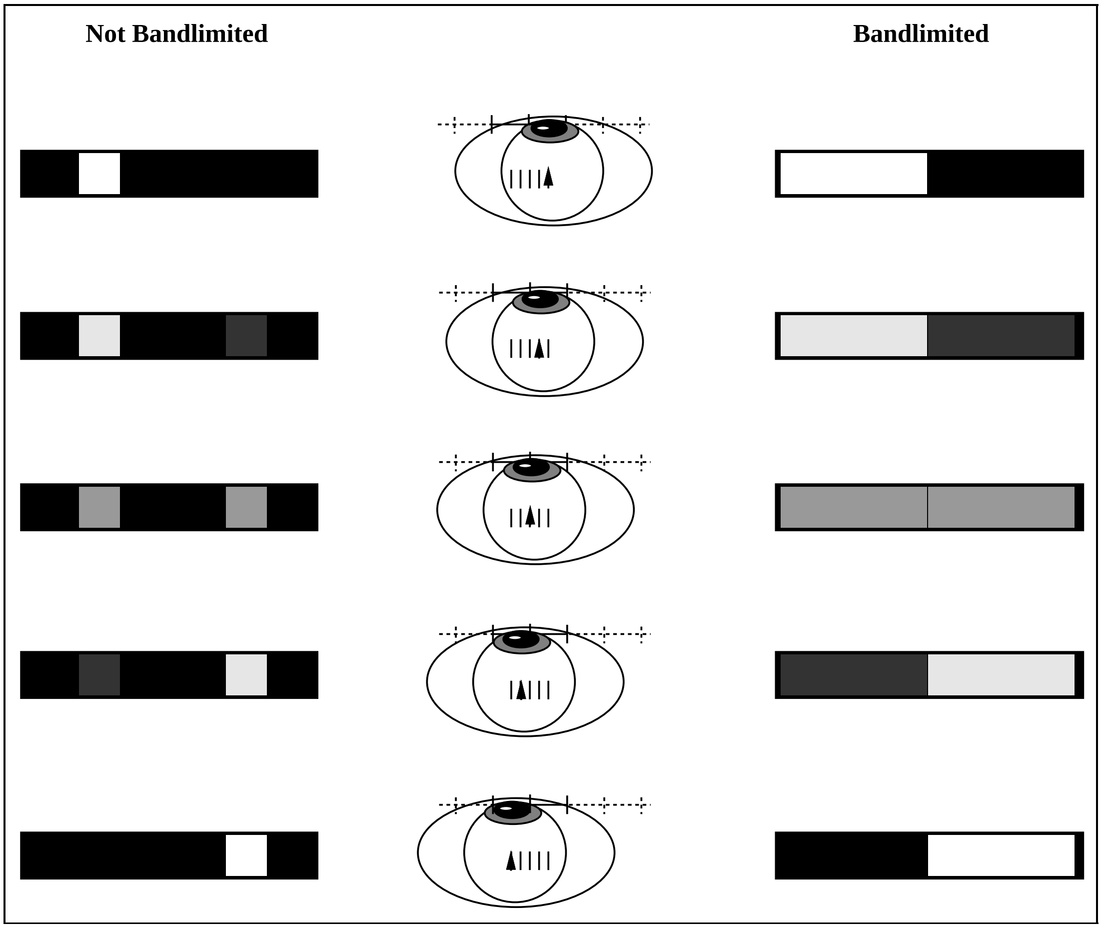

Figure 3.4: Bandlimiting of a stereogram image eliminates the gaps between the images of object points when the viewer's eye falls within more than one slit.

from the rightmost side of the slit. The center camera view, which until now has been the only image representing the slit, produces an image in the middle of the two endpoints. The required point width represents the object point as seen from all camera positions within the horizontal confines of the slit.

In the stereogram model, all slits have the same fixed width, so the distance that the image of an object point hops from view to view is a constant over all views. Because this distance is equal to the width of the filter required to bandlimit the point, the filter width is also a constant for any point. Just as hop size is directly proportional to the distance between the object point and the projection plane, so too is the filter width. In short, stereographic bandlimiting requires a filter whose width varies with the depth of the object point being filtered.

Only at the projection plane, where the stereogram's wavefront approximation exactly matches the wavefront of the actual object point, can the extent of an image point be arbitrarily small. The image of objects at all other depth planes must be filtered by some amount in the direction of sampled parallax so as not to exceed the maximum spatial frequency allowed by the stereogram. The amount of filtering required for object points at a given depth plane can be measured in two ways. The first, using the approach discussed above, is to measure the minimum size of the point's image on the projection plane. This technique implicitly incorporates camera perspective when calculating the filter size. For an object D_{obj} from the viewer, the minimum width of the object's image is given by the equation shown in Equation 3.1, which is the same one that describes view-to-view hop distance.

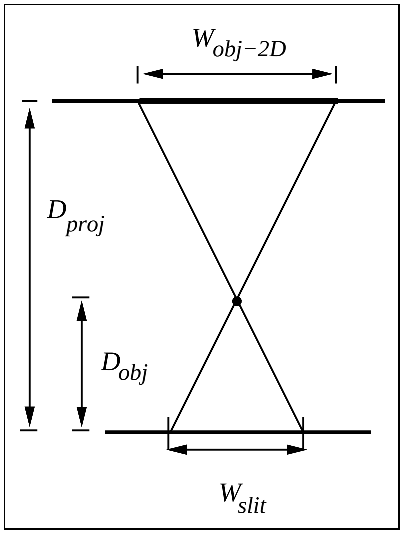

$$W_{obj-2D} \geq \left| \frac{D_{obj} - D_{proj}}{D_{obj}} \right| W_{slit} \qquad (3.1)$$

The other way of measuring the required bandlimiting is to define the size of an object's extent in three-dimensional space, instead of the size of its two-dimensional projection on the projection plane. Such a method eliminates the effect of camera perspective. One useful landmark in determining this extent is that an object on the projection plane can have points of infinitely small size. Another landmark can be found by positioning an object point directly in front of and at the same position as the opening of the taking camera's lens, so that $D_{obj} = 0$. Equation 3.1 states that such an object's projection must be infinitely wide, a result supported by the reasoning that, were an object point to be located at the plane of a viewer's eye, that object point would obscure all other objects and cover the entire visual field. To obscure all other information, the object's three-dimensional extent must be at least as wide as the opening of the camera lens, or correspondingly, the width of the stereogram slit. Sampling theory also supports this reasoning: an object resolution at the camera/slit plane higher than the size of the slits is not sampled often enough to be recorded without loss of information.

The relationship between amount of filtering and distance between object and projection plane is linear. Given that infinite resolution is possible at the projection plane, and only slit-sized resolution is allowed at the slit plane, the general solution at any depth can be found. Equation 3.2 gives the formula for the three-dimensional extent of the object.

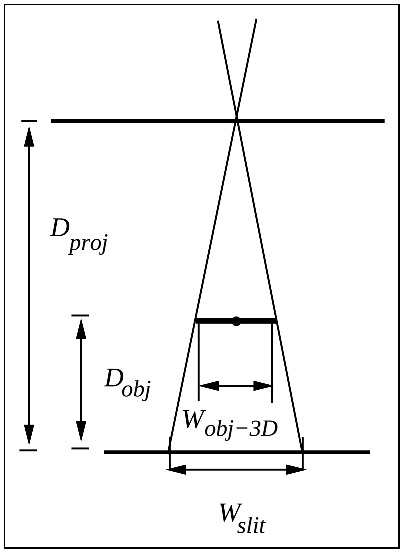

$$W_{obj-3D} \geq \left| \frac{D_{proj} - D_{obj}}{D_{proj}} \right| W_{slit} \qquad (3.2)$$

The validity of this formula can be shown by relating it to Equation 3.1. While W_{obj-2D} is the union of the images of an object point as seen from all camera locations across the width of a slit, W_{obj-3D} can be thought of as the width that an object point must be to subtend W_{obj-2D} on the projection screen when seen from a camera located on the slit plane. To project an image of size W_{obj-2D} onto the projection plane, W_{obj-3D} must be as shown in Equation 3.3. Combining Equations 3.1 and 3.3 yields Equation 3.2, assuring that the three are self-consistent.

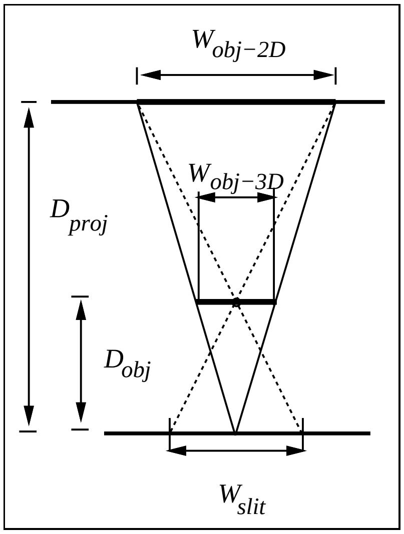

$$W_{obj-3D} = -\frac{D_{obj}}{D_{proj}} W_{obj-2D} \qquad (3.3)$$

Bandlimiting and wavefront approximation

If all points in the image field of a stereogram have spatial frequencies lower than the maximum permitted for the depth plane on which they lie, then the stereogram is bandlimited. Analysis of the wavefront of a bandlimited stereogram demonstrates why bandlimiting eliminates wavefront anomalies. In the previous chapter, the approximate wavefront for a single point emitted by the stereogram had discontinuities in curvature where the wavefront segments joined. In contrast, in a bandlimited image, each slit presents each object point not as a single wavefront, but as an incoherent sum of many wavefronts, each wavefront emitted from a portion of point's bandlimited image. Continuity of slope between the wavefronts emitted by adjacent slits is guaranteed because the far rightmost portion of a point's image in one slit contributes a wavefront identical to that of the leftmost portion of the next slit. In effect, bandlimiting replaces each wavefront with the smallest range of wavefronts that will smooth out the cornerlike discontinuities of the stereogram's wavefront approximation. The effect of bandlimiting on the wavefronts of two slits is shown in Figure 3.5.

Practical bandlimiting

Several observations can be made about the bandlimiting equations. First, the formulae and the reasoning behind them apply to points on either side of the projection plane. Furthermore, the equations include only implicit references to the camera's location in depth; the same bandlimiting factors apply to objects behind the camera as well as in front of it. At this point in the discussion, the camera and the viewer are located at the same plane, so an object located behind the camera would also be located behind the viewer and as a result would not be a very intuitive or visually useful object. However, the types of stereograms to be discussed in the next chapter allow the slit plane and the viewer plane to be different, so the viewer may indeed be interested in seeing objects that appear behind

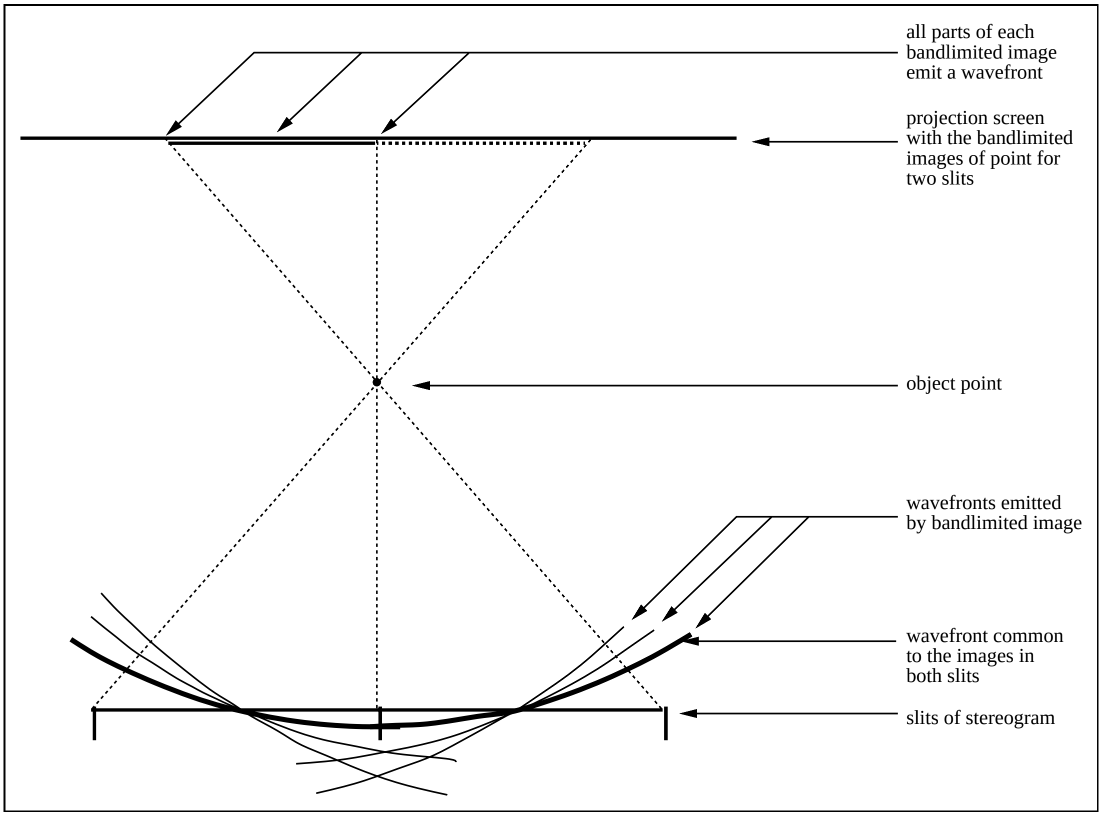

Figure 3.5: A bandlimited image of a point for each slit emits a range of wavefronts. This range is wide enough so that the far extreme wavefront of one slit matches the other extreme wavefront of the the next slit. The shared wavefront for the two slits in the above picture is emboldened.

the slit plane.

Bandlimiting and depth of field

Another observation about bandlimiting that can be drawn from the above equations is that a stereogram's depth of field is not only similar, but identical to, the depth of field of a camera lens. A camera lens captures the range of perspective views that enter its lens aperture or "entrance pupil". The lens aperture fixes the effective diameter of the optic, determining depth of field. In the image that the lens forms on its focal plane, every point in the scene is replaced by (or, more correctly, convolved with) an image of the aperture of a size that varies linearly with the distance of the point to the plane of best focus. To achieve proper bandlimiting in a stereogram, all points in the imaged field must be convolved with an image of the viewing aperture (the slit), of a size linearly related to the distance from each point to the projection plane, in the direction of sampled parallax (here, horizontally).

The finite aperture camera

The similarity between the effect of a finite-sized aperture on photographic capture and the effect of a finite-width slit on stereogram viewing suggests one way to perform the necessary stereogram bandlimiting using a physical camera. If the capturing camera has an effective aperture width equal to the slit width of the stereogram to be produced, and the lens is focused on the plane of the object corresponding to the eventual projection plane, all object points will be imaged as the appropriately sized disk on the camera's film plane. The width of each disk is precisely the correct width to bandlimit the stereogram because of the lens-slit aperture correspondence. Seen another way, all possible perspectives of the object as seen from any location along the camera track now enter the camera lens and are recorded in some camera view. Perspectives that the camera sees simultaneously are summed to form a single image, and no perspective contributes to more than one sum. This summation satisfies the stereogram requirements that all perspective views that could be seen in a slit area be presented together, and that no views be presented twice. Once again, correspondence between capture and viewing controls the accuracy of the final image with respect to the original, in this case with regard to image artifacts.

For practical “real camera” stereography, however, such a literal correspondence between capture and recording geometries is often inconvenient and difficult. Matching lens aperture to slit sizes is hard for several reasons. First, the size of the taking camera’s aperture controls not only depth of field but also the amount of light passed to the film; one method of exposure control is lost by fixing the aperture size. (In HPO stereograms, the vertical size of the aperture can still be used to control exposure.) And while a fixed pupil’s-width camera aperture may be simple to develop for a stereogram with a fixed 1:1 object to image magnification, a stereogram that scales the object down by a factor of ten or more would require an aperture several centimeters in width, and correspondingly expensive and unwieldy optics. For the range of scale factors and exposures where it does work, though, the finite aperture camera is a simple way to bandlimit.

The sweep camera

A more general, if approximate, approach to using large apertures to bandlimit natural scenes is to simulate the effect of a large lens opening by using a small, moving one. Recall that when a photograph is taken, the different perspectives that enter the lens are recorded incoherently on the film, with no memory of the direction from whence they came. Except for diffraction and film reciprocity effects, each perspective could be recorded sequentially on the film to produce an identical image. So, a slit-sized aperture could be simulated using an arbitrarily shaped smaller aperture as the camera is swept through the slit area. Mathematically, the resulting synthesized aperture would be the shape of the small aperture convolved with the path through which it travels during exposure. The effect of a simple sweep across the slit width is not exactly equivalent to the effect of a hard-edged aperture, but at least it includes all the perspectives of the slit. Using small lens apertures approves the approximation to the intended aperture.

A more serious issue is how to force the projection plane of the object to be the plane of arbitrarily high sharpness in the resulting images. After all, camera motion is usually associates with the general blurriness often seen in consumer photographs. A recentering camera is the key to solving this problem. A recentering camera is useful in stereography because it assures that the eventual projection plane of an object is imaged to the same position on the taking camera's film plane through all stereogram views. For the simulated aperture technique to work, the projection plane of an object must be imaged to the same position on the film *during a single exposure* as different perspectives are accumulated. If a recentering camera is used, both needs can be satisfied. As the camera moves along its track taking pictures, it no longer opens its shutter only at the middle of each slit. Instead, the shutter remains open as the camera sweeps through a slit area. When the camera reaches the end of the slit area, the shutter closes, the film is advanced, and the next image is begun. The camera's lens translation mechanism assures that the images of points on the projection plane fall on top of themselves throughout one slit's exposure. An object point located off of the projection plane will produce a line across the film as the camera and lens move through the slit, forming a bandlimited image of that object. The sweep of the image is only horizontal, so the vertical extent and resolution of objects is not affected.

Bandlimiting in synthetic images

If the stereogram's images are generated synthetically using a computer graphics camera instead of a physical camera, the "sweeping camera" bandlimiting method is even more straightforward. Computer graphic cameras are typically modeled upon ideal pinhole cameras, offering infinite depth of field and thus no camera blur. Because computer graphics are not necessarily bounded by the laws of optics, the film reciprocity and diffraction of the camera aperture that complicate the physical sweeping camera do not exist in the synthetic sweeping camera. The camera sweep can be done by rendering recentered synthetic camera views of the object from many camera viewpoints across the slit area, then combining them by averaging the intensities of the different images. The resulting images will be have the desired bandlimiting properties. This method is only practical if the different perspective views of the object can be computed with little computational expense. Some image artifacts may also be apparent because only a finite number of views of the object contribute to each slit's image. Alternatively, any other physically accurate computer graphics methods of simulating a finite width aperture camera can be used to produce synthetic bandlimited images.

Volume prefiltering

The synthetic sweep camera works by producing many high resolution views and combining them to produce one image of selectively reduced resolution. This process can waste much of the work that the renderer spends producing exact results of high quality and detail. Another synthetic bandlimiting technique, called volume prefiltering, performs the necessary bandlimiting on the three-dimensional representation of the object just once. If every point in an object volume is filtered using a depth dependent filter of width described in Equation 3.2, then the projection of any object point as seen from any camera location in the view zone by definition satisfies bandlimiting constraints. Once the volume is bandlimited, the stereogram capture camera can take one picture per slit, from the center of each slit, and obtain a set of correctly bandlimited views. Not all object primitives are conducive to volume prefiltering. Polygonal databases, for example, are poor choices for this method, while regularly sampled volume data sets are good candidates. Even more efficient methods of bandlimiting synthetic scenes may be possible.

Raster images and bandlimiting

A final convenience of computer generated images is that they are usually inherently bandlimited. If the image is composed of finitely sized pixels, scanlines, or other image primitives, no image detail can be smaller than the size of the primitive. Resolution may be further reduced because of limitations of the device that records the image for coherent projection in the holographic setup. Similarly, images can intentionally be rendered at a reduced horizontal pixel resolution, then interpolated back to the correct aspect ratio to guarantee a fixed maximum resolution. The drawback of this method is that it filters the image by a constant amount independent of depth, removing resolution from points that were already sufficiently bandlimited for their depth. Pixel-based bandlimiting techniques are simple and computationally inexpensive, but limited and crude. The quantized bandlimiting that rasterization implies, however, is an image property generally worthy of attention.

Changing slit size

Another constraint imposed by the simple stereogram model can now be relaxed. Until now, the size of the slits of the stereogram was restricted to be about the same as the size of the pupil of the human eye, as is commonly done in practice. This choice is a practical compromise. First, were the slits any wider, the object would noticeably jump from one position to the next as the viewer moves from side to side. Also, very wide slits would limit the resolution of points far from the projection plane. The pros and cons of sub-pupil width slits, on the other hand, are more subtle and have not been researched in depth. If the eye is filled with several slits horizontally, the image presented by each slit will simultaneously be visible. The effect of seeing several slits at the same time is not an averaging of the different views, however, because the different views come from different directions. This increased directionality is more faithful to the light from the original object. Wavefront analysis also shows that the wavefront approximation of a stereogram using sub-pupil width slits is more accurate because the piecewise segments are very small. Diffraction effects, on the other hand, establish a practical limit to how small slits can be made.

The gain that comes from this increase in accuracy seems to be an ability of the eye to focus on points off the projection plane as if they were at their true position in space. Recall that for pupil-width slits, all points appear to the eye's focus as if they were on the projection plane. As always, resolution far from the projection plane is improved with smaller slits due to the higher sampling rate. The tradeoff for this improved approximation is increased calculation and more complicated holographic exposure. Stereogram technology seems to currently be at the point where computation time is more precious than image fidelity, so little work with very fine slits has been done. Also, for horizontal parallax systems, the benefits of selective focus are dubious because of the eye's inability to tolerate astigmatism. And full parallax stereograms, where true focus would be useful, are still too difficult to generate to afford the luxury of many small stereogram elements. Further experiments must await more efficient image generation tools and more refined slit designs.

Changing slit shape

Another area that deserves further study is the prospect of stereographic slits without hard edges. The intensity profile of a slit need not be a binary function; rectangular slits are used mostly for simplicity of manufacture. From an image processing point of view, overlapping Gaussian-like slit intensity profiles make more sense than rectangular profiles because they lack high frequency edges, which is also desirable when the hologram is made because of diffraction effects of the slit mask. Bandlimiting synthetic images for non-rectangular slits is straightforward. Instead of simply averaging all perspective camera views, or convolving a volume with a variable width box filter, the averaging or filter function can be weighted to reflect the intensity profile of the slit. If adjacent slit profiles overlap when exposed on the stereogram, some perspective views will be included in the contribution for two (or possibly more) slits. These experiments are impractical for natural scenes; further research using computer graphics and careful holographic study may prove the merit of non-rectangular slits.

The photograph, the hologram, and the stereogram

The likely focusing ability of a finely slitted hologram and the relationship between stereographic and photographic depth of field both support the stereogram's role as a continuum between a photograph and a true hologram. Stereograms with wide slits cannot accurately present much three-dimensional information about an object; in the extreme, a stereogram with a single slit can present only a single plane and thus no depth at all. Bandlimiting compensates for limited three-dimensional resolution by incoherently blurring together the different perspectives seen through each slit so that all perspectives within the view zone will be represented in the stereogram. As more and more slits are added to a stereogram, more information is stored in a direction-preserving manner. Less information is recorded incoherently and so less filtering of the image is required. Ideally, if enough slits are recorded, the amount of directional information may be sufficient to allow the focusing of points at their true location. If so, incoherent or photographic blur is replaced with true three-dimensional blur. If other factors did not limit the minimum width of a stereogram's slit, the width could be reduced until the number of perspectives approximated that of a true hologram. At that point, all the optical properties of the original object would be replicated. Unfortunately, boundary effects of the slit aperture, such as diffraction, prevent a straightforward reduction of slit size. Clearly, though, the stereogram designer has a wide latitude in deciding how accurately the hologram should approximate a continuous three-dimensional image.

Chapter 4

Separating the Slits from the Viewer

The previous chapter built on the simple stereogram model, but relaxed it in two ways: the viewer's eyes could be positioned at an arbitrary position within the view plane (not just centered in a slit), and the stereogram's slit size could be different from a human pupil's width. This chapter will relax a different constraint on the simple stereogram model: the one that requires the viewer to be positioned at the plane of the slits of the stereogram. The goal of removing this restriction is to allow the creation of stereograms with an arbitrary view plane distance, freeing both the image designer and the viewer from optical constraints imposed by the holographic recording apparatus.

Viewer positions behind the slit plane are possible in the simple stereogram, but viewing locations between the slits and the projection plane are possible only through the use of two-step holographic techniques. Distortion-free two-step holography does not, however, alter any other relevant properties of the viewing geometry, so without loss of generality the beginning of this chapter will implicitly assume that the viewer is able to pass through the slit plane to viewer locations close to the projection plane. Issues peculiar to two-step stereography will be reserved for discussion at the end of the chapter.

Understanding the optical effects of moving the viewer to a different view distance requires another means of optical analysis called *ray tracing*. While wavefront analysis is useful when determining the small changes in the direction of light that proved significant in the stereogram's wavefront approximation, ray tracing's strength is in illustrating the general paths of light from large areas, overlooking small differences in direction and completely omitting phase variations. Ray tracing can be used to determine the image that each camera along a track sees, and thus what each projection screen should look like when each slit is exposed. It also shows what part of the projection screen of each slit is visible to a viewer at any one position. Distortion-free viewing requires that the rays from the photographic capture step and the viewing step correspond to each other.

Understanding stereogram distortion

The single point test pattern used in previous chapters is too simple to display the distortions experienced as the viewer changes position in depth. Instead, a new test pattern will be used. The three dimensional object used in this test pattern consists of three hollow squares all having the same height and width but located at different depths. One square is located on the projection plane, another is located one unit of distance directly behind the first, and the last is located one unit directly in front of it. The capturing camera's track is located three units from the projection plane, and two units in front of the front square. A recentering stereogram camera moves down this track, which is four units long and centered on the central axis of the squares. The stereogram in this example has many slits, so the camera will capture a long sequence of views, but the pictures captured from five judiciously chosen camera positions will be illustrated here. One view is taken from the center of the camera track. This view will be called view *C* for "center". Two more views are captured one unit to each side of view *C*; these views are called view *NL* (for "near left") on the left and view *NR* ("near right") on the right. The camera position for views *NL* and *NR* was chosen to be collinear with the edges of the squares of the test pattern so that in those images, the edges on one side of all three squares would appear to line up. The final two camera views, *FL* on the left and *FR* on the right (for “far left” and “far right”), are taken from positions two units from the center of the track. A diagram of the geometry of the test pattern stereogram is shown in Figure 4.1.

Viewing from the slit plane

A viewer located at the slit plane of the stereogram sees an image much like one of the five shown and, with the help of stereoscopic disparity between two views, correctly interprets the different apparent sizes of the objects as due to differences in the depth of the squares. The correspondence between the images that the camera captures and the image that the viewer sees from a particular position is especially simple when the viewer is at the slit plane: at any one time, the viewer sees an image of a single projection screen captured entirely from one camera viewpoint, as illustrated in Figure 4.2. If the viewer moves backwards, slightly away from the slit plane, this direct correspondence no longer holds. A single slit is no longer capable of filling the entire field of view of the eye. To the viewer, the left part of the projection screen will appear through one slit, the center part through another, and the right part through still another. The viewer sees, through three different slits, parts of three different projection screens captured from three different camera positions. Conversely, no projected image captured by any one camera is visible in its entirety from any viewer position.

Viewing from far away

The visual effect of moving the viewer away from the slit plane is greatest, and hence the easiest to understand, when the viewer is positioned very far from the stereogram and centered in front of it. From this distant vantage point, assuming the holographic plate and

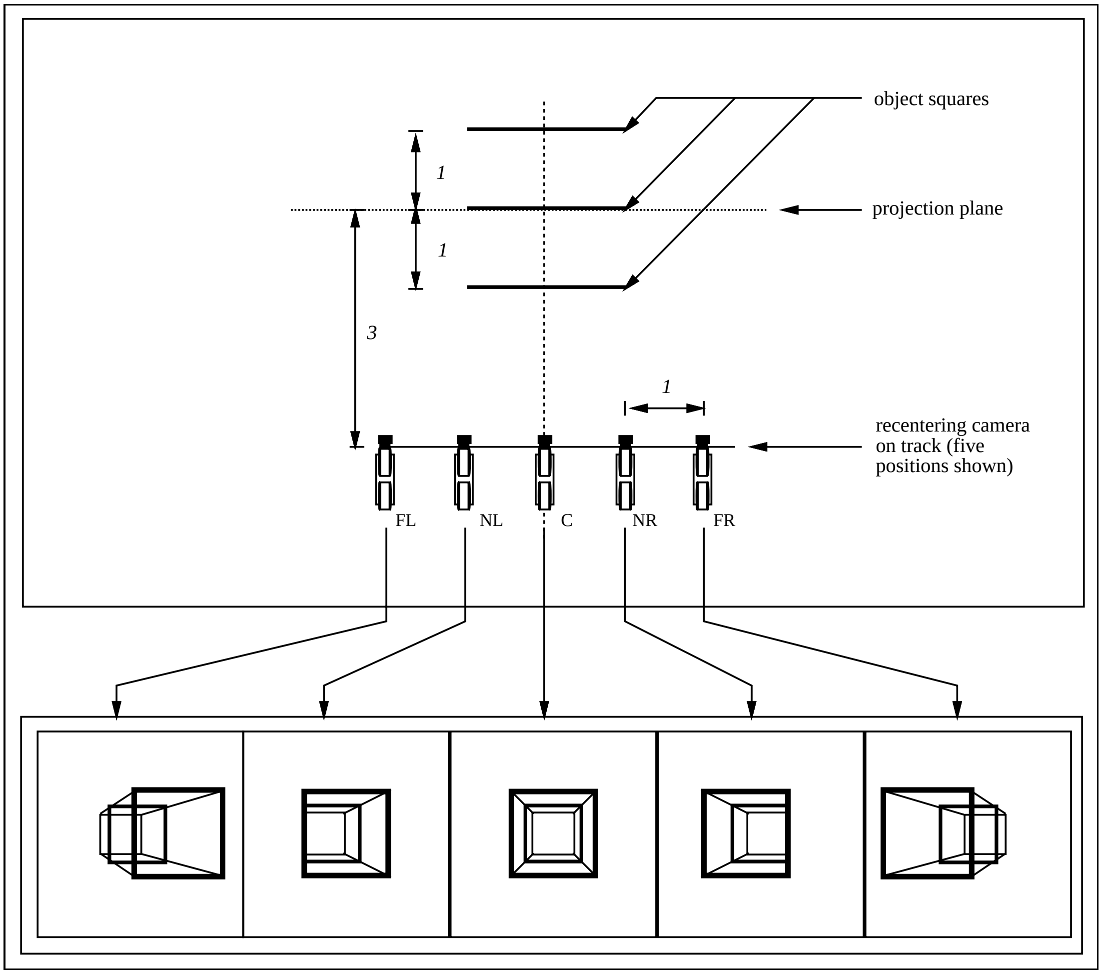

Figure 4.1: The camera geometry used to capture the stereogram is shown at the top. The five selected camera locations capture the images shown below. Different parts of these camera views will be seen at different view distances, forming the composite image that the viewer sees at a particular position.

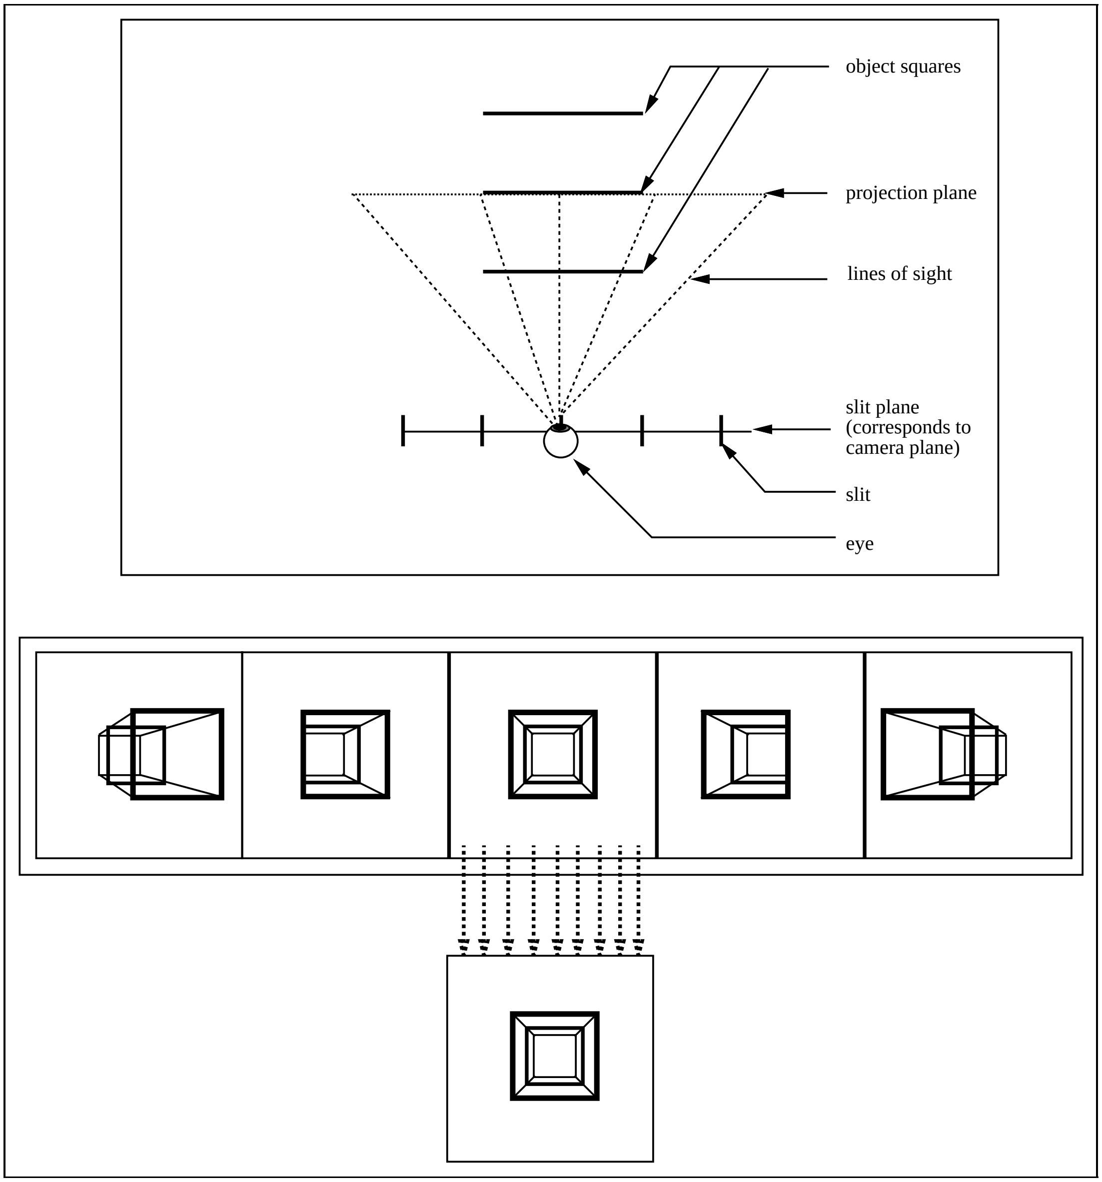

Figure 4.2: When the viewer is located at the slit plane of the stereogram, only one camera view is visible. The presented image, shown at the bottom, is identical to one of the images of the sequence.

the projection screen are the same size, the viewer will see the leftmost slit-width part of the screen through the leftmost slit, the next slice of the screen through the second slit, and so on over to the rightmost slit, which shows the righthand sliver of the projection screen. In other words, the part of the projection screen that is visible through any slit is the part directly in front of that slit; the viewer's great distance implies that the lines of sight passing through the slits are parallel to each other. To figure out what the viewer sees from the distant viewpoint, the pieces of the different views must be assembled. In this example, a limited construction will be done with the five camera views shown; this partial view will be used to guess the appearance of the complete image. Figure 4.3 shows the viewing geometry and the image as seen from the distant viewpoint.

The narrow slice of the projection screen that the viewer sees through the middle slit is at the center of the screen. This slice is the only part of the image that the viewer sees that looks the same as it did at the original viewer depth. In fact, the viewer saw this center particular slice, along with the rest of the projection screen, from the slit plane. For slits *NL* and *NR*, recall that the camera's position on the track was chosen so that the edges of the three squares would line up with each other right in front of the camera. A viewer looking through those slits, then, sees the edges lined up on the left side of the the squares through *NL*, and lined up on the right side of the squares through *NR*. Finally, no part of the squares is seen through slits *FL* and *FR*; there are no objects at the left side of the leftmost camera view or the right side of the right most view because the front square has shifted to the other side of the image from its original position, and the rear square has not shifted far enough. The image seen through these slits, and through the rest of the stereogram, is shown at the bottom of Figure 4.3.

If a viewer actually were to have moved from the slit plane of a stereogram to a position far away, the image would gradually change from its original undistorted appearance to its final distorted one. More and more disparate views would become included in the viewer's image as the viewer receded, and the pattern of the rays from the viewer to the projection

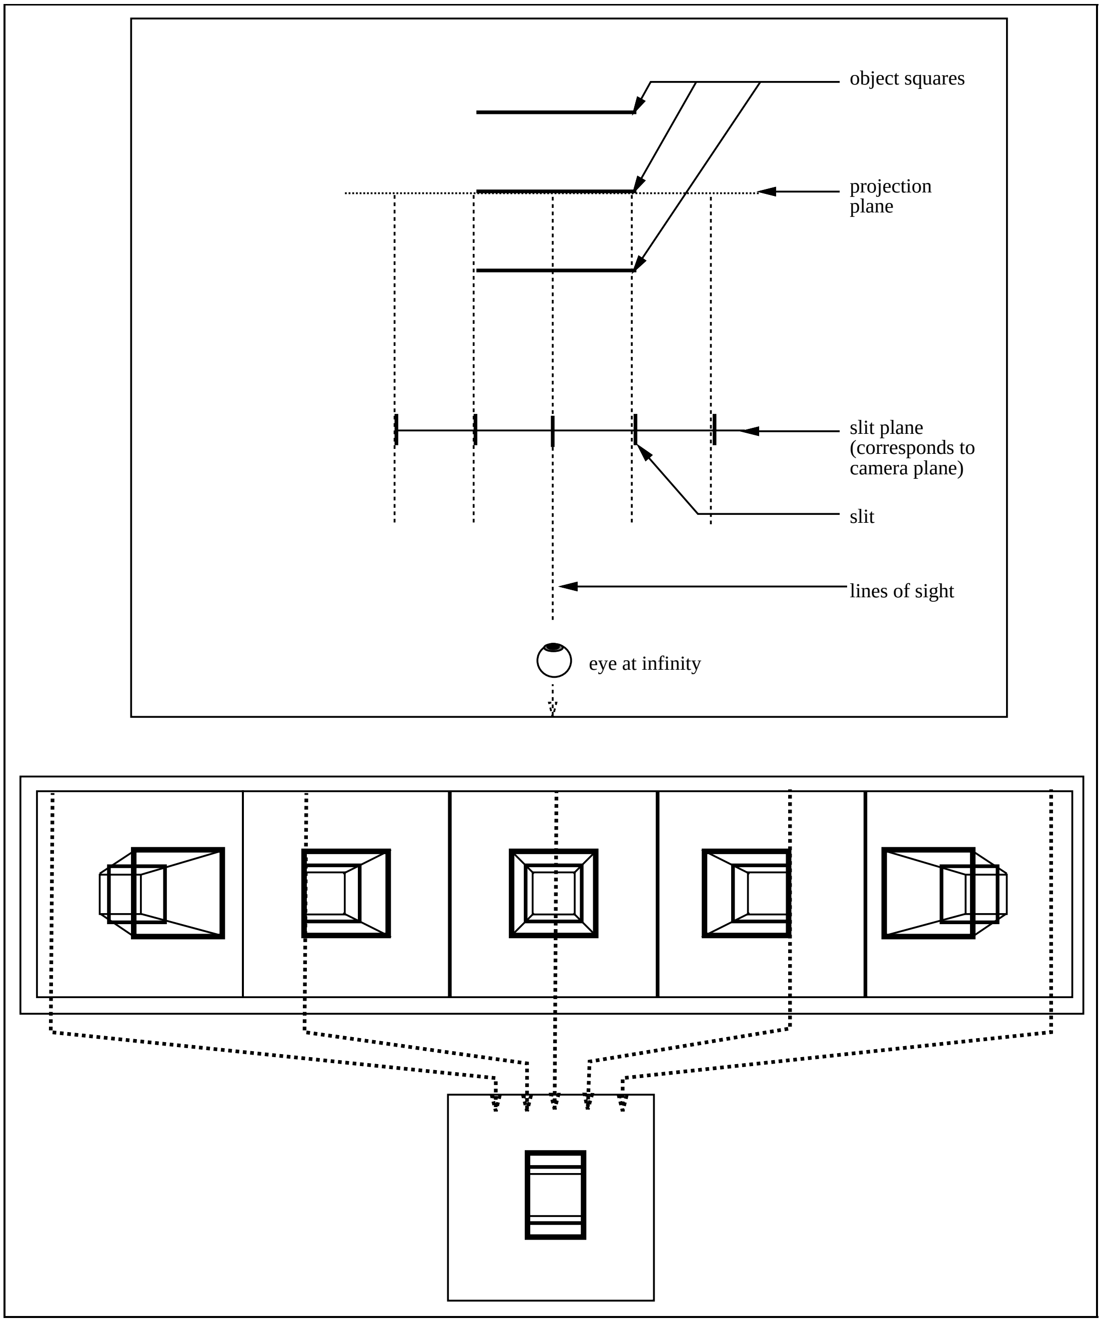

Figure 4.3: When the stereogram is viewed from an infinite view distance, the viewer sees the parts of the projection screen located directly in front of each camera during capture. The result is that the three squares have the same horizontal extent: the front square has shrunk and the rear square has stretched horizontally.

screen would change from fan shaped to parallel. Each ray gradually sweeps across the projection screen for the slit into which it falls. Just as for a real object, small changes in viewer position produce correspondingly small changes in image appearance. Additionally, objects located on or close to the projection plane such as the middle square are immune to distortion, because their size remains constant independent of viewer position (as shown in the diagram).

Viewing from in front of the slit plane

Instead of moving away from the hologram, the viewer can instead pass through the slit plane closer to the projection plane (assuming the stereogram has been transfered). The viewer's location is shown in Figure 4.4. In this case, analysis is done by tracing the path of the light from the projection screen to each slit, and seeing which rays are intersected by a viewer at a particular location. For example, a viewer centered in front of the hologram and located half way between the slit plane and the first square will see the same part of the projection screen through the center slit as was visible from the infinite viewer location. Once again, then, the center part of the viewer's image is unchanged. The viewer's eye also falls upon a line between the projection screen, an edge of the front square, and either slit *NL* or *NR*. So the part of the viewer's image that forms the right side of the front square comes from perspective *NL*, and the left side of the front square originates from *NR*. Again, no part of the object is visible through slits *FL* or *FR*. In the resulting image, shown in Figure 4.4, the front square appears stretched horizontally compared to the undistorted view, while the back square is slightly compressed. Note that the scale of the distortion is the opposite of that seen from far away.

From this viewer position, if the viewer moves slightly closer to the projection screen, the eye intercepts rays passing from the projection screen, through the edges of the front square, and through the far slits *FL* and *FR*. So, while the middle part of the image is again

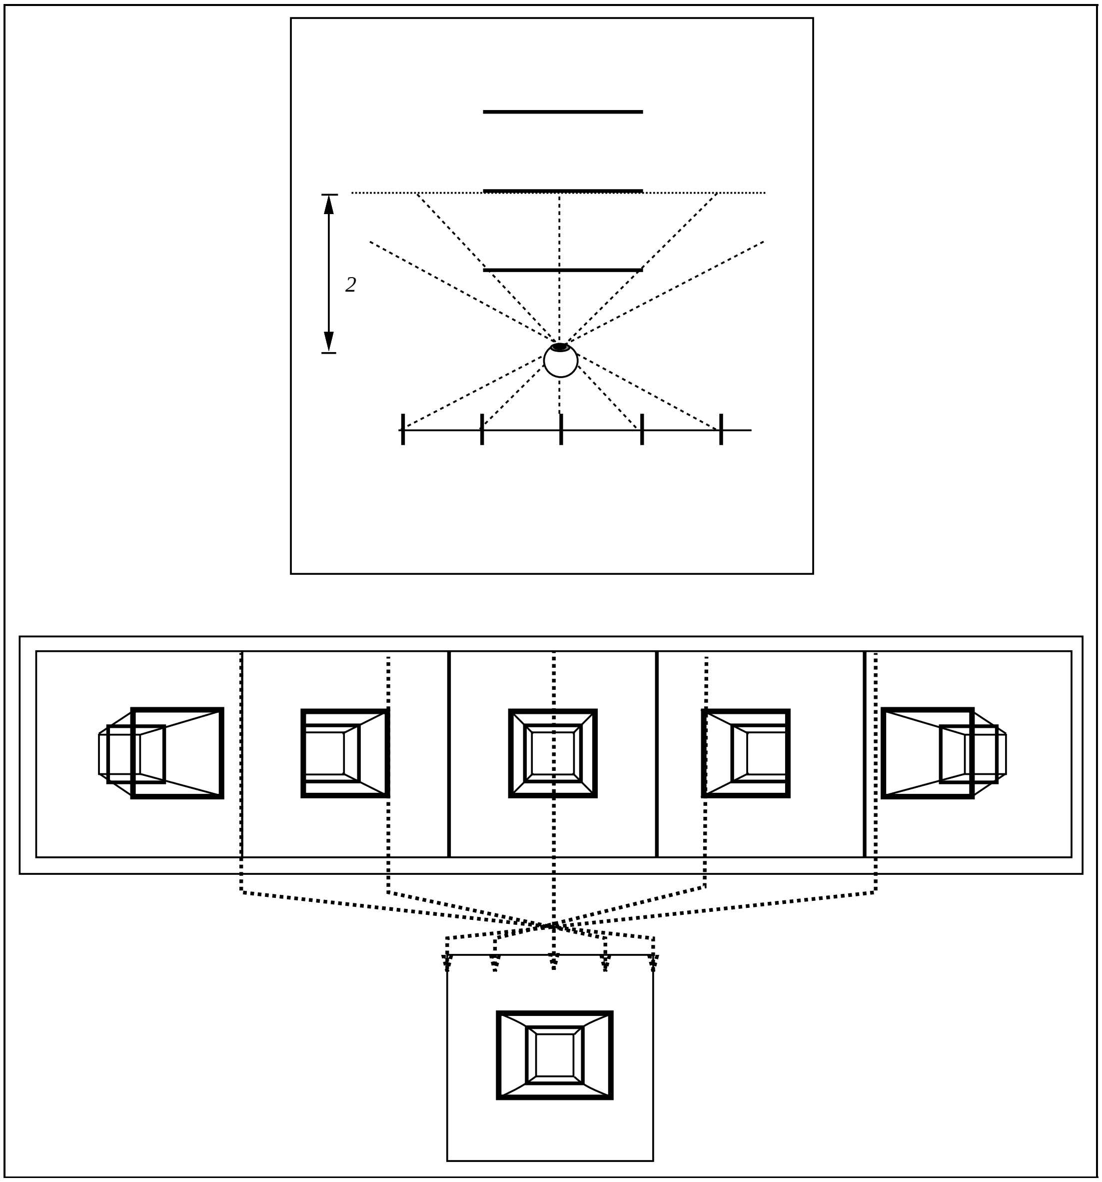

Figure 4.4: When the viewer moves inside the view plane, left camera views control the appearance of the right side of the object. Here, the eye falls on lines connecting the edges of the front square with the two intermediate slits. In the resulting image, the front square widens horizontally, while the back square shrinks. This view position requires that the slit hologram be a projected image.

unchanged, the edges of the front square in the final image are located as seen from the extreme camera viewpoints. The resulting image of the front square is even more stretched horizontally, while the back square is somewhat more compressed. A diagram of this geometry is shown in Figure 4.5.

The stereogram as an HPO display

The trend of the image distortion in the stereogram as seen from different distances may now be understood. In the horizontal direction, the stereogram accurately approximates the behavior of a continuous parallax system. Horizontally, objects change in perspective in a natural way as if they were truly three-dimensional. Vertical detail is not three-dimensional, however; all vertical information in the image is photographically recorded and “stuck” to the projection plane. No change in viewer position can alter the relative positions of vertical details. Instead, vertical perspective remains fixed to that seen by the camera when the images are recorded. The stereogram’s discrete structure is irrelevant to this type of distortion: in this sense, a stereogram is prone to the same distortions as a continuous parallax HPO display. Were the stereogram instead a full parallax array, both the horizontal and vertical information in a scene would change naturally as the viewer moved towards or away from the stereogram.

Prerequisites for undistorted stereograms

What guarantees that the horizontal perspective changes of the stereogram will match those of the object is the correspondence between the location of the camera plane of best focus and the location of the slit when the stereogram is viewed. A slit of the stereogram is the place where all the rays from that slit’s projection screen converge, cross, and then diverge away from the hologram. This horizontal focus must be accurately modeled, using

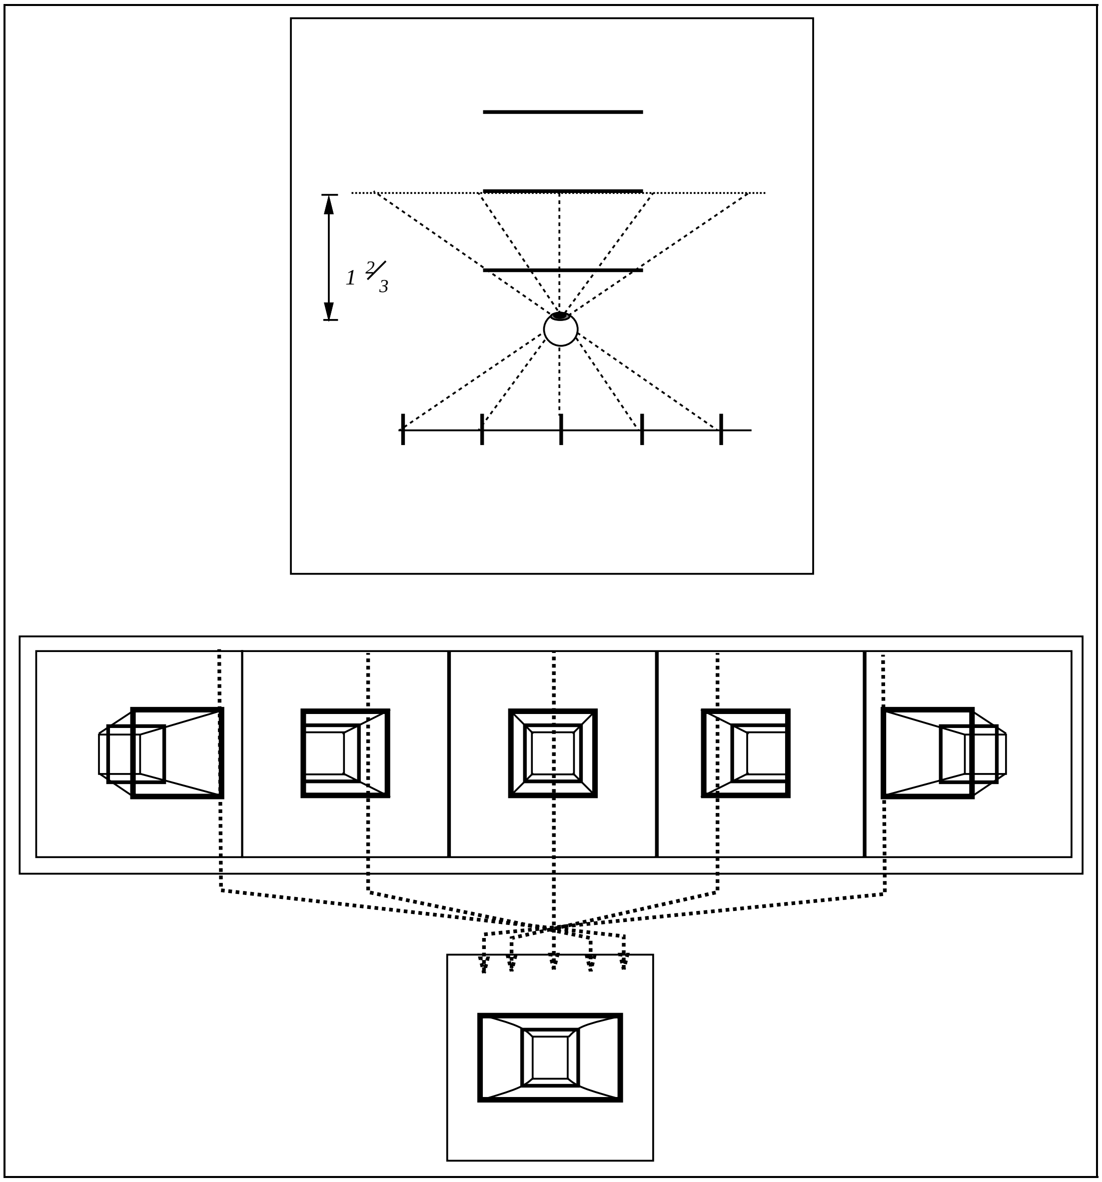

Figure 4.5: A closer view distance produces even more extreme distortions. The eye now receives information about the edges of the front square from the extreme slit on the opposite side from the edge. Note the non-linearity of the image distortion, as shown by the curve connecting the the corners of the squares.

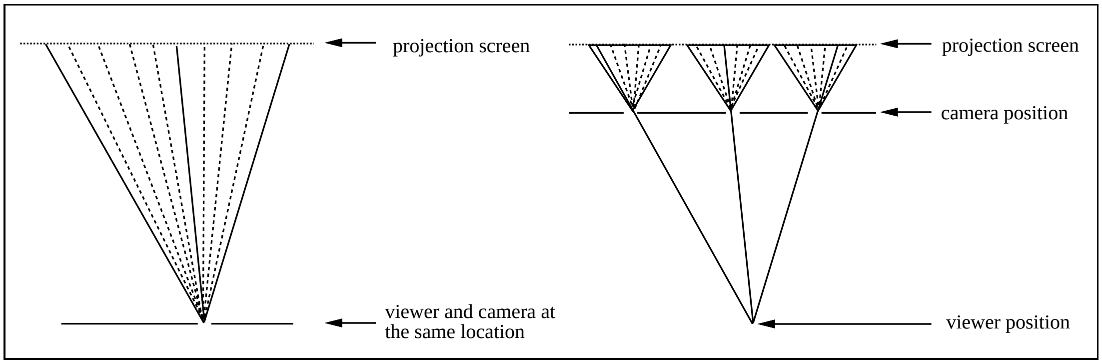

Figure 4.6: As long as the position of the image capture camera and the slits coincide, a viewer will see the same image rays coming from the stereogram independent of the location of the slits.

a physical or synthetic camera, when the image is captured. If the horizontal focus is correctly modeled, a viewer will see the same rays independent of the slit position, as shown in Figure 4.6. To perform this modeling accurately, the camera's geometry must be equivalent to the slit's geometry; otherwise, intensity information about the ray directions that the camera captures will be used to modulate the intensity of rays traveling to the viewer at slightly different directions, resulting in compression or expansion of all or parts of the object in depth. In short, the camera and the slits must coincide.

The only viewer location where the viewer can see an undistorted image while looking at a simple HPO stereogram is at the plane where the camera captured the scene. This is the only viewer location where the stereogram's variable horizontal perspective and its fixed vertical perspective are matched. So, while the camera and slit planes must coincide to avoid depth distortions, the viewer and camera planes must coincide to match vertical and horizontal perspectives. Therefore, a stereogram captured with a conventional camera must be viewed at the distance where the slits appear to be if the image is to be undistorted.

The nature of the distortion introduced when the viewer and the slits are not at the same depth, however, suggests a way of modifying the camera's optics to capture a sequence of images that, while appearing distorted when examined on a frame by frame basis, will produce an undistorted stereographic image when viewed from a fixed distance other than the slit plane. This modification involves separating the positions of the vertical and horizontal foci of the camera's lens to match the differences between the camera's slit location and the eventual viewer location. The stereogram requires the position of the vertical focus of the camera match that of final viewer. It also demands that the horizontal focus of the camera lie at the same location as the slits of the stereogram. Both of these requirements can be satisfied by the same image capture camera if that camera's lens is anamorphic. We call stereograms that use image distortion techniques to change the slit or viewer position *advanced stereograms*, or ULTRAGRAMS, because of the added power and image configurability these methods can provide.

Anamorphic cameras

An anamorphic physical camera can be created with a standard spherical-surfaced lens coupled with a cylindrical optic; alternatively, two crossed cylindrical lenses can be used. The horizontal and vertical details of a scene are displayed completely independently by a stereogram; this independence can only true if, as in a conventional stereogram, the camera faces straight forward during the capture of all views. The independence of the two axes is especially important if the image capture camera is not physical but instead synthetic. For a computer graphics camera, horizontal and vertical independence means that perspective calculations can be altered in one direction without affecting the other, which greatly simplifies the process. The configurability of synthetic anamorphic cameras, compared to the expense and complexity of the corresponding physical optics, has resulted in their exclusive use for the work described here.

The availability of an anamorphic camera frees the stereogram design of any constraints on slit plane location with respect to the viewer or to the object. For example, the object can pass through the plane of the stereogram, or can appear completely out in front of it. If the plane of the slits of the stereogram is inside the object, the camera geometry is somewhat unintuitive; it dictates that the horizontal focus of the capturing camera must also be inside the object. The image captured, however, is not of the inside or the back side of the object; the location of the focus merely decides where the rays that the camera captures must cross. The viewer's position still determines what objects or parts of objects are visible.

As the images of the distorted views of the square test pattern show, however, a nonlinear distortion is produced by a change in viewer position. The inverse distortion required to produce an undistorted image, then, also must be nonlinear. For physical cameras and ray tracing computer graphics cameras, this nonlinearity can be produced with a cylindrical optic or its computational equivalent. However, most scanline-based computer graphics renderers, though, rely on the linear properties of image transformations, so that the effect of a cylindrical lens is difficult to model. This difficulty is especially acute if the object passes through the slit, or horizontal camera plane, a situation that often occurs in one-step holography. If the object is intersected by the slit plane, the part of the object on the projection screen side of the slit is greatly magnified, the part on the opposite side of the slit is magnified and reversed left to right, and the part that falls on the slit plane has infinite extent, causing a singularity in standard computer graphics perspective calculations. Instead of performing the extensive modifications to a standard computer graphics renderer necessary to deal directly with these mathematical problems, a post processing method of synthesizing the desired perspective views from other, more easily obtainable views was used.

Approaches to synthetic anamorphic cameras

The essential property of the final projected images is that the rays that they describe form a fan that spans the projection screen and crosses at the slit. All the techniques presented here for synthesizing computer generated images presented here involve the following general technique: calculate a sequence of images containing the intensity of all rays needed for all the slit images, but where the information for any one slit is distributed in parts of many images of the sequence. The parts of the input views are then rearranged to form a new sequence of images, in which each image is one required to expose a slit. The camera geometry used to capture the input sequence of images can be chosen to simplify the image rearrangement process. The vertical focus of the camera always remains at the location that corresponds to the final viewer position.

The infinite viewpoint camera

One of the simplest approaches to this technique moves the camera's horizontal perspective away from the slit plane and positions it at infinity, forming a *horizontally orthographic projection* of the scene. From this position, the horizontally-displaced rays of the camera travel parallel to each other. The number of horizontal rays, which corresponds to the number of horizontal pixels in the computer graphic image that the camera generates, is matched to the number of slits of the stereogram. The camera's horizontal position determines the angle at which the camera's rays strike the plane of the hologram. Only those rays that pass through the stereogram are of interest; therefore, the camera's view is always recentered to include the horizontal window of space defined by the holographic plate. If the camera moves horizontally through space, a collection of ray fans is swept out, each fan centered on the location of a slit. Each fan consists of the rays needed to expose its particular slit. Figure 4.7 shows how rays from three different camera positions cross to trace rays through three slits.

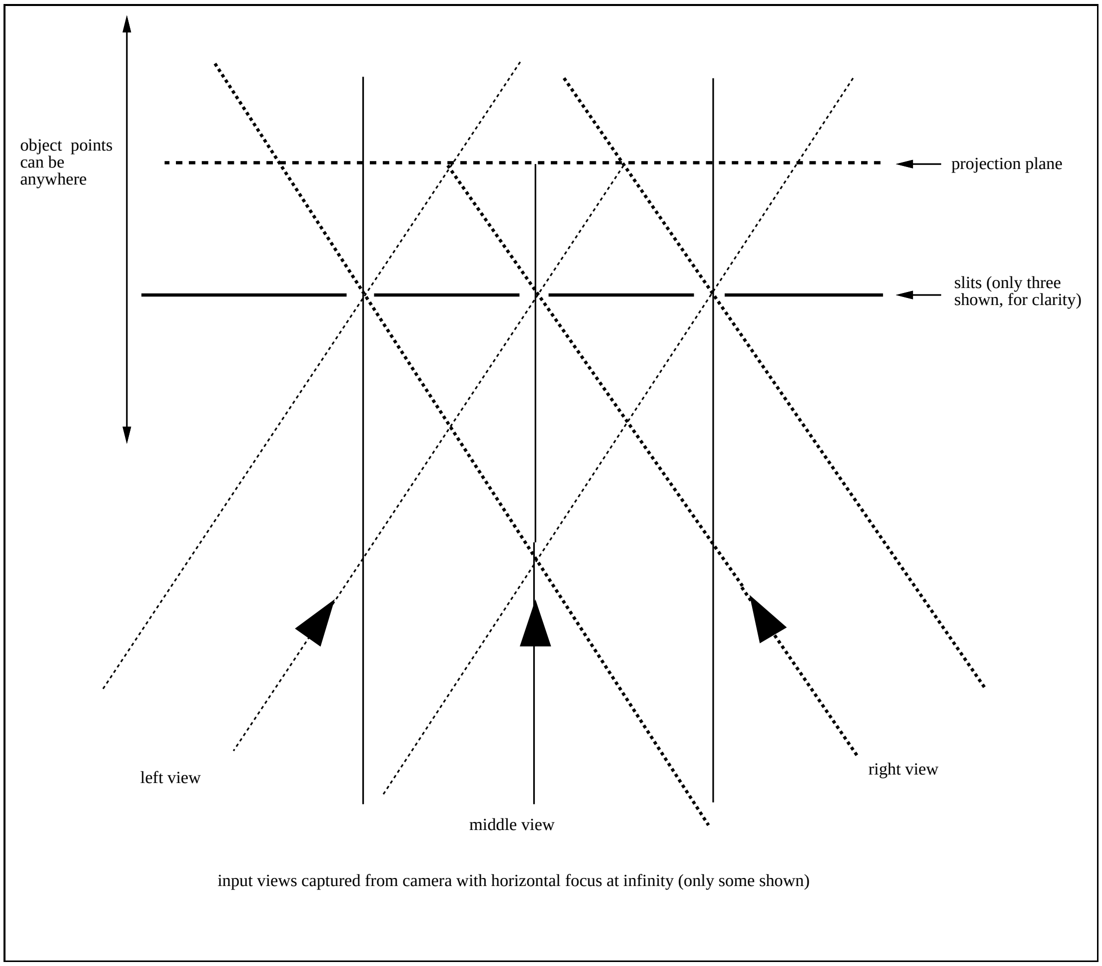

Figure 4.7: The rays that three cameras at infinity trace through a scene cross at each slit when the cameras recenter on the stereogram plane. This collection of rays can be formed into a view that appears to be from a camera with a horizontal focus on the slit plane.

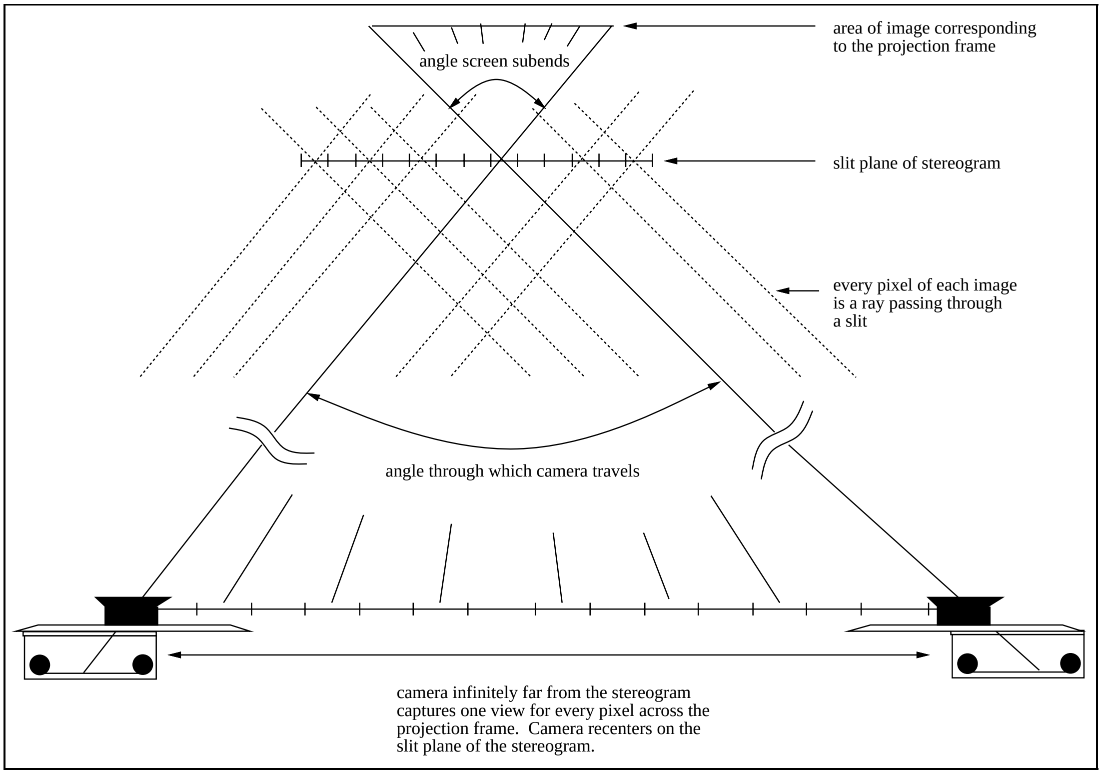

Figure 4.8: The image to be projected onto the projection screen for a slit is made up of one ray from each camera position. The angle that the projection frame subtends with respect to the slit must match the angle through which the camera travels.

To further simplify the example of this technique, first assume that the stereogram to be made will be a “simple” or “non-recentered” camera stereogram, so the projection frame used to expose each slit is centered in front of that slit. If this is true, the shape of each fan of rays needed to expose each slit is exactly the same for all the slits. To capture the sequence of input images, then, the camera moves along a horizontal track of a length proportional to the width of the projection frame, capturing one view for every column of pixels in the projected image. Figure 4.8 diagrams the camera geometry needed, using one slit as an example. In the final sequence, The collection of first columns from every camera view is assembled to form the first slit, the second column for the second slit, and so on. This reorganization is accomplished by forming a pixel volume by stacking the input views with the leftmost camera view in back, the next view in front of it, and so on, with the rightmost view in front. The volume is then rotated ninety degrees around its vertical axis and resampled into the output frames. This process is summarized in Figure 4.9. Pseudocode that implements this volume-rotation based predistortion approach is shown in Figure 4.10.

The advantage of placing the synthetic camera at infinity is that all the objects in the scene are very far from the camera plane and are thus unlikely to intersect it. A disadvantage to this predistortion scheme is that instead of an arbitrarily high resolution at the projection plane as allowed by bandlimiting constraints, the maximum resolution available is the pixel width of the input views, which necessarily corresponds to the slit width. Computer images, however, seldom have arbitrarily high resolution, so that if the smallest pixels capable of being projected onto the projection screen are about the same width as a stereogram slit, nothing is lost by using this technique. An added bonus of this resolution limitation is that the images produced are certain to be bandlimited at a depth at least up to the plane of the slits because the slit plane is sampled at a rate of one sample per slit.

Rendering recentered views

A straightforward extension of the infinite camera predistortion technique can be applied to make recentering-camera stereograms: the camera at infinity just has to move further horizontally to capture an image of all the rays that pass through a slit and strike the now stationary projection frame. Much calculation is wasted, however, because at extreme camera positions only a small fraction of the computed image is used. Figure 4.11 shows the worst case of this waste, when the camera is imaging the last column of the first slit and only one column of the entire image is used. Another technique is more efficient for computing recentered stereogram views.

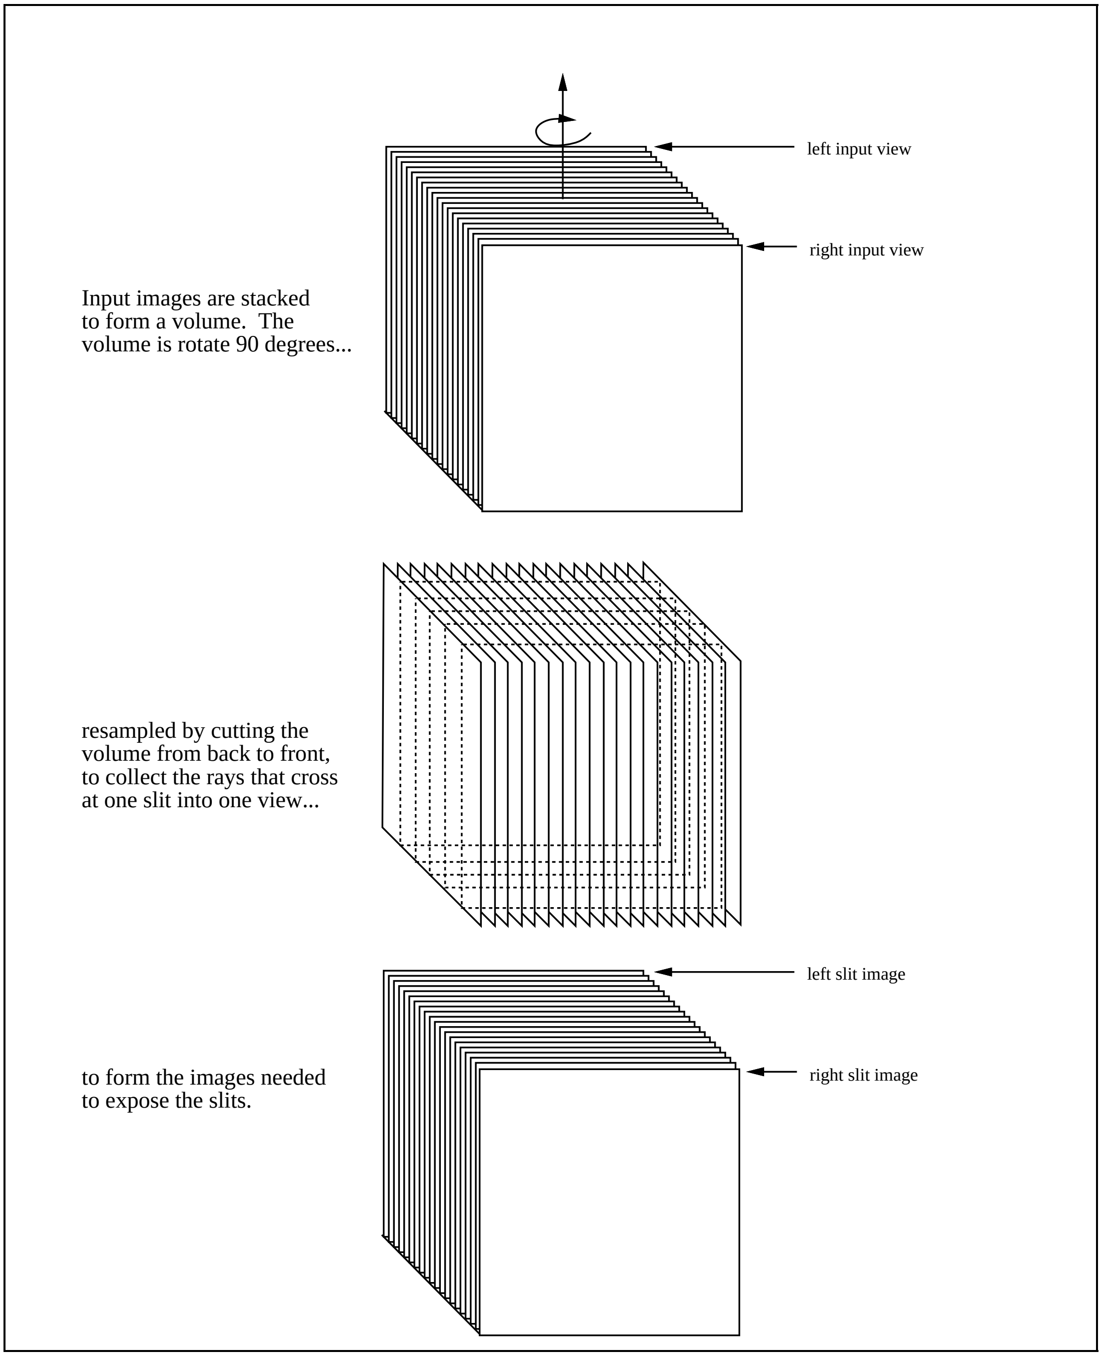

Figure 4.9: Input views from a camera at infinity can be resampled in computer memory to form predistorted views for exposing the slits of a stereogram.

```
program rotate_volume

declare parameters
    n_screen_cols (horizontal resolution of each slit's image)
    n_slits (number of slits to be exposed)
    n_rows (vertical resolution of each slit's image)

declare variables
    slice[n_screen_cols][n_slits]
        (each row of array holds a row of each input view)
    t_slice[n_slits][n_screen_cols]

for row ← 0 upto n_rows { (for every image row)
    for input_view ← 0 upto n_screen_cols {
        (input a row of pixels and put it in a row of slice)
        slice[input_view][] ← read_row(row, input_view, n_slits);
    }
    t_slice ← rotate_slice(slice);
    for output_view ← 0 upto n_slits {
        (output each newly formed row to the image for each slit)
        write_row(row, output_view, n_screen_cols, t_slice[output_view][]);
    }
}

function read_row(row, view, number)
(read number pixels from row row of image view)

function write_row(row, view, number, data)
(write number pixels to row row of image view, using data from data)

function rotate_slice(array)
(rotate the slice array 90 degrees around its center, the same as
        transposing the array and then flipping each row)
```

Figure 4.10: Pseudocode for the volume rotation used in the infinite anamorphic camera method.

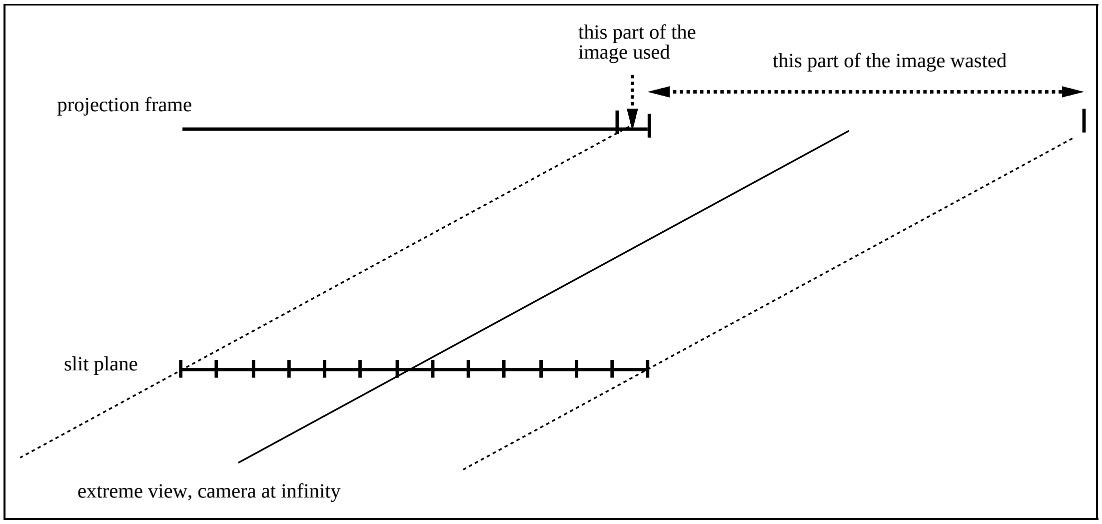

Figure 4.11: The extreme camera view needed to capture the last column of the first slit of a recentering camera stereogram with a camera at infinity wastes the computation of all but that slit because no other part of the image is used in any view.

Perspective slicing

A viewer standing behind a stereogram sees segments of many projection screens abutted together, each projection screen segment seen through one slit. Each slit subtends some width on the projection plane, a width given by Equation 4.1, where $W_{slit-proj}$ is the width subtended, D_{proj} is the viewer's distance to the projection plane, and D_{slit} is the distance from the viewer to the slit plane.

$$W_{slit-proj} = \frac{D_{proj}}{D_{slit}} W_{slit} \tag{4.1}$$

If the viewer moves horizontally, the segment of the projection screen that was visible through each slit will move out of view at a rate proportional to the viewer's velocity. After a certain fixed distance, the old segments of the projection screen will have completely moved out of the slits, and non-overlapping new ones will be visible. One way to generate the views for the slits of the stereogram is to use a horizontally moving recentering camera

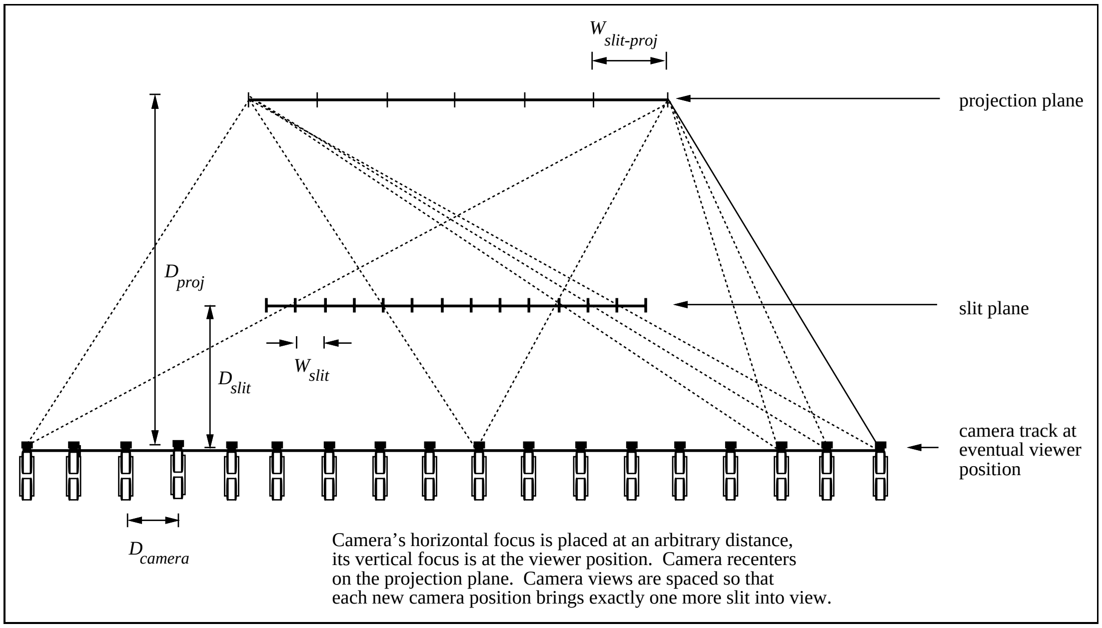

Figure 4.12: The camera geometry used for the perspective slicing predistortion method.

that takes a picture of the projection plane every time the camera has moved this fixed distance, as given in Equation 4.2.

$$D_{camera} = \frac{D_{proj}}{D_{proj} - D_{slit}} W_{slit} \tag{4.2}$$

The length of the camera track, W_{track}, is the given by Equation 4.3.

$$W_{track} = \frac{W_{proj} D_{slit} + (n_{slits} - 2) W_{slit} D_{proj}}{(D_{proj} - D_{slit})} \tag{4.3}$$

The number of camera views is given by Equation 4.4.

$$n_{views} = n_{slits} + \left\lceil \frac{W_{proj}}{W_{slit-proj}} \right\rceil - 1 \tag{4.4}$$

Figure 4.12 shows the capture geometry. The image for one slit is formed by piecing together all the segments of the projection plane that were subtended by that slit's boundary

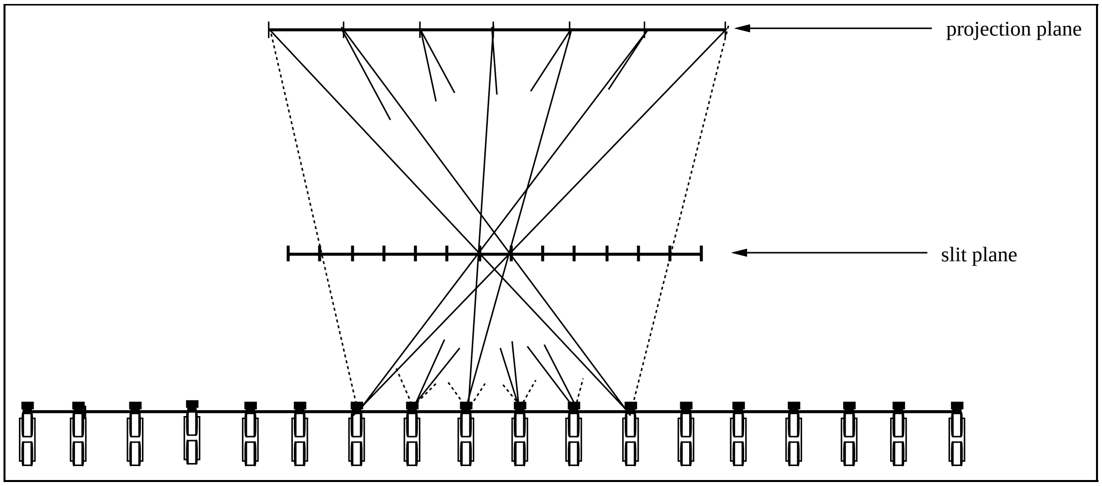

Figure 4.13: In perspective slicing, the projection screen for each slit is formed by collecting the pieces of each camera view that pass through that slit.

in the sequence of input views. In other words, the camera perspective required for a single slit is approximated from slices of many images captured from a camera at an arbitrary depth. These slices may span many pixels. Figure 4.13 shows how one slit's image is composed of slices of many different views. Pseudocode that implements perspective slicing is shown in Figure 4.14. Nothing assures that a slice boundary will fall on a pixel boundary when the input images are distributed, so care must be taken to carefully combine the edges of overlapping slices.

The big advantage of this method is it permits objects near the projection plane to be rendered at a high resolution. This property is contrast to the infinite camera technique, which limits the resolution of the entire image field to the width of the stereogram's slits. As mentioned above, perspective slicing also can reduce the amount of computation, compared to infinite camera distortion, that is spent rendering parts of a scene that will never be visible. This savings occurs because the input camera always renders only the area of the recentered projection frame from all camera locations. Unused information is still rendered to produce the output frames for the slits on the edges of the stereogram, but the percentage of unused

```
program perspective_slice

declare parameters
    w_slit      (width of slit)
    n_slits     (number of slits)
    d_proj      (distance from camera to proj. screen)
    d_slit      (distance from camera to slit plane)
    pw_proj     (pixel width of proj. screen)

w_slit_proj ← w_slit·(d_proj/d_slit);     (width of slit proj. on proj. plane)
n_slit_projs ← w_proj/w_slit_proj;     (# slits subtending one proj. screen)
pw_slit_proj ← pw_proj/n_slit_projs;     (pixel width of slit proj.)
n_input_views ← n_slits + n_slit_projs - 1;
                        (number of input views needed to cover all proj. screens)

for row ← 0 upto n_rows {
    (input a row from all input views)
    for input_view ← 0 upto n_input_views {
        in_data[input_view][] ← read_row(row, input_view, pw_proj);
    }
    start_view ← 0;
    for output_view ← 0 upto n_slits {
        input_view ← start_view;
        for slit_proj ← 0 upto n_slit_projs {
            offset ← slit_proj · pw_slit_proj;
            copy(in_data[input_view][], out_row[], offset, pw_slit_proj);
            input_view ← input_view + 1;
        }
        write_row(row, output_view, pw_proj, out_row[]);
        start_view ← start_view + 1;
    }
}

function copy(src, dst, offset, number)
    (copies number pixels from src to dst starting at offset from the
            beginning of src and dst)
```

Figure 4.14: Pseudocode for the perspective slice predistortion method.

data decreases significantly as the number of stereogram slits increases. This fact saves considerable computation when small, high resolution images with a large number of small slits are produced. Finally, this method permits a non-anamorphic camera to produce the input views for a stereogram that requires predistortion, allowing many existing image creation methods to be incorporated into advanced stereographic displays.

The disadvantages of this method arise from the discrete nature of the slices of the images. When the slits of the stereogram are very narrow, the pixels on the edges of the slits may often be cut by a slice boundary, leading to excessive filtering and loss of spatial resolution. So, this method is not well-suited to cases in which the slit and projection frame pixel size are approximately the same. Furthermore, much as a stereogram approximates a wavefront by using pieces of other, different wavefronts, the perspective slicing method approximates one perspective by using pieces of significantly different ones. While the net difference in the desired perspective and the approximation is small, the local, intra-slit differences can be large. The perspective within a slice of an image has a horizontal perspective corresponding to the input camera's location, not the slit plane's. This difference makes a perspective-sliced image vulnerable to aliasing artifacts, which are similar to but distinct from the artifacts produced by the stereogram's discrete aperture structure. These artifacts can be reduced by matching as closely as possible the locations of the slit plane and the input camera. A slit plane on the other side of the projection plane from the viewer, which can happen in direct-illuminated two-step holograms, cannot easily be matched with perspectives from a camera whose viewpoints are easy to calculate. In this case, a camera at infinity should be used. Whatever the position of the input camera, the exact consequences of the sampling artifacts caused by perspective slicing have not yet been quantitatively explored.

Predistortion and two-step stereograms

As parts of this discussion have mentioned, two-step holography offers an increase in image flexibility and variety when compared to the simple stereogram example of this text. Two-step holograms let the viewer be positioned at or in front of the slit plane. Such a stereogram uses a holographic transfer step in which the slit hologram is illuminated with a phase conjugate illumination source that projects an image of the projection screen out into space. True phase conjugation is difficult to achieve, though, and improper illumination during either the transfer or the viewing step can produce an image of the projection screen where the horizontal and vertical foci of the screen do not fall at the same depth; that is, the transfer process can introduce astigmatism. Shifts in wavelength between the light sources used in stereogram mastering, transfering, and viewing can also result in astigmatism. Similarly, an anamorphic optical element may take the place of the projection screen in a holographic setup. Predistortion techniques can compensate for the major effects of an anamorphic projection screen by invoking a generalization of the rules for normal predistortion.

Guidelines for distortion-free stereograms

The constraints for producing undistorted horizontal parallax only stereograms can be summarized in the following list. These points better define exactly how the capture, recording, and viewing geometries must correspond to each other.

- The horizontal and vertical positions of the stereogram's slits and the projection screen should be found relative to the final viewing geometry, after all optical distortions have been accounted for.

- The location of the horizontal focus of each slit must correspond to the location of the camera that captured that slit's projection screen image.

- The area of space imaged by the image capture camera must match the area of the projection frame with respect to the slit being exposed.

- If the projection screen (or its image) is astigmatic, the position of its horizontal focus should be used in determining the geometrical relationship between the slit and the projection screen.

- Aliasing artifacts are minimized by positioning the scene so as to straddle the plane of the projection screen's horizontal focus.

- The vertical focus of the capture camera must lie at the final viewer's position with respect to the vertical focus of the projection screen.

- If the final hologram is to be viewed in white light, the projection screen's plane of vertical focus should coincide with the plane of the hologram in order to minimize chromatic and source-size blurring.

- HPO stereograms will always suffer from some distortion as the viewer moves away from the intended view distance, but this distortion can be minimized by positioning the object near the plane of vertical focus of the projection screen.

Advantages of ULTRAGRAMS

By following these guidelines and using the necessary image predistortion techniques, stereograms with many interesting properties can be made. For example, while most two-step stereograms require error-prone phase conjugate illumination, ULTRAGRAMS can be made in which the image of the slits is behind the plane of the transfer plate, allowing

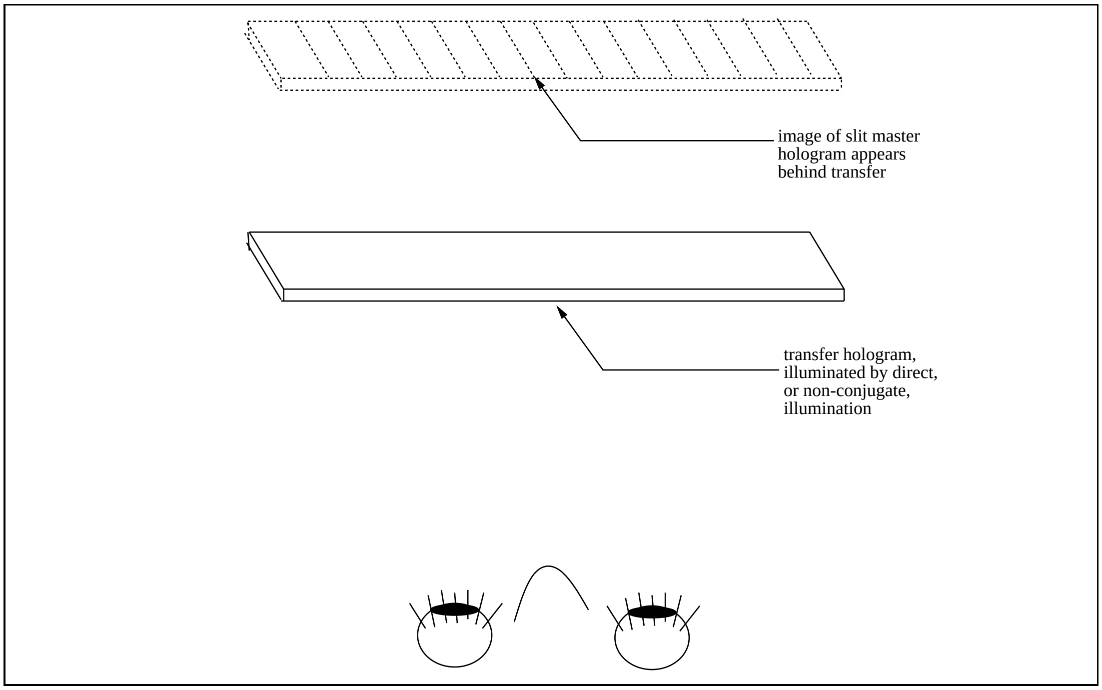

Figure 4.15: Predistortion allows stereograms to be lit with direct illumination, so the slit master floats behind the transfer hologram.

the use of a diverging illumination source that closely matches the reference source. The viewing geometry for this type of stereogram is shown in Figure 4.15.

The ability to position the plane of the stereogram anywhere in space not only increases the flexibility of the medium, it also permits a tradeoff to be made between spatial resolution and the size of the stereogram's view zone. If a two-step's master and transfer holograms are placed close to each other, a viewer can move very far horizontally and vertically and still be able to see through at least part of the two windows defined by the plates. This increase in view zone size is shown in Figure 4.16. However, the depth of field of the stereogram, and thus the maximum spatial resolution possible at any depth plane, is reduced as the slit plane is positioned closer to the projection plane. From an information content point of view, the size of the view zone may increase as the planes get closer together, but the amount of

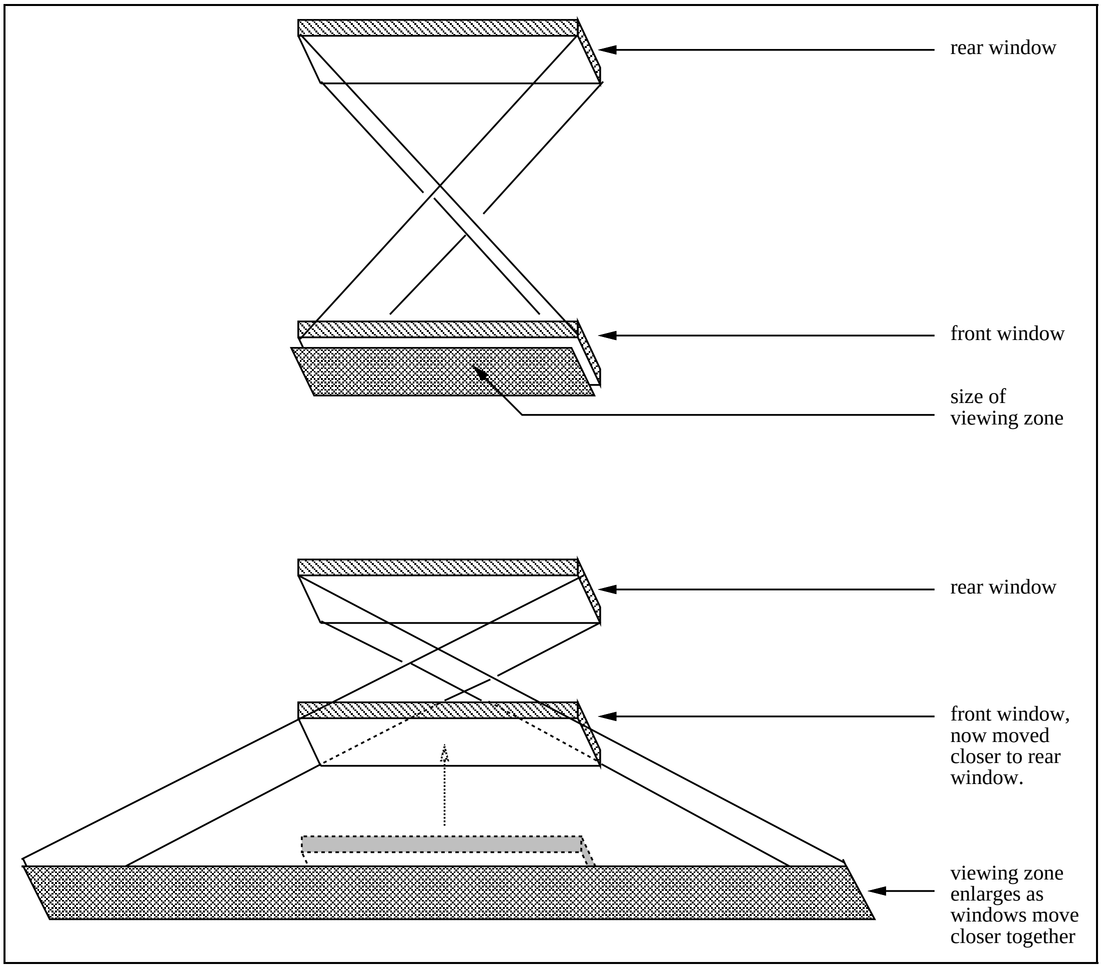

Figure 4.16: The view zone of a stereogram is defined at the region of space where some part of the information of the hologram can be seen. If the slit plane of the stereogram is at the viewer's eye, the view zone is the size of the slit master. Moving the slit plane closer to the transfer hologram widens the size of the view zone. Not shown is the similar increase in the vertical size of the view zone.

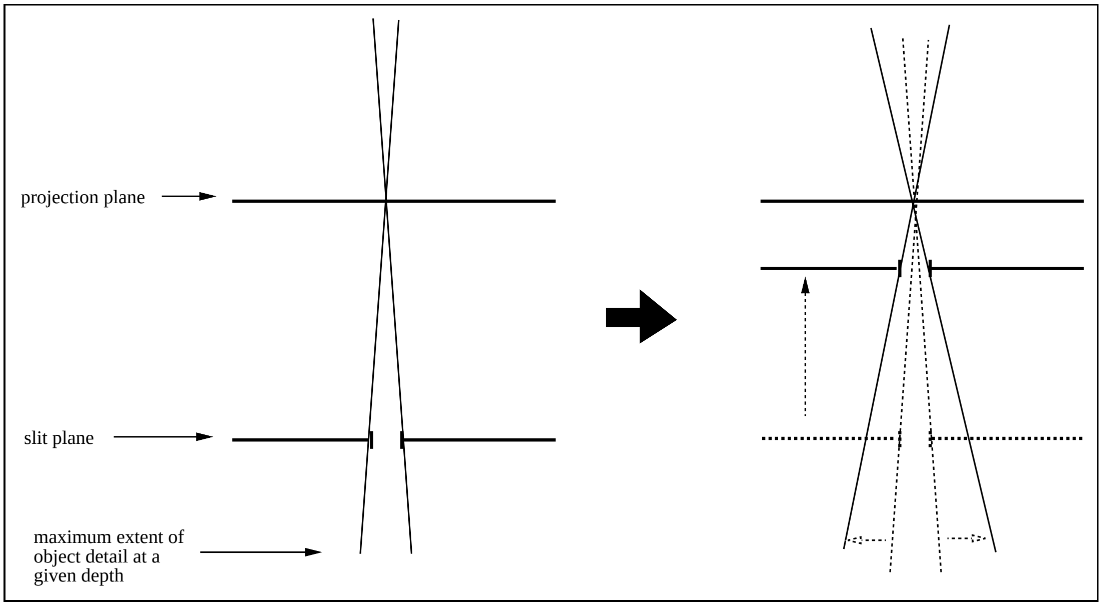

Figure 4.17: As the slit plane and the projection plane get closer together, the size of the smallest allowable detail at any a given depth gets larger. This loss of spatial resolution corresponds to an increase in view zone size.

information captured by the camera and put into the stereogram remains the same. Spatial resolution is lost by the same factor as horizontal view zone size is gained. The change in stereographic depth of field is shown in Figure 4.17.

In summary, with proper attention to details of capture and viewing geometries, stereograms may be created that have a view distance that is independent of the location of the stereogram's slit plane. Stereograms in which the viewer is not at the plane of the slits suffer distortion due to the fact the stereogram has parallax in only one direction. This distortion can be compensated for by generating a sequence of views with horizontal foci that lie on the slit plane. These images have the inverse distortion to that which the stereogram introduces, so that their combination presents an undistorted view to the viewer. The vertical perspective of the camera must match that seen from the eventual viewing zone. Several ways of producing the correct predistorted images have been presented in this chapter. A sequence of conventional, undistorted images can be used to generate a new sequence of distorted ones, so that existing image sequence creation methods can be used. Finally, the size of the stereogram's view zone and the amount of spatial detail the stereogram can store are inversely related. The next chapter will present a case-by-case description of the advanced stereograms produced to date. These stereograms provide useful examples of the wide applicability of these computer predistortion techniques.

Chapter 5

Practical Examples

The previous sections have outlined a variety of basic and advanced holo-stereographic concepts. These concepts have emerged from a year and a half of work, the goal of which was to find practical ways of improving the quality and usefulness of holographic stereograms. The material in this chapter is a history of advanced, flat format stereography at the Spatial Imaging Group (MIT Media Laboratory). The name ULTRAGRAM was coined to describe advanced stereograms that offer significant improvements in viewing quality and production technology through the use of image predistortion. The research direction of the group, and of the Media Laboratory in general, is towards improving the quality of human–computer interaction; naturally then, all of the images used to make the stereograms described here are computer generated, and all are predistorted using computational techniques. At this time, no raytracing programs able to produce distorted images directly as output are used. Such special-purpose programs would require the reimplementation of all the computer graphics techniques required to produce high quality images, a dauntingly large task. The use of standard polygonal renderers, with slight modifications, and a post-rendering predistortion step greatly widens the range of sources from which image data could come.

The types of ULTRAGRAM images made to date include a wide angle of view, two-step reflection hologram; a large scale meter-square one-step transmission hologram; and a small, high-resolution two-step reflection hologram suitable for mass production. The existence of these holograms is testimony to the fact that computer predistortion of input views can be done practically and efficiently, even when the images are composed of tens or hundreds of megabytes of data. The images and the techniques that were used to produce them are described in historical order.

Phone Test Pattern #1

Type: two-step transmission
Transfer illumination: phase conjugation, white light
Master size: 300mm wide, 300mm high
Transfer size: 300mm wide, 300mm high
Projection frame size: 300mm wide, 300mm high
Slit width: 1mm
Number of slits: 300
Master-transfer separation distance:100mm
Resolution of projected images: 480 pixels wide, 480 pixels high
Stereogram-type: simple camera
View Distance: 1 meter

Phone Test Pattern #2

same as *Phone Test Pattern #1*, except for:
Transfer illumination: direct illumination, white light

The *Phone Test Pattern* was designed not as an end product, but rather as a test of computer predistortion methods, to examine the holographic consequences of a very small master-transfer distance, and to see what viewing advantages might come from using the ULTRAGRAM format. This image was not intended to be a high quality display piece; the fact that the master and transfer plates were parallel precluded the use of an achromatic transfer, so that the final image suffers from moderate chromatic blurring.

The views for this stereogram, an image of a telephone handset crossing a grid and surrounded by wireframe cubes, were computed using a rendering library called Rendermatic, developed by students in the Computer Graphics and Animation Group at the Media Laboratory. Predistortion of the images was done using the infinite camera model.

However, instead of the camera's horizontal focus being at infinity as required, it was erroneously positioned at the view zone distance. The first version, *Phone #1*, used conventional phase conjugate illumination of the transfer hologram. Although the diverging illumination source used to light the hologram did not match the conjugate of the transfer's reference source, no attempt was made to correct for the distortions that resulted. *Phone #2* was the first stereogram to demonstrate the viability of direct illumination, in which the image of the slits appear behind the plane of the transfer plate. Test stereograms of both *Phone #1* and *Phone #2*, while somewhat distorted, were surprisingly achromatic, and exhibited sufficiently good image quality and widened view zone to encourage the production of the next image without further theoretical work. Figure 5.1 shows three computer images used for the *Phone* stereograms: an undistorted image, an image distorted for the phase conjugate-illuminated transfer, and an image distorted for the direct-illuminated transfer.

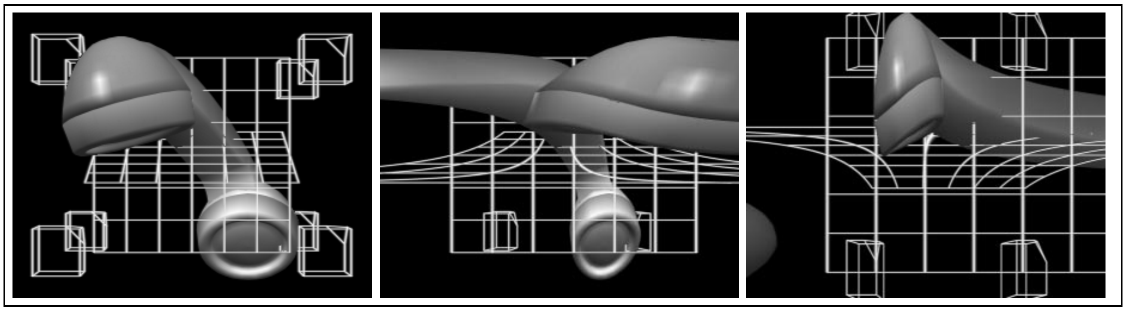

Figure 5.1: Three images of the computer model for *Phone Test Pattern*. On the left is a single viewpoint of the undistorted test pattern. The middle image is a predistorted frame for the *Phone #1*, where the transfer is illuminated in phase conjugate and the master is projected into space. On the right is a frame for *Phone #2*, which use direct illumination for the transfer.

Cadillac Hubcap and Wheel, MIT Version

Type: two-step reflection
Transfer illumination: direct illumination, white light
Master size: 300mm wide, 200mm high
Transfer size: 250mm wide, 200mm high
Projection frame size: 250mm wide, 200mm high
Slit width: 1mm
Number of slits: 300
Master-transfer separation distance: 100mm
Resolution of projected images: 600 pixels wide, 480 pixels high
Stereogram-type: simple camera
View Distance: 1 meter

Cadillac Hubcap and Wheel, Diner Version

same as *MIT* except for subject.

The *Cadillac Hubcap and Wheel* series of stereograms were conceived to show that a first quality two-step display hologram could be made in the ULTRAGRAM format. All of the stereograms in this series used the same holographic mastering setup, shown in Figure 5.2.

The images for this series were generated from a computer aided design database of a Cadillac automobile hubcap and wheel assembly supplied by the General Motors Design Staff. The spline database was converted to polygons and rendered using the Rendermatic library. The same predistortion methods used with *Phone* were used in the *Cadillac Hubcap* series, including the finite instead of infinite horizontal camera viewpoint. The final transfer hologram was a single color Lippmann type reflection hologram in order to minimize chromatic aberrations produced by the parallel plate transfer geometry. In

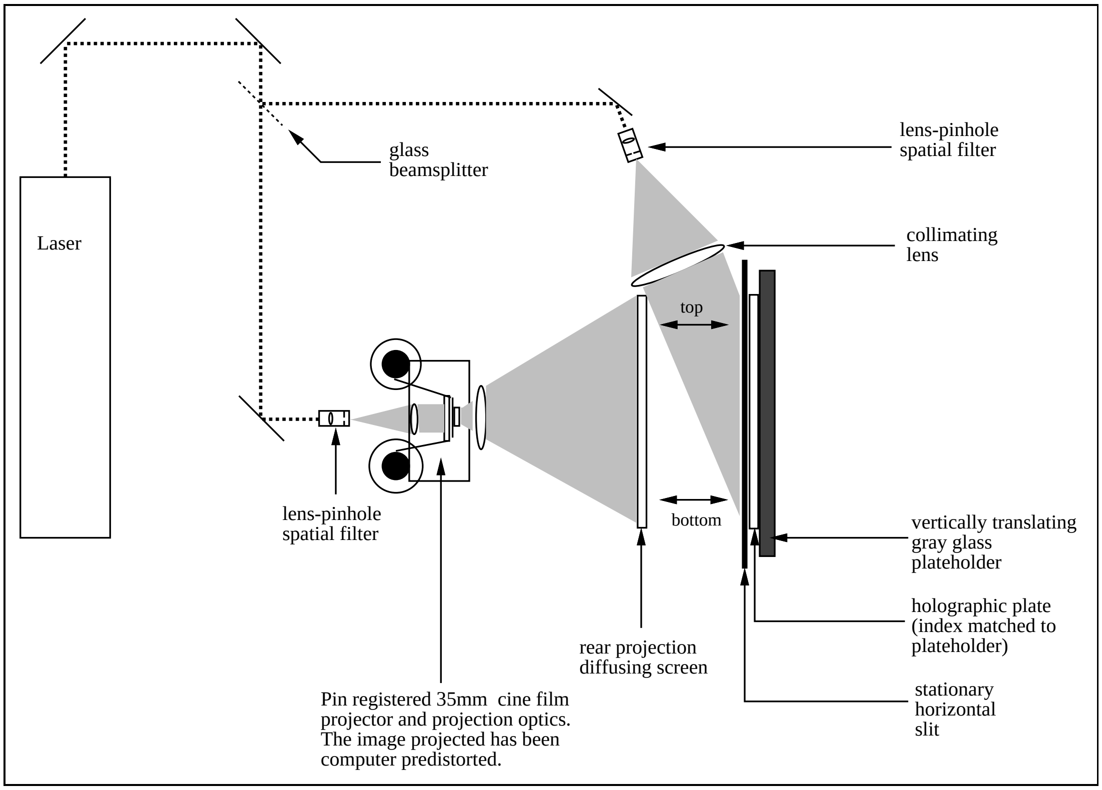

Figure 5.2: The holographic table layout used to make the slit master holograms for the *Cadillac Hubcap* series.

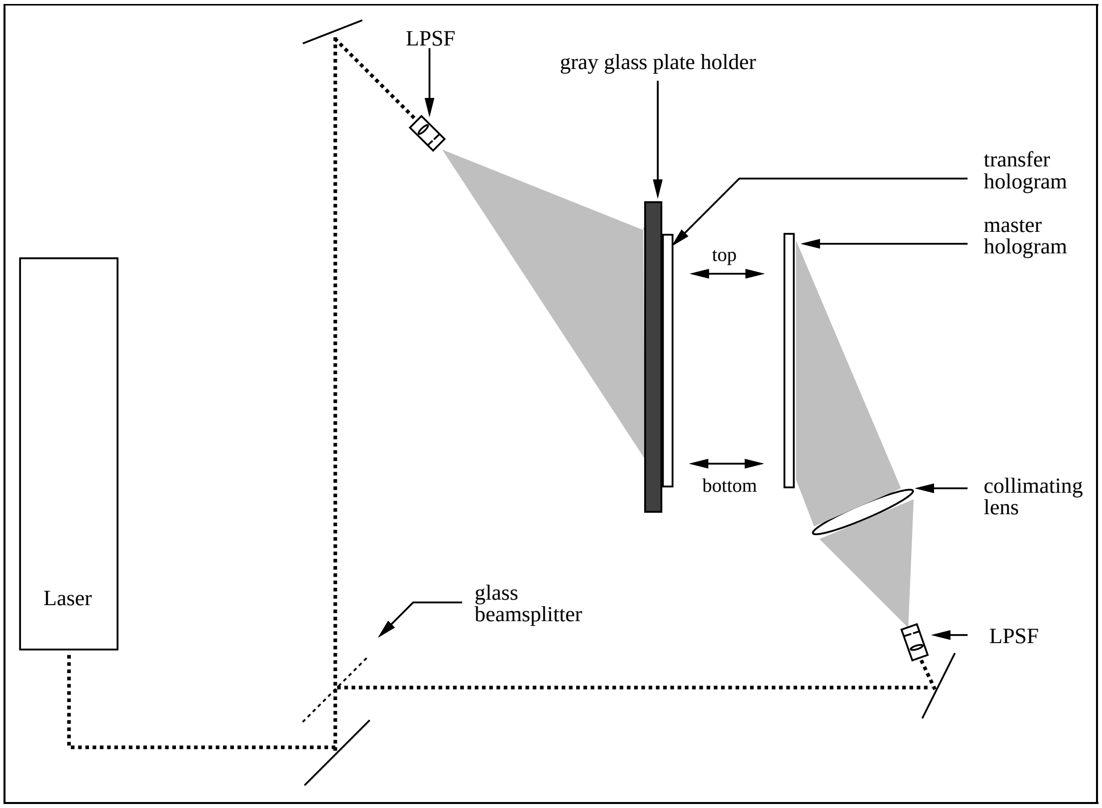

Figure 5.3: Holographic table layout for a direct illumination transfer of an ULTRAGRAM master.

Cadillac Hubcap and Wheel, MIT, the wheel was positioned mostly in front of the image plane and was filled with a reflection of a distinctive view of the MIT campus. The color of the reflection transfer used in *Cadillac Hubcap, MIT* was chosen to be close to the wavelength of the laser used to make it (647nm Krypton) to minimize the geometrical effects of a wavelength shift. No bandlimiting was applied beyond that provided by the infinite camera predistortion technique.

The transfer setup for the *Cadillac Hubcap* stereograms is shown in Figure 5.3. The first transfers of *Cadillac Hubcap, MIT* that were produced showed a very noticeable distortion due to the misplaced horizontal focus of the taking camera. The images appeared

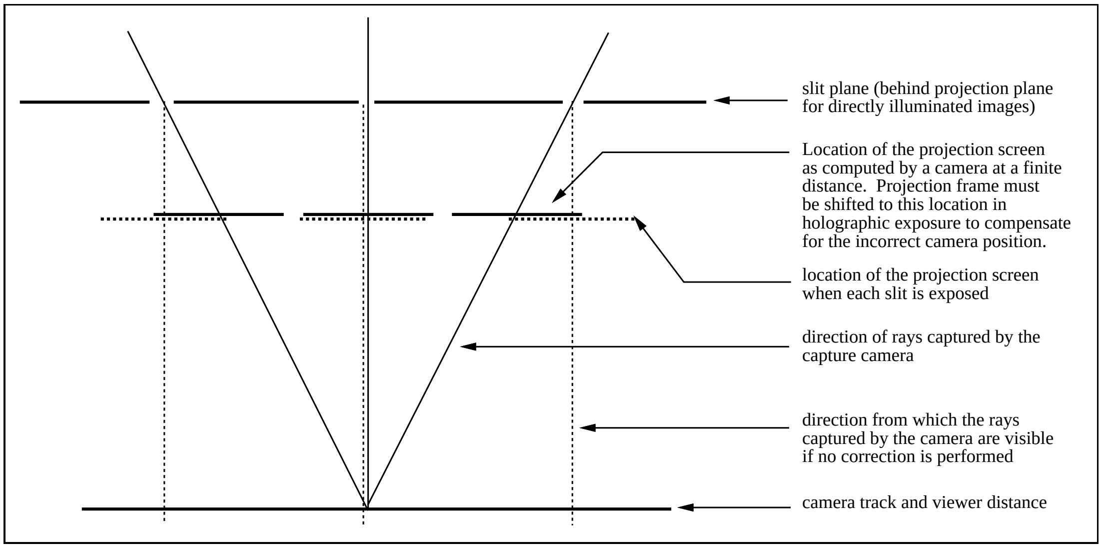

Figure 5.4: When a camera with horizontal focus at the view distance rather than at infinity is used in infinite camera predistortion, the camera captures rays through each slit that are not centered around the perpendicular to the slit, but rather around the line that passes through the slit from the middle camera viewpoint. To correct for this error, the projection frame must be shifted during the holographic exposure of the master to match the region that the camera captured.

compressed horizontally by an amount that changed with the depth of the object. The circle of the hubcap, for example, appeared to be a slightly tall ellipse.

The theoretical origin of this distortion was unknown at this point, so its effects were emperically compensated for during the holographic exposure by slightly translating the images horizontally on the projection screen from one view to the next. This translation partially solved the problem for the following reason, shown pictorially in Figure 5.4. In the infinite camera model, the view taken from the middle of the camera track contains rays that pass perpendicularly through the slits after travelling through the middle of each non-recentered projection frame. If the camera is located instead at a finite distance, only the middle ray of the image traces through to the middle of its frame. All other rays pass through the slits at some angle that becomes greater for slits far from the center of the hologram, and the offset between the center of the middle ray and the center of the screen becomes larger. The computed projection frame for each slit is still centered around the ray from the center camera view, but the offset of the center ray on the projection plane shifts the computed projection frame with respect to the physical projection frame. As a result, the the computed and actual projection frames do not occupy the same region of space. Lens translation slid each actual projection frame over to the location of the computed projection frame, thus partially restoring the correspondence between capture and viewing.

Once this distortion was corrected, the image met or exceeded all expectations. When the transfer hologram was viewed, the fact that the slits of the stereogram appeared behind the image plane was not noticed by most observers. The transfers were bright and had about a ninety degree wide field over which some part of the image was visible to the viewer. The same rendering, distortion, and holographic process was used to create *Cadillac Hubcap and Wheel, Diner*, in which the object appeared larger and closer to the plane of the hologram. The wheel surface reflected a view of the Fell's Point Diner from the movie *Diner*, used without permission. Figure 5.5 is a picture of the stereogram. For the *Diner* ULTRAGRAMS, a range of different colors, from deep red to golden, were used for the reflection transfers with no significant geometric distortions introduced. The *Cadillac Hubcap* ULTRAGRAMS subjectively seem more real and tangible than conventional stereograms of similar subjects. The reason for this tangibility is not easily explained.

The *Cadillac Hubcap* stereograms exhibit a vertical striping artifact that may be due to aliasing. At first such an artifact may seem unlikely because no part of the object in either image extends much further from the projection plane than the projection plane-slit plane distance, and the infinite camera technique guarantees bandlimiting out to at least the plane of the slits. However, the optical effect produced by the reflecting chrome of the hubcaps and wheels is that of concave and convex mirror structures, capable of moving the apparent depth of objects such as the relection map to locations far from the projection plane. The striping artifacts do not appear attached to the objects or the slits, but to the

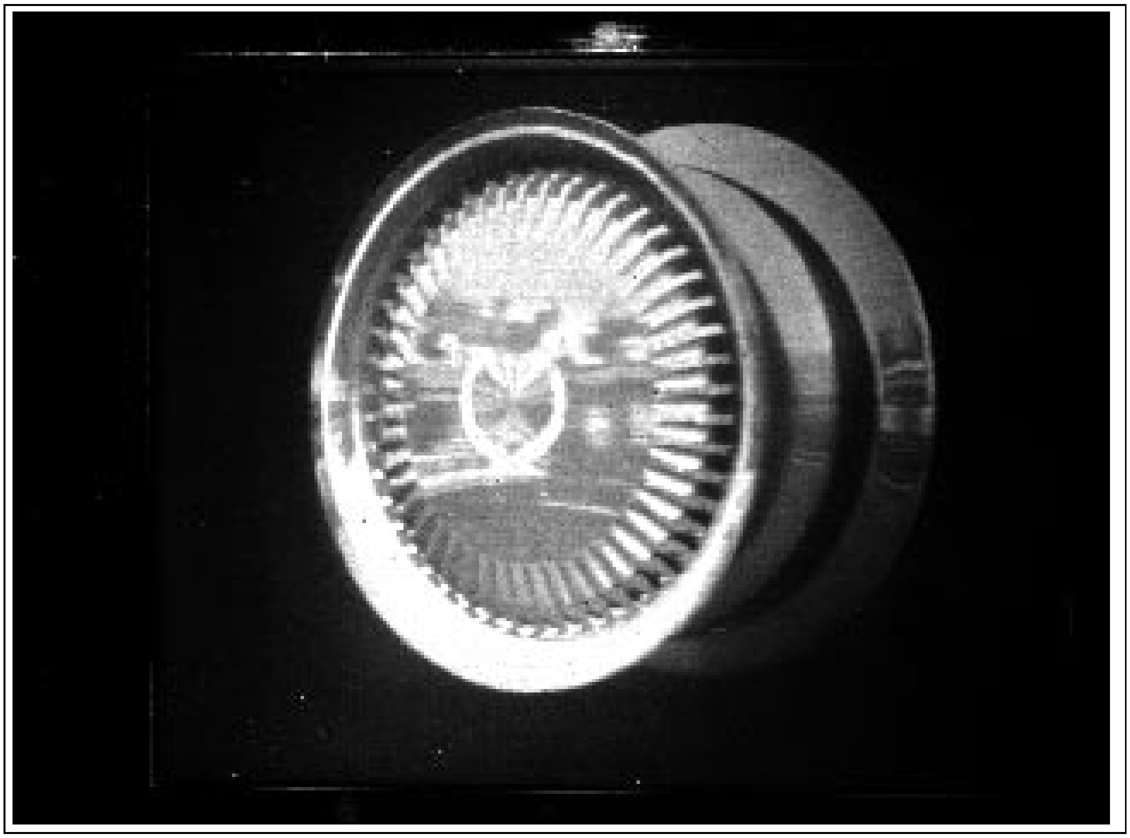

Figure 5.5: *Cadillac Hubcap and Wheel, Diner.*

projection plane, which supports the aliasing hypothesis. The appearance of this aliasing artifact is similar to that of an insufficiently anti-aliased computer graphics image. Other explanations for the artifact have not been ruled out.

Large Chevrolet Wheel

Type: one-step transmission
Transfer illumination: direct illumination, monochromatic (mercury arc)
Size: 1 meter square
Projection frame size: 980mm square
Slit width: 1mm
Number of slits: 1000
Slit-projection screen separation distance: 500mm
Resolution of projected images: 780 pixels square
Stereogram-type: simple camera
View Distance: 2 meters

The next application of advanced stereogram techniques was the production of large one-step stereograms. This project was intended to demonstrate that large scale images with arbitrarily large view distances could be produced in a relatively small laboratory. Other than this work, almost all large scale display stereograms are made using two-step transfer holography, which requires that the viewer stand at the plane of the master. When the transfer is made, then, the master and transfer plates must be separated by the intended view distance. Such a transfer could only be made in a very large, vibration free holography lab. Furthermore, the proper phase conjugation required for these transfer holograms is very difficult to achieve and as a result the images are especially prone to distortions when viewed. A one-step stereogram, on the other hand, requires stability only for the slit being exposed, and has no difficult transfer step. A one-step stereogram with an image that crosses the holographic plate is only possible using predistortion techniques.

The holographic exposure apparatus for the one-step ULTRAGRAM is shown in Fig-

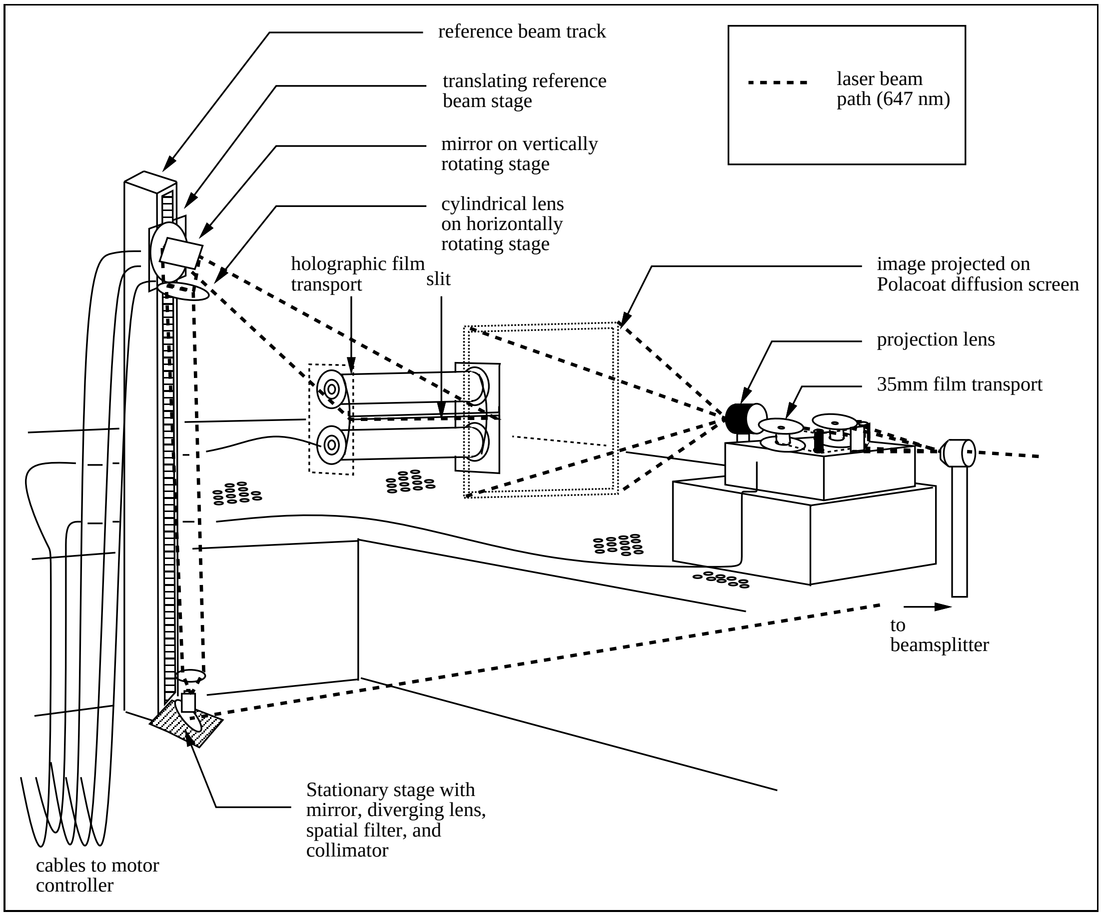

Figure 5.6: The table layout used to expose the large scale ULTRAGRAM.

ure 5.6. Each predistorted image image is projected onto the large rear-projection screen, which is centered in front of a 1mm high slit machined in an opaque mask. Behind the slit is a film transport mechanism capable of advancing the 1.1 meter wide film from the supply roll to the takeup roll in 1mm steps. A vacuum mechanism holds the film stationary against the slit during exposure. Because the slit and the projection screen remain fixed with respect to each other, the stereogram follows the simple, non-recentering camera model. The exposure of all 1000 slits that compose the hologram took approximately twelve hours to complete, due mostly to long settling times between exposures. When the hologram is viewed, the projection screen appears to be one-half meter behind the plane of the slits.

Each slit is visible through an angle of ninety degrees.

The need for direct illumination of the final hologram dictated the use of the most unusual piece of optics in this system: an optical tower used to translate the stereogram's reference beam between every slit exposure. The tower's design models the final illumination of the stereogram. When the stereogram is lit by a point light source above and behind it, light strikes each slit at a slightly different horizontal angle. This angle must be matched during each slit's exposure in order to avoid image distortion. The tower translates the reference source to precisely the same location with respect to the slit as the illumination source occupies when the slit is lit. The illumination source can then provide exact direct illumination for all slits simultaneously.

This large scale ULTRAGRAM was the first to correctly implement the infinite camera capture geometry. The infinite horizontal viewpoint was produced by warping the object's points to new locations in space that just cancel the normal effects of camera perspective. Once correctly implemented, the infinite camera technique yielded an undistorted final image. Two stereograms were produced; both images were of another automobile wheel, this time from a Chevrolet, with a view of the sign in front of the General Motors Technical Center reflected by its surface. The first stereogram placed the object almost completely in front of the plane of the hologram, protruding about one-half meter into space. Figure 5.7 is a picture of this hologram. This image's most significant shortcoming is that it suffers noticeable aspect ratio changes as the viewer moves closer or further away from the hologram than the intended two meter view zone distance. This distortion can be seen in Figure 5.8. However, the realistic, intriguing quality of the aerial image of the hubcap is little short of astounding. In the second image, most of the hubcap is placed between the plane of the slits and the projection screen. This stereogram is at least as realistic as the first, but is not as visually fascinating as the aerial image. The second stereogram does, however, more closely maintain the correct aspect ratio of the object as the viewer changes positions. The slit structure on the surface of both images is visible but not intrusive; some slits are

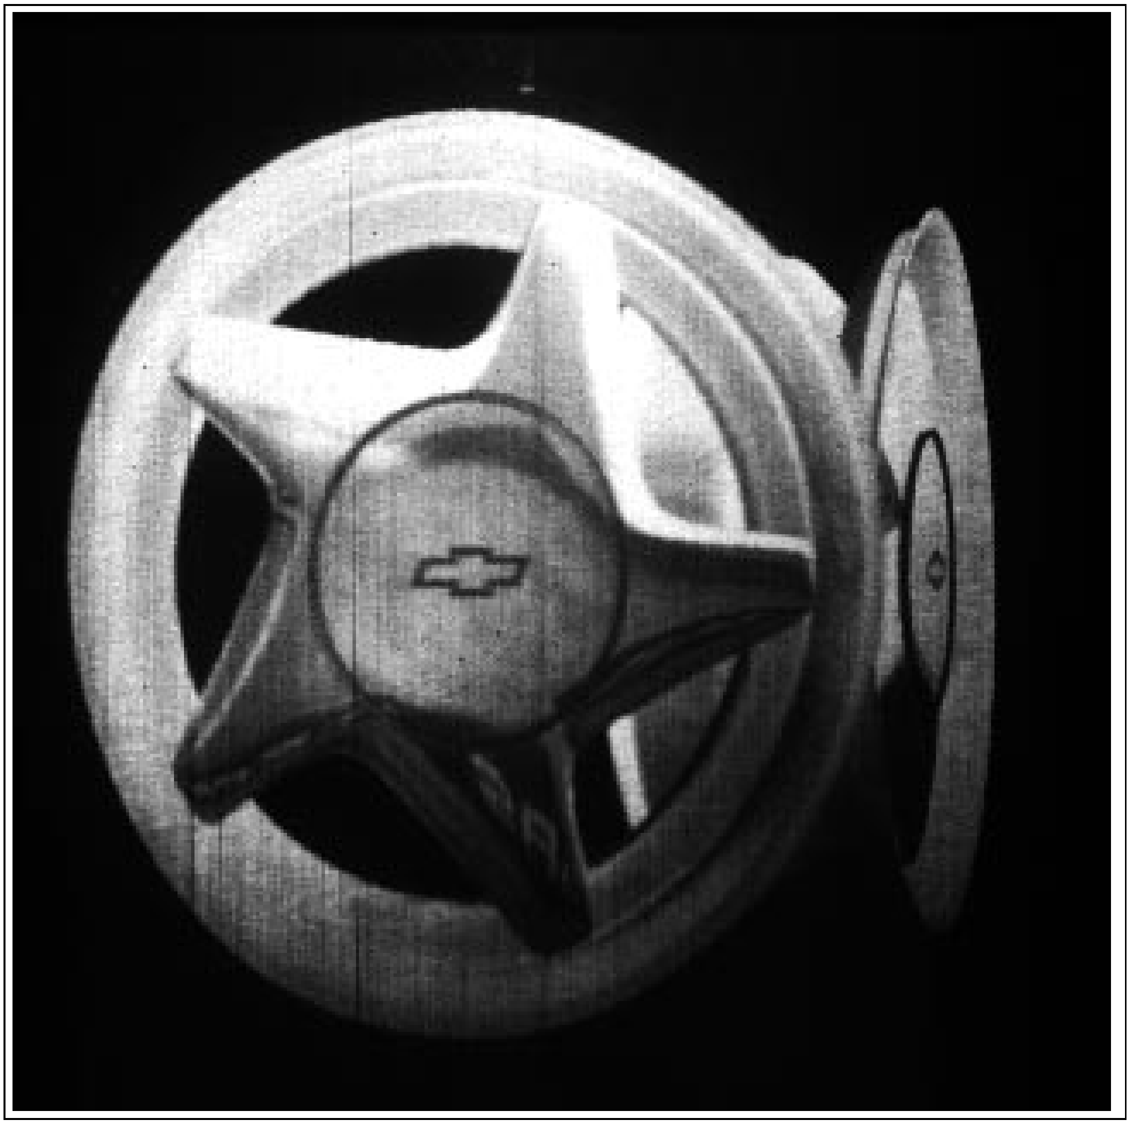

Figure 5.7: *Large Chevrolet Wheel*, as seen from the intended view zone, about 2 meters from the hologram.

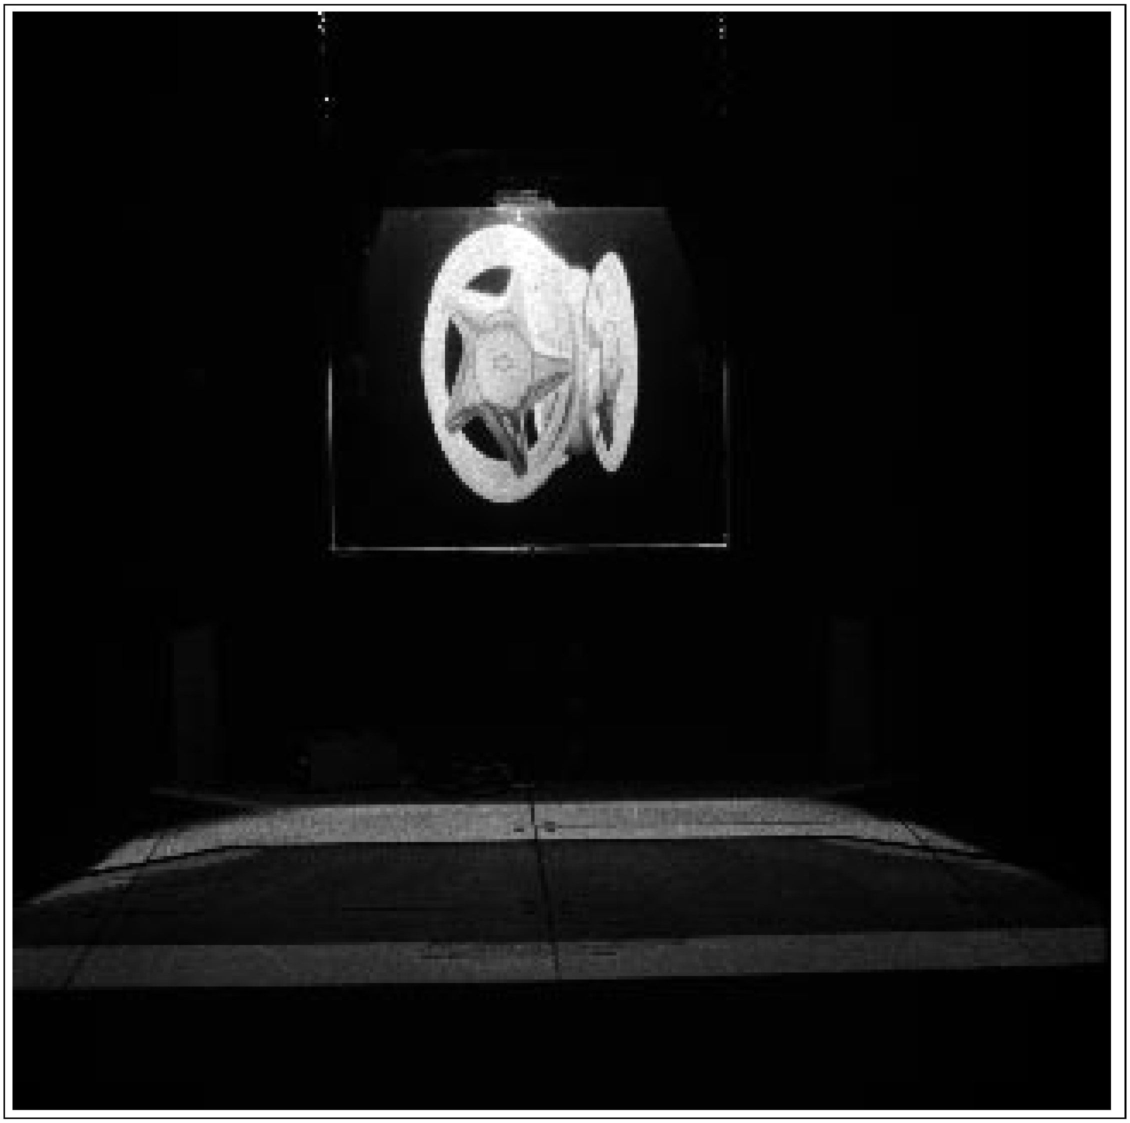

Figure 5.8: *Large Chevrolet Wheel*, seen from about 8 meters. The image exhibits aspect ratio distortion when it is not viewed from the correct view distance.

dimmer than others due to irregularities in holographic exposure. Because the projection and slit planes are so disparate in the large one-step ULTRAGRAM, illumination is provided by a monochromatic mercury arc lamp so as to avoid chromatic blur. Work continues on making one-step white light viewable stereograms, including ones producing deep aerial images insensitive to changes in viewer location.

Breakfast Attempt, Ultragram Version

Type: two-step reflection
Transfer illumination: direct illumination, white light
Master size: 300mm wide, 200mm high
Transfer size: 100mm wide, 100mm high
Projection frame size: 100mm wide, 100mm high
Slit width: 1mm
Number of slits: 300
Master-transfer separation distance: 250mm
Resolution of projected images: 880 pixels wide, 880 pixels high
Stereogram-type: recentering camera
View Distance: 500mm

The most recent ULTRAGRAM is *Breakfast Attempt, Ultragram Version*. This stereogram shows a photorealistic computer graphics rendering of the highest quality in a directly illuminated, high resolution reflection two-step hologram. It was particularly intended to be suitable for contact copying *en masse*. The fanciful computer graphic scene in the hologram, depicting a failed early-morning cooking and serving extravaganza of the Scandinavian persuasion, was designed by two graduate students in the Spatial Imaging Group and rendered using the Rendermatic computer graphics library. The mastering geometry for this stereogram is very similar to that used in the *Cadillac Hubcap* series. The plate separation was increased to 250mm to retain as much spatial resolution far from the image plane as possible. While the size of the master hologram was the same as *Cadillac Hubcap*, the transfer plate was much smaller (100mm wide by 100mm high). *Breakfast Ultra* was

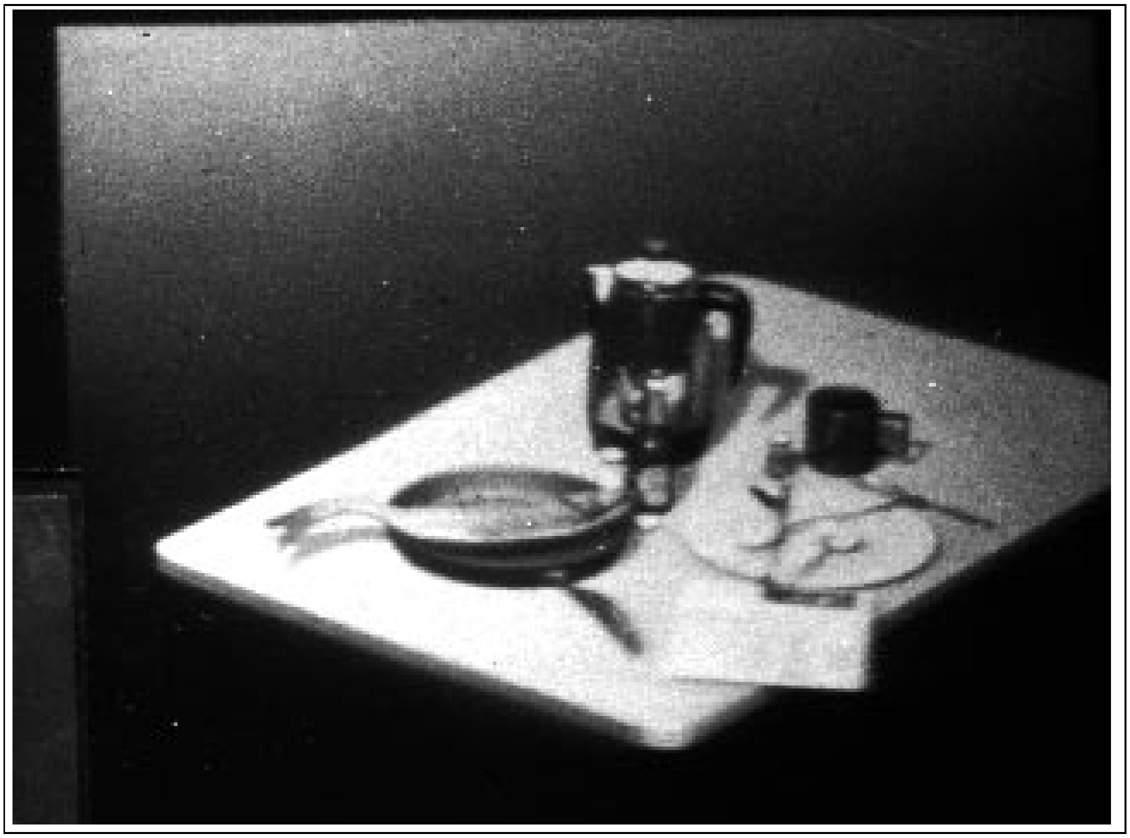

Figure 5.9: *Breakfast Attempt, Ultragram Version*, recorded on silver halide film.

the first stereogram to use a recentering camera geometry, with image predistortion done using the perspective slicing technique. Both the horizontal and the vertical foci of the input camera were located at the intended view distance of 500mm. The difference in wavelength between the laser and the intended color of the transfer was taken into account during predistortion. Recentering was done in the holographic setup using a lens translation mechanism to move the projection of the image to the correct location with respect to the slit being exposed.

The success of the perspective slicing predistortion method was evident even before the final transfer of *Breakfast Ultragram* was made. The master, when viewed is a collimated laser beam, projected a tiny, solid looking three-dimensional scene a quarter meter in front of the stereogram plate. The final transfer image, shown in Figure 5.9, exhibits no significant distortions and is fairly insensitive to changes in viewer depth. Although the size of the transfer hologram's view zone was deliberately sacrificed in favor of improved spatial resolution, all of the image can be seen over a 0.8 meter-wide viewing zone located 0.5 meters from the hologram, with parts of the image visible over 1.4 meters. The image exhibits some blurriness, possibly due to poor focus on the projection screen. As true as with the other ULTRAGRAMS, *Breakfast Attempt, Ultra*'s three dimensionality was extremely believable. Unfortunately, the image has not yet been mass produced.

The broad gamut of hologram types and sizes to which ULTRAGRAM techniques have been applied demonstrates their wide applicability to stereography. From large, one-step transmission holograms to small, wide angle two-step reflection ones, ULTRAGRAPHIC techniques can reduce the complexity and expense of producting conventional stereogram formats, and can be used to make new images that are impossible to achieve by any other means.

Chapter 6

Conclusion and Future Work

One of the most powerful tools of scientific thinking is analogy: an unknown system can be better understood by comparing and contrasting its various properties with those of better-understood systems. Systems cease to be isolated; instead, they fall into general classes about which sweeping statements can be made. The previous discussion has removed the holographic stereogram from the isolated context of three-dimensional display and classified it within the broader scope of discrete and continuous optical systems. The finite width and number of slit apertures makes the stereogram discrete. As a result of this discreteness, the holographic stereogram is susceptible to aliasing and sampling artifacts. Because of the large number of apertures, though, the holographic stereogram behaves in some ways like a continuous system. For example, horizontal parallax only stereograms experience image distortions equivalent to continuous HPO displays when the viewer strays from the intended view distance. Understanding a general class of system is also often easier than understanding any one specific case, just as understanding a continuous HPO display is easier than understanding an HPO stereogram.

Once classification has been made, methods that apply to the general class of problems can be brought to bear on the specific case. Bandlimiting, for instance, can be used to eliminate the aliasing artifacts in holographic stereograms, and anamorphic imaging can correct for problems of astigmatism. While the details of any specific system, such as slit size or viewer distance, may vary, the problems and solutions that apply to the general case still hold. Thus, a general holographic stereogram model, free of all but essential constraints, can be produced by adhering the lessons of the general models of discrete and anamorphic optical systems. Attention to the correspondence between the pieces of the model, namely the capture, recording, and viewing geometries, "fine tunes" the general model to fit specifics such as viewer position, location of the plane of the slits, or required spatial resolution.

The economy and insight provided by an understanding of both the optical properties of a specific stereogram and of the relationship of that stereogram to the general model are now indispensable tools of the Spatial Imaging Group. ULTRAGRAM techniques as presented here are one aspect of those benefits. Because of them, displays of fundamentally different format and improved quality can be produced in a smaller, less expensive laboratory. But future work comes not just from the specific ULTRAGRAM techniques but from extending this conceptual approach to new display geometries.

Some extensions of ULTRAGRAMS are relatively straightforward. The large scale one-step ULTRAGRAM can be scaled in size to 1×3 meters to provide full size holographic models for designers and engineers. The size of the ULTRAGRAM view zone can be increased by using larger, more Lambertian projection screens. Improvements need to be made in the abutting of adjacent stereogram slits when the slit plane is clearly visible. And although the difficulty of the full-parallax process does not currently outweigh the result, the application of ULTRAGRAM predistortion to full-parallax stereography is direct.

Some basic research still needs to be done. The advantages and difficulties of holographic slits with non-rectangular intensity profiles should be explored. Perceptually, the eye's tolerance to aliasing artifacts in stereograms and the usefulness of focus in HPO systems should be used to set lower and upper limits on the number of slits required in a stereogram. Lastly, an improved understanding of three-dimensional computer graphics rendering for stereograms may greatly improve the speed with which the many perspective views needed are generated.

More interesting are the possibilities for new display formats. Work proceeds on edge-lit stereograms that benefit from ULTRAGRAM distortion correction. Color ULTRAGRAMS are possible by predistorting each color separation slightly differently, based on the effect of wavelength-related aberrations. Other research is working to replace the conventional rear-projection screen in the stereogram exposure apparatus with a large holographic optical element to increase the efficiency of the exposure process.

The combination of these technologies will make possible a computer peripheral similar to today's laser printers but capable of printing out white light viewable, three-dimensional images. Such a device, using three laser diodes and an HOE projection screen for exposure, could produce full color images that could be edge-lit in a special "walkman-sized" viewing station, or conventionally lit in normal room light. Such a device will finally bring the benefits of truly high quality three-dimensional image creation and display to our dimensionally starved culture.

Appendix A

Glossary of Terms

holographic stereogram A hologram composed of a series of apertures, each recording a two-dimensional view. Projectional information of a scene can be recorded in the two-dimensional views. When displayed, a holographic stereogram forms a discrete approximation to the original three dimensional scene. In this text, a holographic stereogram is often called a **stereogram** for brevity.

one-step hologram A hologram created using a single holographic plate. A **one-step stereogram** is a single holographic plate exposed in many slits.

two-step hologram A holographic copy of another hologram. The copy is called a **transfer hologram**, and is made from a **master hologram**, which is made from the object. A **two-step stereogram** consists of a holographic transfer of a stereogram master.

horizontal parallax only (HPO) Having true parallax only in the horizontal, and not the vertical, direction. If a scene is displayed on a **HPO display**, all horizontal detail in the scene appears to be located at its true depth, while all vertical detail appears to lie at the same plane. In an **HPO stereogram**, all vertical detail lies at the projection plane. Such a stereogram is made using a series of long, thin, vertical apertures, called **slits**.

full parallax Having parallax in both horizontal and vertical directions. Laser transmission holograms usually have full parallax. a **full parallax stereogram** consists of a two-dimensional array of apertures.

capture geometry The geometrical relationship between camera and subject.

holographic recording geometry The geometrical relationship between a stereogram's slits and the projection screen when the stereogram is exposed. Details of holographic transfer, reference and reconstruction beam direction, and laser wavelength may affect the effective holographic recording geometry.

viewing geometry The geometrical relationship between the viewer and the final hologram. The illumination of the hologram may alter the effective holographic viewing geometry.

physical camera A mechanical or electronic camera used to take pictures of natural scenes. Usually made of metal, plastic, and glass.

synthetic camera A computer analog of a physical camera, used to make images of computer databases of three-dimensional objects. Also called a **virtual** or **computer graphics** camera.

scene The group of objects of which a stereogram is an image. A **natural scene** is a scene of "real world" object. A **synthetic scene** is a three-dimensional computer database captured with a synthetic camera.

object Something in a scene. An object is composed of **object points**.

object point The smallest three-dimensional detail on the surface of the object.

image A two-dimensional projection of a three-dimensional object. The image of an object point is its projection in a two-dimensional view.

view Similar to and often interchangeable with **image**. A view, or **perspective view**, is captured by a camera looking at a scene. A view is seen from a **camera position**, also called a **viewpoint**.

simple camera a camera whose lens axis is fixed with respect to the film plane. Standard still and cine cameras are simple cameras.

recentering camera A camera whose lens stage can translate with respect to the film plane. Such a camera can be used to select a window in a scene that will always appear at the same position on the film plane. A view camera with a translating or “shifting” lens stage is an example of a recentering camera.

projection screen The physical object onto which the perspective views of a stereogram are projected. In this thesis, the projection screen is a rear-projection screen. Such a screen is usually made out of ground-glass or Polacoat. The projection screen lies at the projection plane.

projection frame The two-dimensional area of the projection plane that is fill by the image projected when a slit is exposed. The projection frame may be at different locations for different slit exposures. The depth at which the projection frame itself appears appears to be located is called the **apparent projection frame distance**. This depth need not correspond to the projection plane, but in a recentered camera stereogram the two do coincide.

projection plane The infinite plane in which the projection screen lies. A stereogram of a scene forms a wavefront approximation to that scene composed of pieces of wavefronts emitted from this plane. The **holographic projection plane** is the plane of the projection screen in the holographic apparatus. The **camera projection plane**

is the plane of the object that corresponds to the the holographic projection plane. The **viewer projection plane** is the plane at which the image of the holographic projection screen is reconstructed when the final hologram is viewed.

simple camera stereogram A stereogram that requires the images produced by a simple camera to expose its slits. In this type of stereogram, the projection screen is located directly in front of each slit during exposure. In the perspective images used to expose this type of stereogram, objects far from the camera remain at a fixed location from image to image, while objects closer to the camera translate across the camera's field of view. A correct holographic apparatus for exposing this type of stereogram would be one that held the projected image and slit fixed and translated the holographic plate behind the slit.

recentering camera stereogram A stereogram that needs the images from a recentering camera to expose its slits. The projection frame of this type of stereogram remains fixed in space relative to the holographic plate and thus translates relative to each slit during exposure. The projection frame defines a "window" of information at the apparent projection frame distance that all slits record. This distance usually corresponds to the projection plane. The recentering camera required to produce the views for this stereogram type translates its lens stage to keep objects on the camera projection plane stationary through all perspective views.

continuous system A system represented by a signal of arbitrarily high frequency.

discrete system A system composed of periodic samples of a continuous signal. Discrete systems are prone to **aliasing**.

spatial velocity In a stereogram, the rate at which an object point moves through the sequence of projectional views. All object points have a constant spatial velocity based on their distance from the apparent projection frame plane.

spatial frequency A measure of the amount of detail in an image. The higher the spatial frequency, the more detailed the image, and the smaller the minimum extent of any features in the image.

depth of field In relation to a lens, the relationship between and the distance of an object point from the lens' plane of focus and the sharpness of that image on the lens' focal plane. The image of such a point is called a **circle of confusion**. Depth of field is regulated by the effective diameter of the lens.

bandlimiting Removing high frequency information from a signal to avoid aliasing artifacts. A signal with no such high frequency information is called **bandlimited**. Bandlimiting is done using a **filtering** process.

aliasing An unwanted, periodic artifact that occurs in discrete systems when the rate that a signal is sampled is not high enough to capture all the signal's information.

ray tracing A technique of optical analysis. Also, a computer graphics technique to produce projectional images of three-dimensional scenes.

anamorphic camera A camera equipped with a lens whose horizontal and vertical foci are located at different planes.

infinite anamorphic camera A computer image-distortion method that produces a sequence of images with horizontal foci at one plane from another sequence of images that have horizontal foci at infinity.

orthographic projection An image of scene produced so as not to exhibit any foreshortening in one or both directions. The effect is that of a viewer an great distance from an object. A **horizontally orthographic projection** is used in the infinite anamorphic camera: the camera produces images of normal perspective projection vertically but orthographic perspective projection horizontally.

perspective slicing A computer image-distortion method that produces a sequence of images with horizontal foci at one plane by combining vertical slices of images from another sequence with an arbitrary horizontal focus location.

conventional stereogram A holographic stereogram in which the plane of the viewer and the image of the stereogram's apertures coincide. The images for a conventional stereogram can be captured using a non-anamorphic camera lens. Sometimes called a **traditional stereogram**.

advanced stereogram A holographic stereogram in which the plane of the viewer and the image of the stereogram's apertures can be located at arbitrary positions with respect to the plane of the final holographic display. Advanced stereograms are a superset of conventional stereograms. In general, advanced stereograms require that the camera that captures the perspective images have an anamorphic lens to compensate for image distortions produced by the stereogram. Advanced stereograms are also called ULTRAGRAMS.

Appendix B

Two Extended Examples

This appendix will present two examples of ULTRAGRAM creation: one using the infinite focus anamorphic camera technique, the other using perspective slicing.

Infinite camera

The hologram to be produced is a master hologram for a direct-illuminated, white-light-viewable transfer. The process is summarized in Figure B.1. The master stereogram is 250mm wide and 200mm high. The slits on the master are 1mm wide; there are thus 250 slits on the master. The master plate appears to be 100mm behind the transfer, so the projection plane-master separation and the master-transfer separation must both be 100mm. Assume that collimated light is used to both reference and reconstruct the master hologram. The projection screen is 300mm wide and 200 mm high. The stereogram uses a simple-camera geometry, so the projection screen is located directly in front of each slit during exposure. The raster image projected on the projection screen has a resolution of 768 pixels horizontally by 512 pixels vertically.

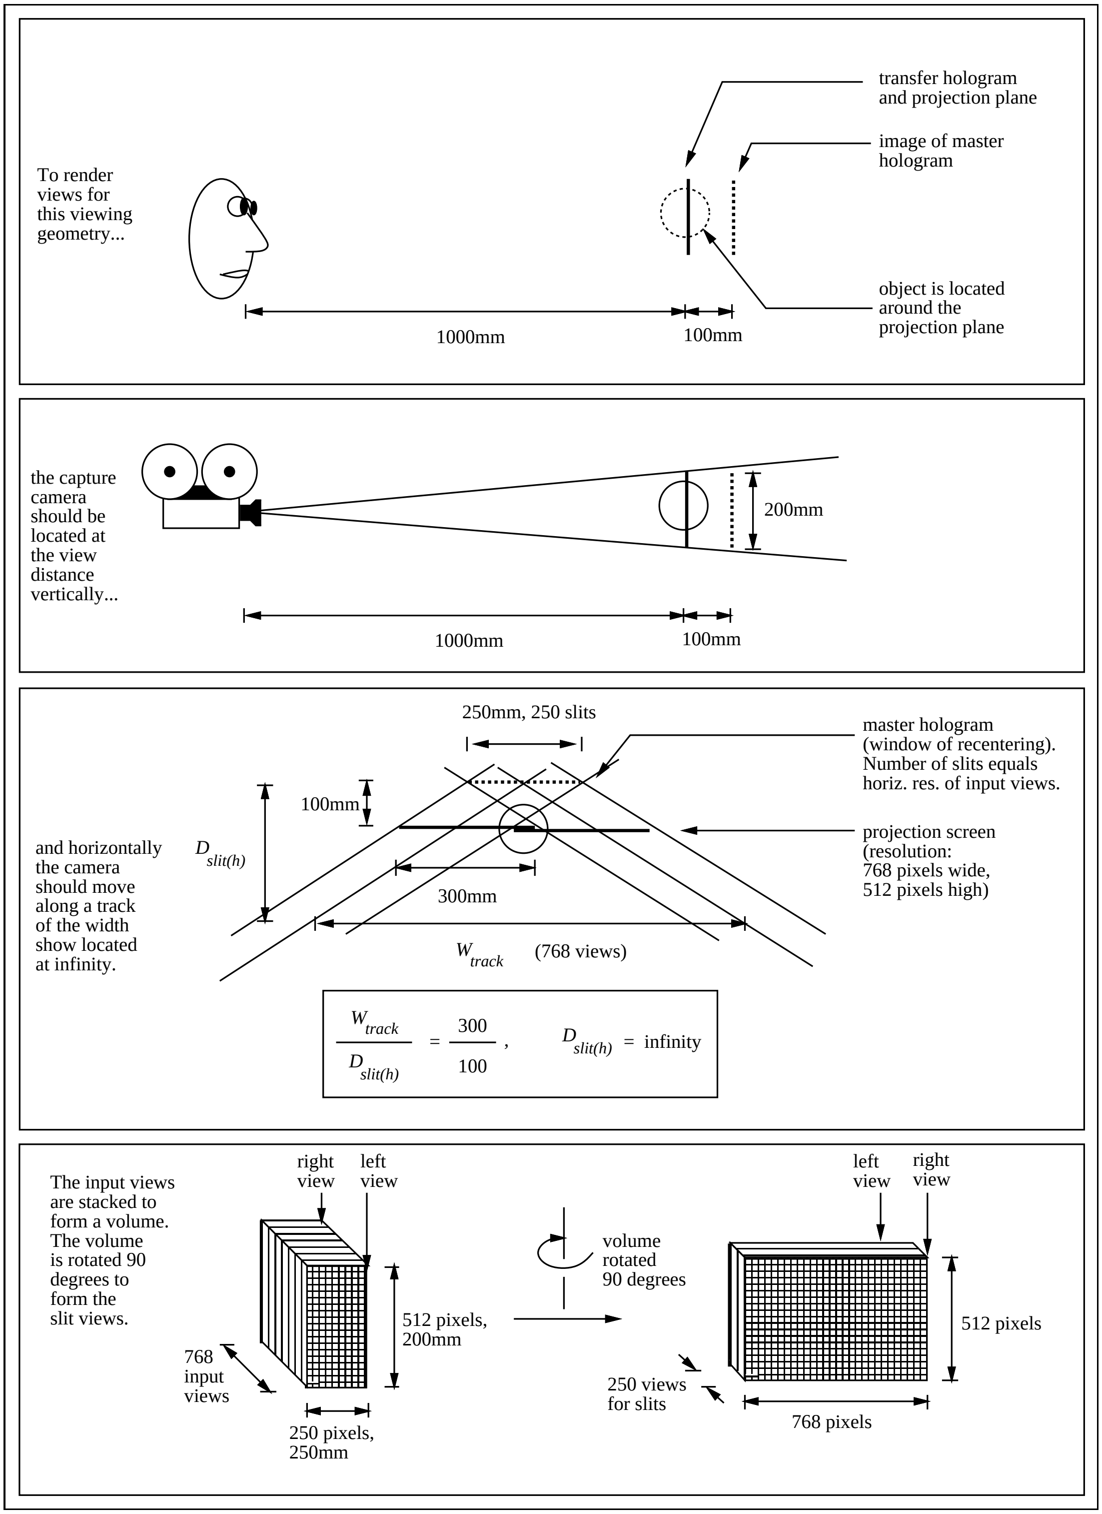

Figure B.1: The infinite camera ULTRAGRAM creation process.

Image generation begins as follows. The vertical focus of the computer graphics camera is located at a distance corresponding to the final view distance, 1 meter from the projection (and transfer) plane. The scene is centered on the projection plane. The horizontal focus of the camera is located at infinity. The ratio of the camera track's width to the camera's distance from the slit plane is equal to the ratio of the projection screen's width to the screen's distance from the slit plane, in this case 3 to 1. Thus, if a camera with a horizontally orthographic projection were on a track 1 meter from the slit plane, the track must be 3 meters wide.

The camera recenters around the master plane as it moves down the track taking views. A view is captured for every pixel of horizontal resolution of the projection screen image, in this case 768. The 768 views are spaced equally on the track. The vertical resolution of the captured views equals the vertical resolution of the projection screen, 512 pixels. The camera captures the entire region of the master stereogram, but the horizontal resolution of the images equals the number of slits in the master, here 250. The aspect ratio of the rendered images is not the same as the aspect ratio of the region that they portray (in other words, the images are anisotropic).

Once all the input images are rendered, they must be predistorted. The predistortion consists of rotating the volume composed of the input views and resampling them to yield the set of output views. The **rotate_volume** code used shown in Figure 4.10 can be used for this rotation. The output views are then suitable for slit exposure.

Perspective Slicing

This hologram is also a direct illuminated two-step master, but the master plate is designed to be located 250mm behind a 100mm square transfer plate, which is 250mm from the viewer. Figure B.2 shows the production process. As in the previous example, the

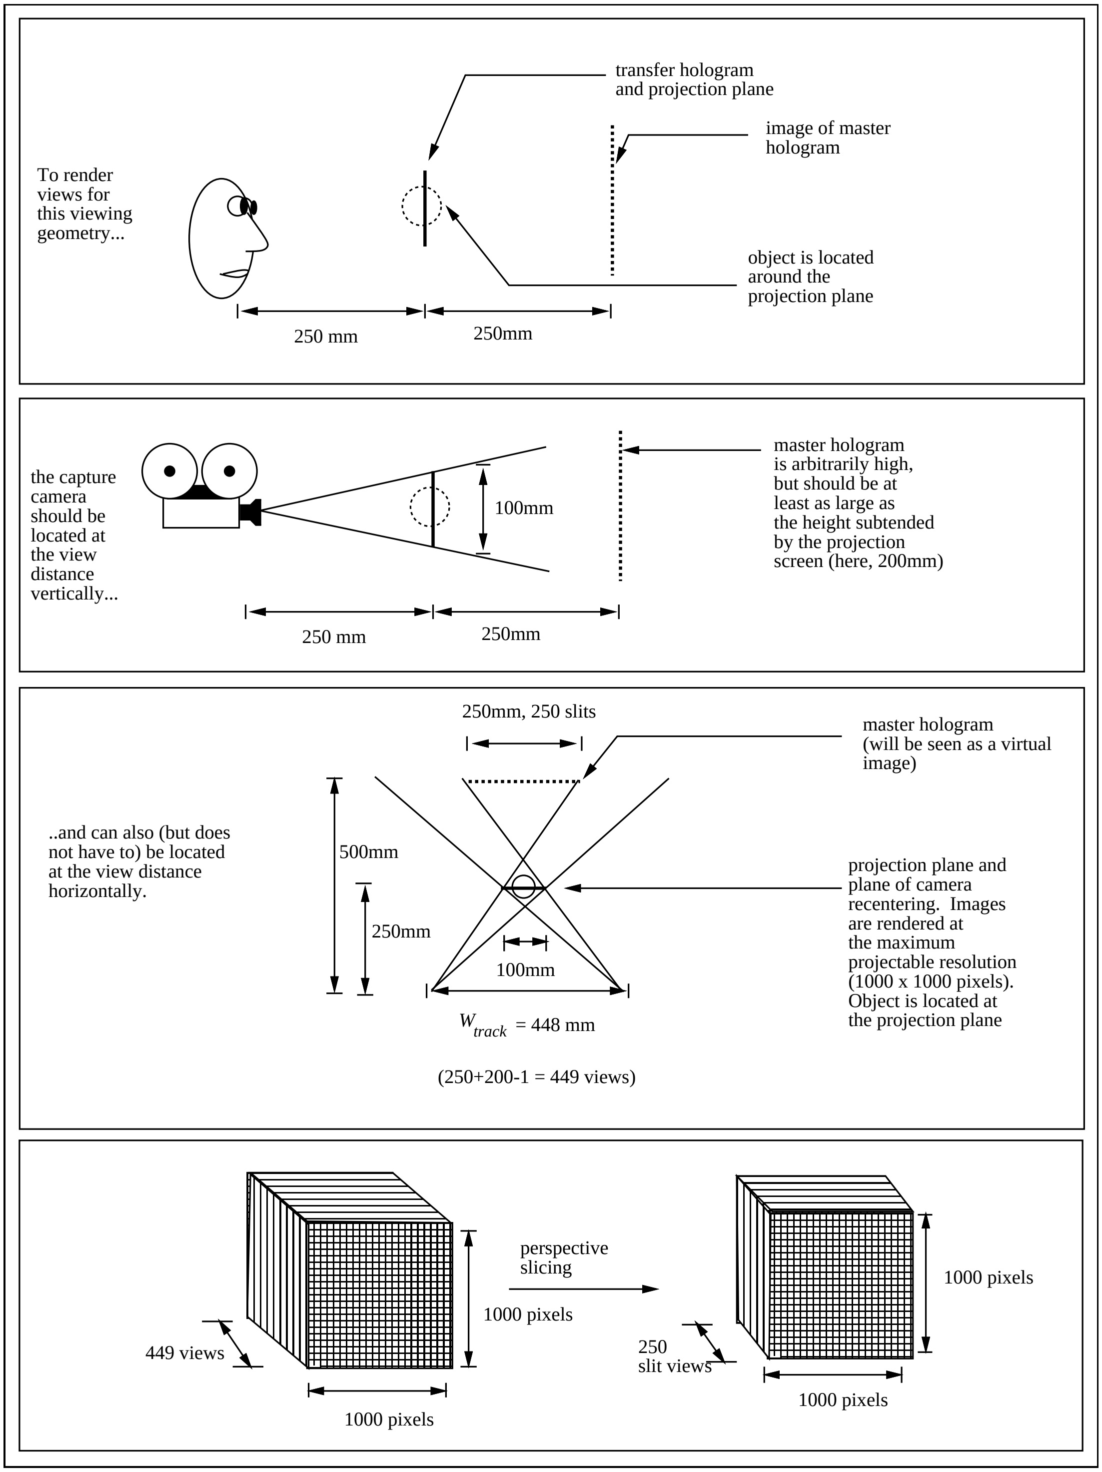

Figure B.2: Perspective slicing predistortion.

master plate consists of 250 1mm slits. This stereogram is of recentering-camera type, so perspective slicing is used to predistort the images. A image of resolution 1000 pixels square is displayable on the projection screen. Both the undistorted and distorted images are rendered at this resolution.

The camera used to produce the input views has both lens foci at the intended view distance with respect to the capture projection plane, 250mm. The camera recenters its lens to always frame the area that corresponds to the projection frame. Using Equation 4.3, the camera track width is found to be 448mm. The number of views that the camera captures can be found using Equation 4.4; here it is 449. Once again, the camera track is divided equally among views. After the input images are rendered, the **perspective_slice** code in Figure 4.14 may be used to produce the output images.

Bibliography

[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] N. D. Haig. Three-dimensional holograms by rotational multiplexing of two dimensional films. *Applied Optics*, 12–2:419–420, February 1973.

[6] Mark Holzbach. Three-dimensional image processing for synthetic holographic stereograms. Master's thesis, Massachusetts Institute of Technology, September 1986.

[7] L. Huff and R. L. Fusek. Color holographic stereograms. *Optical Engineering*, 19–5:691–695, September/October 1980.

[8] 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.

[9] E. N. Leith. White light holograms. *Scientific American*, 235:80–95, October 1976.

[10] E. N. Leith and J. Upatnieks. Photography by laser. *Scientific American*, 212:24–35, June 1965.

[11] 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.

[12] Herbert C. McKay. *Three-Dimensional Photography*. American Photographic Publishing Company, New York, 1953.

[13] William J. Molteni. Star wars. Holographic Stereogram.

[14] Michael A. Teitel. Anamorphic raytracing for synthetic alcove holographic stereograms. Master's thesis, Massachusetts Institute of Technology, September 1986.