Improved Rendering of Parallax Panoramagrams for a Time-Multiplexed Autostereoscopic Display

Adam Kalai[a,*] and Mel Siegel[b]

[a]Department of Computer Science and [b]The Robotics Institute
School of Computer Science, Carnegie Mellon University, Pittsburgh PA 15213

ABSTRACT

Parallax-barrier panoramagrams (PPs) can present high-quality autostereoscopic images viewable from different perspectives. The limiting factor in constructing PP computer displays is the display resolution. First, we suggest a new PP display based on time multiplexing in addition to the usual space multiplexing; the barriers move horizontally in front of the display plane. The time multiplexing increases the horizontal resolution. In addition, it permits us to use wider barriers than are acceptable for static displays. We then analyze these displays, showing that wide-barrier PPs have advantages relating to depth-resolution and smoothness, and we present a novel algorithm for rendering the images on a computer.

KEYWORDS: stereoscopy, stereo cameras, 3D-TV, disparity, occlusion, view interpolation

1. BACKGROUND

Parallax panoramagrams (PPs) are capable of displaying stereoscopic 3D images, viewable without the aid of glasses[3]. These autostereoscopic displays also provide horizontal motion parallax, or "lookaround," meaning that the viewer sees around objects as she moves her head from side to side. Recall that a PP is a two-plane device in which the front plane is a fence-like array of thin vertical barriers, and the back plane is a standard two-dimensional display such as a piece of paper or an LCD. As Figure 1 illustrates, from above, different vertical strips are visible from different eye positions. Note that a single strip may be visible to the viewer's left eye from one viewpoint and visible to the viewer's right eye from another. Therefore, the strips correspond to different eye positions and cannot, in general, be labelled as left-eye or right-eye views. However, at any particular location, each of the viewer's eyes sees a different strip, and thus the viewer perceives a stereoscopic image. The multiple views of PPs provide the viewer with motion parallax; studies indicate that motion parallax is more important than static binocular perspective for certain visualization tasks[10].

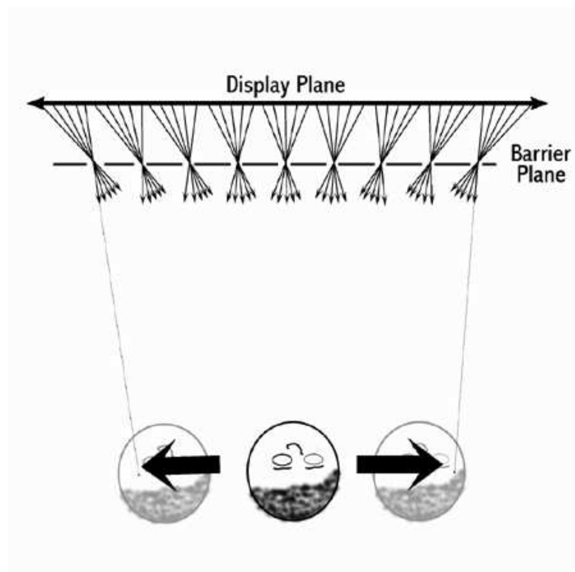

Figure 1. A standard PP shown from above. Different strips of the display plane are visible from different eye positions, creating a stereoscopic effect.

There are three main drawbacks of PPs. First, since the available horizontal resolution has to be shared among the views, horizontal resolution is reduced for any single view. Second, wide vertical barriers are visible and distracting. Third, the barriers reduce the average brightness by occluding much of the light from the display. Typically, high-resolution screens combined with very thin barriers or lenticules are used to address these problems.

2. INTRODUCTION

In this paper, we show how time multiplexing combined with PPs can significantly improve the perceived resolution and decrease the visibility of the barriers. Towards this end, we describe a new PP display in which the barrier plane is an LCD,

* Additional author information --
Adam Kalai (corresponding): akalai+@cs.cmu.edu, http://www.cs.cmu.edu/~akalai.
Mel Siegel: mws+@cmu.edu, http://www.cs.cmu.edu/~mws, voice 412 268 8802, fax 412 268 5569.

and the barriers are opaque columns of pixels. We move the barriers horizontally and simultaneously update the display plane. This creates a high-resolution autostereoscopic 3D display with motion parallax. Analogously, when driving by a picket fence, a viewer sees a full stereoscopic image with motion parallax, albeit with diminished contrast; the pickets, which are moving rapidly relative to the observer, are not distracting. Similarly, our PP barriers, if moved quickly enough, will not be distracting to the human eye.

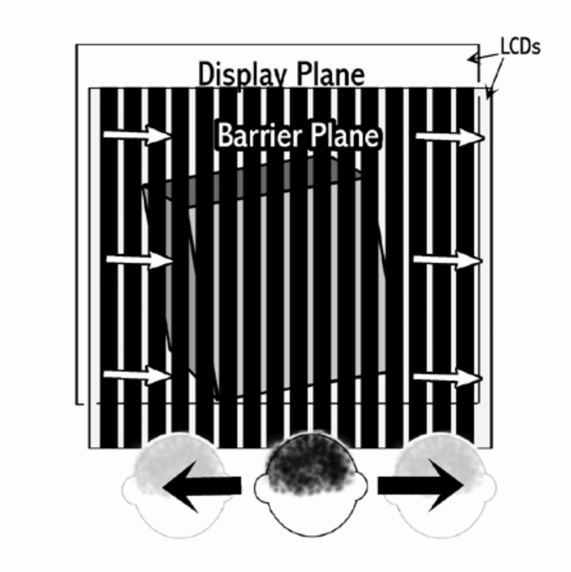

Figure 2. The time-multiplexed PP shown from the front. As the barrier plane, displayed on the front LCD, changes its position, the display plane must be updated accordingly.

There are advantages of using thin barriers and narrow gaps, and advantages of using wider barriers and gaps. Current displays have slits that are narrow enough to show just one column of pixels; thus an eye at any point will be presented with a single view or, at most, a combination of two views. This is pleasing because the thin barriers are hardly visible to the eye. On the other hand, with larger slits, which we call gaps, and larger barriers, we could present an eye at any point with a smooth combination of many more views. This produces a more continuous feel as the viewer moves her head side to side, as in hard copy displays generated by similar photographic techniques[6]. These wide barriers are more suitable for the time-multiplexing display than static PPs because, as mentioned, they are less distracting when moved quickly.

In Section 5 we introduce a new rendering algorithm for large-barrier displays. This algorithm renders a complete (appropriately distorted) image from the perspective of each gap, rather than the usual separate image for each view.

This time-multiplexing display allows for choosing better horizontal resolution at the expense of a faster refresh rate. As the barriers move more quickly, the perceived resolution increases, but a faster refresh rate is required for the display plane. For example, if the barriers don't move at all, we have a traditional PP, with a fraction of the display plane resolution. On the other hand, if the barriers move quickly enough so that each point on the display plane is frequently visible to each viewer location, a viewer at any location sees a picture with full horizontal resolution. The lower refresh rate of interlaced displays suggests that we might not even need to double the refresh rate to double the perceived resolution, particularly if we use smaller barriers.

3. RELATED WORK

Aside from traditional PPs and lenticular displays[5], several related time-multiplexing systems exist today. The cyclostereoscope, demonstrated in Paris in 1949, consists of a giant circular fence rotating around a large screen, with a narrow triangular viewing zone. The design allows for only two distinct views, so there is no motion parallax. As with all two-view display devices (without head tracking), the viewer's head must be fixed in place, a characteristic which the general public has not accepted. Reports indicate that an Australian company has been working on a modern version of such a device[4]. From its description, their conception appears to be a two-view device like the cyclostereoscope, where LCDs are used in place of the fence. This time-multiplexing approach seems to be similar to ours, with the essential difference that it does not provide for motion parallax. It is not clear that time multiplexing is preferable to using higher resolution displays for two-view devices. By providing only two views, it is simpler to engineer but more restrictive to the viewer.

The parallactiscope and other moving-slit methods are similar to our approach, but have a single slit open at any time[8]. By opening only one slit, they can provide a very wide viewing zone, but the slit motion and display drawing must be extremely fast. The Cambridge display, now being sold by Infinity Multimedia, is a form of a moving slit display where, via optics, a virtual slit is placed right in front of the viewer's eyes[2]. This sound technique can produce high-quality stereoscopic images but requires a refresh speedup of a factor of at least the number of viewpoints presented and has consequent impact on the required display brightness.

Lastly, Dimension Technologies Inc. reports working on stereoscopic laptops that have similar geometrical properties to parallax barrier displays, using bright vertical strips of light placed behind the display plane[1]. While they focus on two-view devices, they also report prototypes for time-multiplexing displays that provide motion parallax and are more optically efficient than their PP counterparts. These displays contain proprietary lighting technology.

4. TIME-MULTIPLEXING DISPLAY

In this section, we propose a simple way of making high-resolution PPs using standard-resolution displays and time multiplexing. When compared to super-high-resolution displays, time multiplexing is advantageous because displays of the required resolution are very expensive (or nonexistent) and also because we can use larger barriers to create smoother motion parallax. Alternatively, it can be combined with high-resolution displays to further boost the resolution.

We suggest multiplexing over time by changing the horizontal position of the barrier plane, which is accomplished by moving the fence-like image on the barrier LCD. As shown in Figure 2, one possibility is to move the planes horizontally to simulate physical motion. However, it is not clear that linear motion of the barriers is the least distracting choice. With several possible barrier positions, it might be better to follow a nonlinear pattern or randomly choose positions, which may or may not overlap. While these alternatives would be difficult to implement with a mechanical barrier, they are easy to try with an LCD barrier.

There are several reasons to believe that moving barriers are less distracting. For one, the human eye is not distracted by the black screen displayed 72 times per second on a movie screen. Secondly, as mentioned, we see an acceptable image as we drive by a picket fence.

There are a several design parameters for this display. The most important are the parameters of the LCD, since these cannot be changed easily. For building a prototype, we could use two LCDs of size 25 cm by 18 cm, standard resolution, say 1024x768, that address their pixels at a fast rate such as 120 Hz. Second, we assume the viewer sits at a distance of about 80 cm from the display and the viewing zone is 40 cm wide. This leaves us flexible to try different placements of the barrier plane and different barrier widths. By placing the barriers at a separation of 6 mm from the display plane, we can calculate that they must be 3 mm wide. We could choose a gap width of 0.6 mm so that the gaps cover one sixth of the screen. Our analysis in Section 6 indicates that to have an acceptable amount of jumping as the viewer moves her head side to side, we must keep the virtual depth within a small range of 48 mm. While this may seem like a severe restriction, large amounts of parallax are problematic for other reasons, causing headaches on some devices and severe convergence-accommodation conflict. Furthermore, even small amounts of parallax are enough to convey valuable depth information[9].

5. RENDERING

In this section, we introduce a computer rendering algorithm for large-barrier PPs. Several lenticular computer displays exist, and they use the same images as parallax barrier displays[5]. However, with these static, small barriers or lenticules, there are a limited number of columns of pixels behind each thin barrier; two columns of pixels per strip is typical, four is impressive. The images are constructed from a small number of views (2-4 ordinary pictures) which are then horizontally multiplexed, column by column. The first column comes from the last image, the next column comes from the previous image, and so on, as shown in Figure 3, with six views

Our images, on the other hand, consist of columns from many more views, potentially hundreds when the barriers are very wide. At some point, the barriers are so wide that the images behind each barrier combine more views than there are barriers. It then becomes profitable to render a separate image per barrier, rather than rendering a separate image for each view.

This difference in computer rendering techniques is analogous to a difference in photographic techniques for creating PPs. One technique involves taking a number of photographs and then multiplexing them with a process involving barriers in front of film. This is analogous to the standard technique for creating computer PPs. An alternate technique involves moving a camera horizontally, while the slits are swept across the film, with the camera shutter open during this process. This creates a smoother image with a wider angle of view[6]. This latter technique is similar to our rendering algorithm in that it renders the image behind each barrier as a single, smooth image.

Figure 3 and Figure 5 show example renderings*, Figure 3 with strip width of 6 pixels, and Figure 5 with strip width of 150 pixels. The displays have different properties when observed behind barriers. The barriers used in both examples have a gap width of one-sixth the strip width. As shown in Figure 4, the first display produces non-uniform error. From some vantages, there is no error; one sees exactly what a picture taken from that perspective would look like at that resolution. From other vantages there is considerable error. In the second display the error is almost exactly the same from all vantages, as shown in Figure 6. Each view has error increasing away from the center of the gaps, so no view is "perfect." However, our analysis in

*Rendering was performed by Alice, a 3D graphics package for Windows (free at http://www.cs.cmu.edu/~alice).

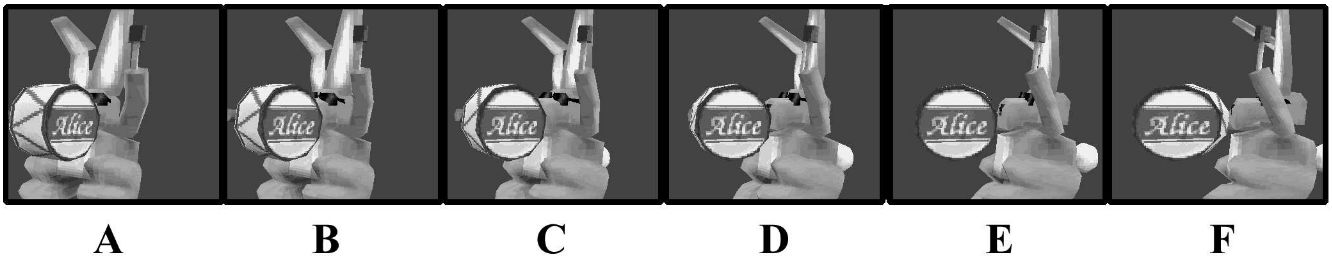

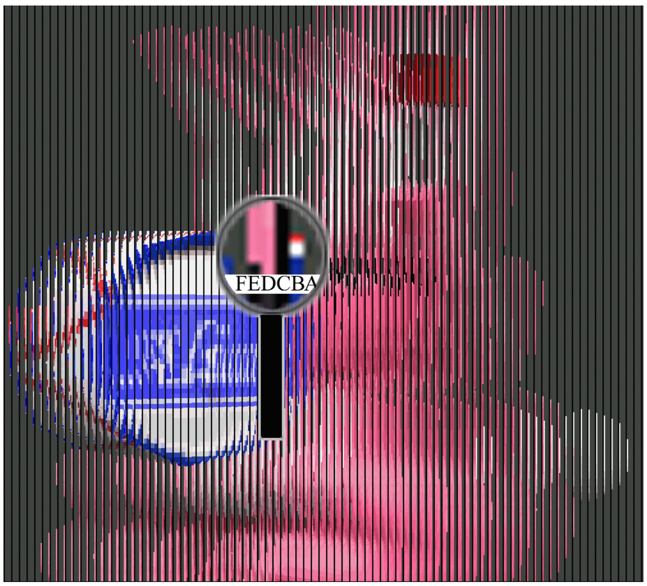

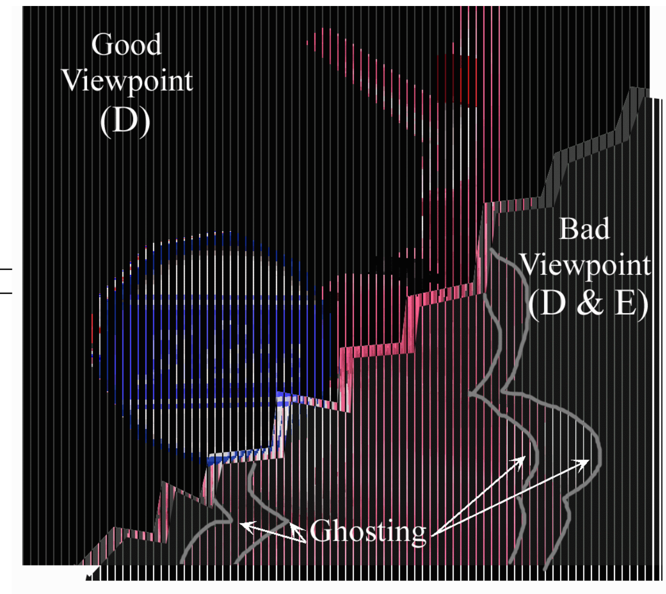

Figure 3. A PP generated from the 6 perspective views A-F. As shown in the magnifying glass, vertical strips are taken, in reverse sequence, from these views.

Figure 4. This same PP viewed behind barriers. From a good viewpoint, only one of the six views is visible. From an inferior viewpoint, two views are visible.

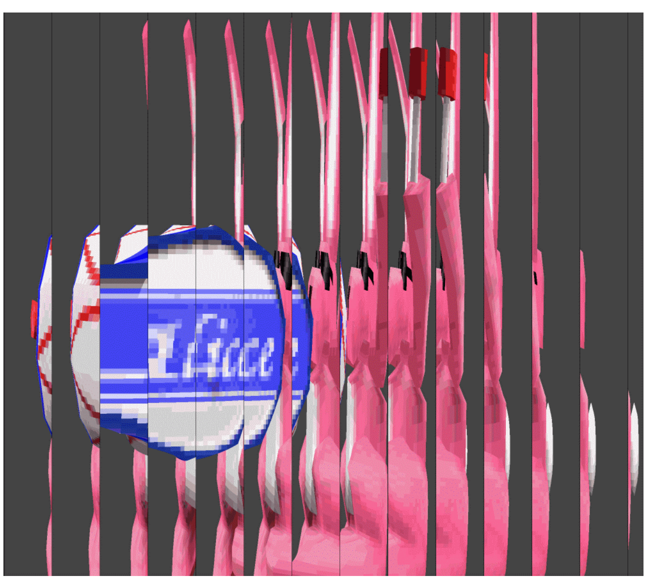

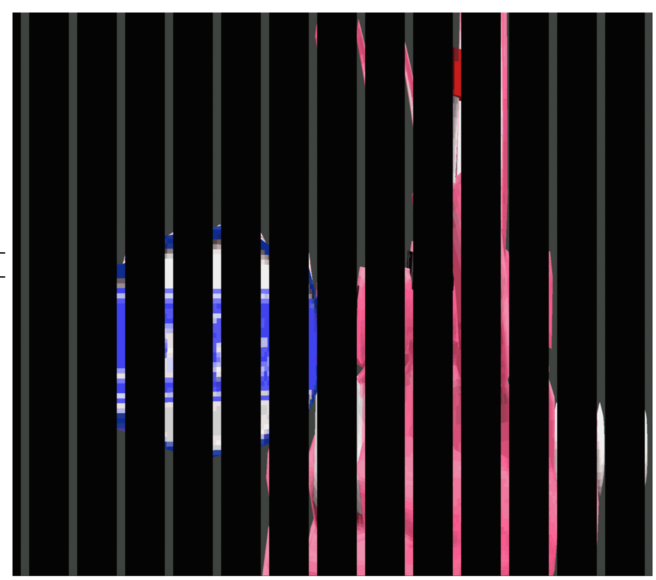

Figure 5. A wide-barrier PP. Each strip is an anamorphic projection of a part of the scene.

Figure 6. The same wide-barrier PP viewed behind barriers. In this case, the error is the same from all viewpoints.

Section 6 indicates that this provides less average error in a sense that we describe there. Furthermore, as the viewer moves her head from side to side, she sees less “jumping” on the second display.

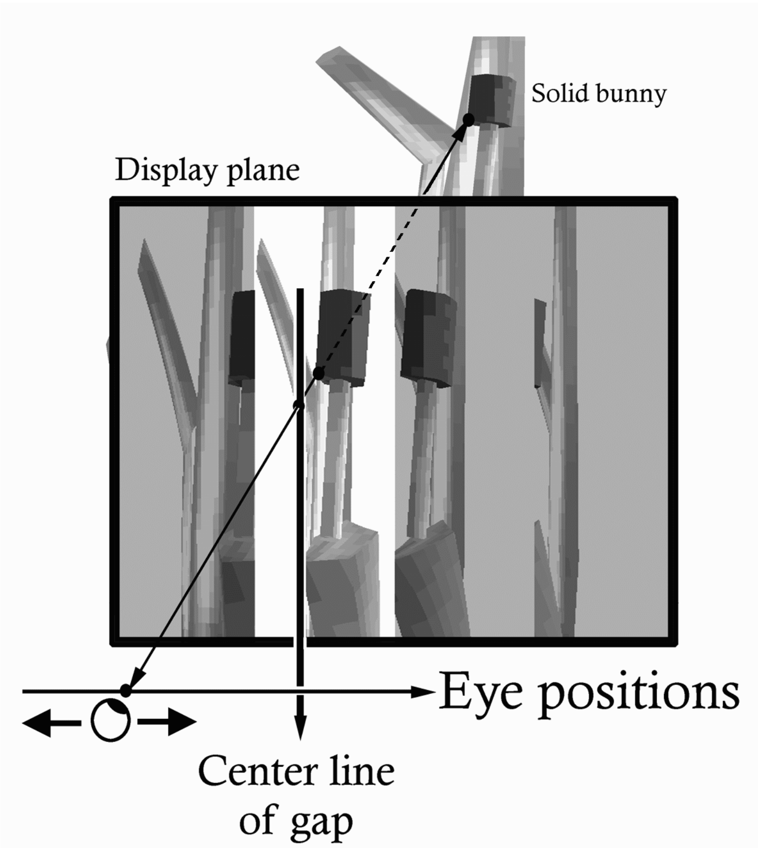

Figure 7. The projection used to generate the wide-barrier image. The line, which connects the point on the display plane and the corresponding point on the solid bunny, must intersect both the eye-position line and the gap line.

The continuous rendering algorithm works as shown in Figure 5. We assume the viewer’s eyes are somewhere on a line segment of a fixed (possibly large) width at a fixed distance from the barriers. Conceptually, we assume that the gap width is infinitesimal. We project the 3D space simultaneously through the center line of the actual gap and the line of the eye positions onto the display plane. For any point on the display plane, there is exactly one line which passes through this point and each of these lines. This line associates a point in 3D space with a point on the display plane. It is interesting to note that such a projection, through two lines onto a plane, is not technically a “projective transformation.” In this projection, lines in the 3D space are rendered as curves, whereas a projective transformation renders straight lines as straight lines.

6. ANALYSIS OF PARALLAX PANORAMAGRAMS

In this section, we consider the various parameters of PPs and the resulting image characteristics. The quantity we will determine is the width of a display pixel at a virtual depth **z** relative to the display plane. As shown in Figure 8, each pixel on the display plane corresponds to a virtual volume of space that could be visible “through” that pixel, depending on the viewer’s head location. We calculate the width of this volume at a depth **z**, and see that it decreases as the barrier plane is moved closer to the user and the barriers are made larger.

The virtual pixel width has several interpretations. First, people recognize lenticular displays or PPs from the familiar small jumps of the image as the viewer rotates the (postcard) display a small amount or moves her head from side to side. This is unlike, say, a true hologram where the virtual points appear to move smoothly across the display. The jumping occurs because the low resolution of the display implies that only a small number of views are represented. Consider what happens as the viewer moves her head from side to side. As each pixel comes into view through a gap, it appears frozen (in terms of motion parallax) at the wrong depth, the depth of the display plane, until it goes out of view. The virtual pixel width is a measure of how much a virtual point at a certain depth should move, while, in fact, it appears frozen at the wrong depth. Second, the virtual pixel width is related to blur at a fixed depth. As seen in Figure 8, these virtual pixels overlap an amount that directly depends on their virtual width.

We assume that the viewer’s head is at a fixed distance from the barriers and within a fixed horizontal range. As the viewer’s eye moves from left to right, the strips of the display plane become unoccluded from right to left. If the eye moves out of that range, say past the right side, the observed image returns to the view from the leftmost point of the range. So, in fact, there are multiple viewing ranges, and several people can view the display at the same time. On the other hand, it is disconcerting to move one’s head from one zone to the next because, via aliasing of the association of particular slits and columns, the image abruptly “flips” between the leftmost and the rightmost views. Even more problematic, when the two eyes are in different viewing zones the viewer experiences pseudoscopic stereo. In this sense, it is preferable to have wider viewing zones.

There are five parameters to our analysis:

d - the distance between the viewer and the display plane.

w - the width of the viewing zone. That is, the distance the viewer's head can move before the image flips.

g - the gap ratio of the display: the ratio of the width of a gap between adjacent barriers to the width of a barrier. A ratio of 0 corresponds to a completely blocked display and a ratio of 1 corresponds to a half-blocked display. This directly affects the average brightness.

s - the separation between the barrier plane and the display plane.

p - the width of a pixel on the display. It is interesting to see what happens as this, in the limit, goes to zero.

Observe that **s**, **d**, and **w** determine the width of the barriers. The viewing zone determines a strip of the display plane that is visible behind each gap between barriers. By similar triangles, the width of each barrier is **w*s/d**, and the width of each gap is **g*w*s/d**. Figure 8 also shows a single pixel and its corresponding virtual volume of space. This width can be broken into two components. By similar triangles,

pixel width at virtual depth **z = |p+z*p/s| +|z/s***gap**| = |p+z*p/s| +|z*g*w/d|**

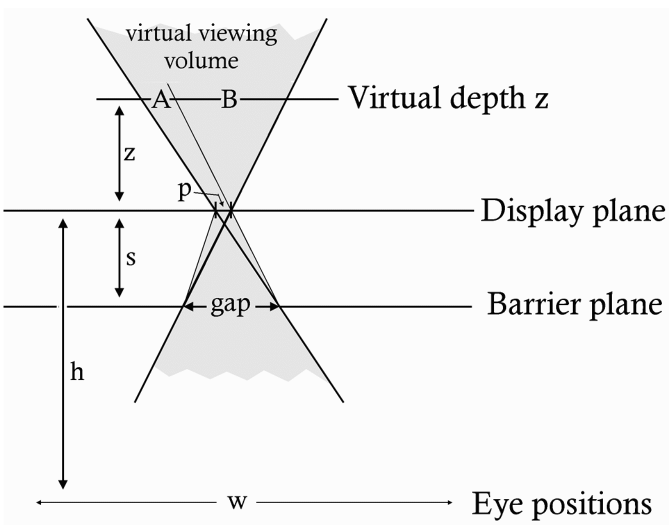

Figure 8. Analysis of viewing zone width at depth z.

As **s** grows, the pixel width decreases, but is bounded from below. Thus, for small **s**, the separation between the planes is in fact the limiting term in the above expression. The excess, the width above **p**, consists of two terms, **|z*p/s|** involving the separation **s** and **|z*g*w/d|** independent of **s** and **p**. Thus the second term is a limiting depth factor based only on **g** and **w/d**. Once **p/s = g*w/d,** the two terms are equal, and the separation is accounting for half of the excess.

As a convenient starting point, we choose **g = 1/6, w/d = 1/2, p = 1/4** mm In this case, a separation of 3 mm means that the separation is accounting for half of the excess pixel width, while a separation of 57 mm means that the separation accounts for five percent of the excess pixel width. In our example, the excess is always at least **|z/12|**, regardless of separation or pixel width. That is to say, a point at depth **z** mm is at least **z/12** mm wide. Thus for an effective pixel width of 2 mm, **z** can be no more than 24 mm in front or behind the display plane -- a total of 48 mm (about 2 inches) deep. This maximum depth is increased as the acceptable pixel width increases, as the brightness decreases, as the head restriction becomes more severe, or as the viewer moves further back. Of course, this last option is the easiest to implement, but the further the viewer is from the display, the less is the importance of stereo.

7. DISCUSSION AND CONCLUSIONS

We have suggested a new autostereoscopic display with motion parallax, which combines time-multiplexing with PPs. The temporal dimension, as with any aspect of 3D-displays, involves costs and tradeoffs. However, we believe that this particular combination of time-multiplexing with PPs is particularly pleasing to the human visual system, based on the examples of interlacing and driving by a picket fence. The biggest drawbacks of this approach are diminished brightness and fast refresh requirements. The computational cost of rendering small-barrier PPs is small and we have presented a more efficient algorithm for rendering large-barrier PPs. To avoid large jumps in the image as the user moves her head a small amount, we must confine the image to a small amount of depth, perhaps about fifty centimeters. While this does not allow for the life-size portrayal of real-word objects, it is more than sufficient to be useful[7] and less likely to cause eyestrain.

8. ACKNOWLEDGEMENTS

We would like to thank Shojiro Nagata for helpful discussions, and Dennis Cosgrove, Randy Pausch, and the Stage 3 research group for their help.

9. REFERENCES

1. Eichenlaub, Jesse, Hollands, Dan, and Hutchkins, Jamie, “A Prototype Flat Panel Hologram-Like Display That Produces Multiple Perspective Views at Full Resolution,” Stereoscopic Displays and Virtual Reality Systems II, SPIE Proceedings 2409, 1995, 102-112.
2. Dodgson, N. A., and Moore, J. R., “Design and Implementation of an Autostereoscopic Camera System,” Technical Report No. 398, University of Cambridge Computer Laboratory, June, 1996.
3. Okoshi, Takanori, “Three-Dimensional Imaging Techniques,” Academic Press, 1976.
4. O’Neill, Graeme, “Into the Third Dimension”, Time Magazine (Australia), October 18th, 1993.
5. [Isono, H., Yasudah, M., Takemori, D., Kanayama, C., Yamada, C., and Chiba, K. 50-inch autostereoscopic full-color 3-D TV display system. Proc. SPIE, 1457:318-321, 1991.
6. Love, Shaun, “3D Hardcopy”, in Stereo Computer Graphics, Editor: McAllister, Princeton, Princeton University Press, 1993.
7. McKee, S. P. 1983. The spatial requirements for fine stereoacuity. Vision Res 23:191-98.
8. Tilton, Homer, “Moving-Slit Methods”, in Stereo Computer Graphics, Editor: McAllister, Princeton, Princeton University Press, 1993.
9. Valyus, N. A., Stereoscopy. London, New York: Focal Press, 1962.
10. Ware, C., Arthur, K, and Booth, Kellog, “Fish Tank Virtual Reality”, InterCHI, 1993, 37-42.