Doing it with Mirrors: Low Budget Stereo Graphics*

Allen Van Gelder and Jane Wilhelms

UCSC-CRL-93-22

Baskin Center for Computer Engineering and Information Sciences
University of California, Santa Cruz 95064

Jan. 7, 1994 (rev.)

Abstract

Interactive stereoscopic images can be viewed on a graphics workstation by producing side-by-side images and viewing through a simple mirror device. However, it is important that the viewing device have pairs of adjustable nonparallel mirrors so large windows can be viewed without the human's sightlines needing to diverge, or "look wall-eyed". Transformations to produce the correct images for this viewing method are described. Previous work applied to the case where both left and right images were to be superimposed and multiplexed in the same region of the screen, often called anaglyphs. Such cases are adequately handled by a translation and an off-axis perspective transformation. The same kind of transformation can be used with a parallel-mirror device, but such devices have practical limitations. This paper shows that nonparallel mirrors require a somewhat more complicated transformation involving scene rotations as well. Derivation of the correct angle of rotation is the main difficulty in computing this transformation. The transformation can be implemented by a sequence of graphics library procedures. Advantages and disadvantages of nonparallel mirror methods are discussed.

*Presented at *SPIE Conference on Visual Data Exploration and Analysis*, Vol. 1278, San Jose, CA, Feb. 1994.

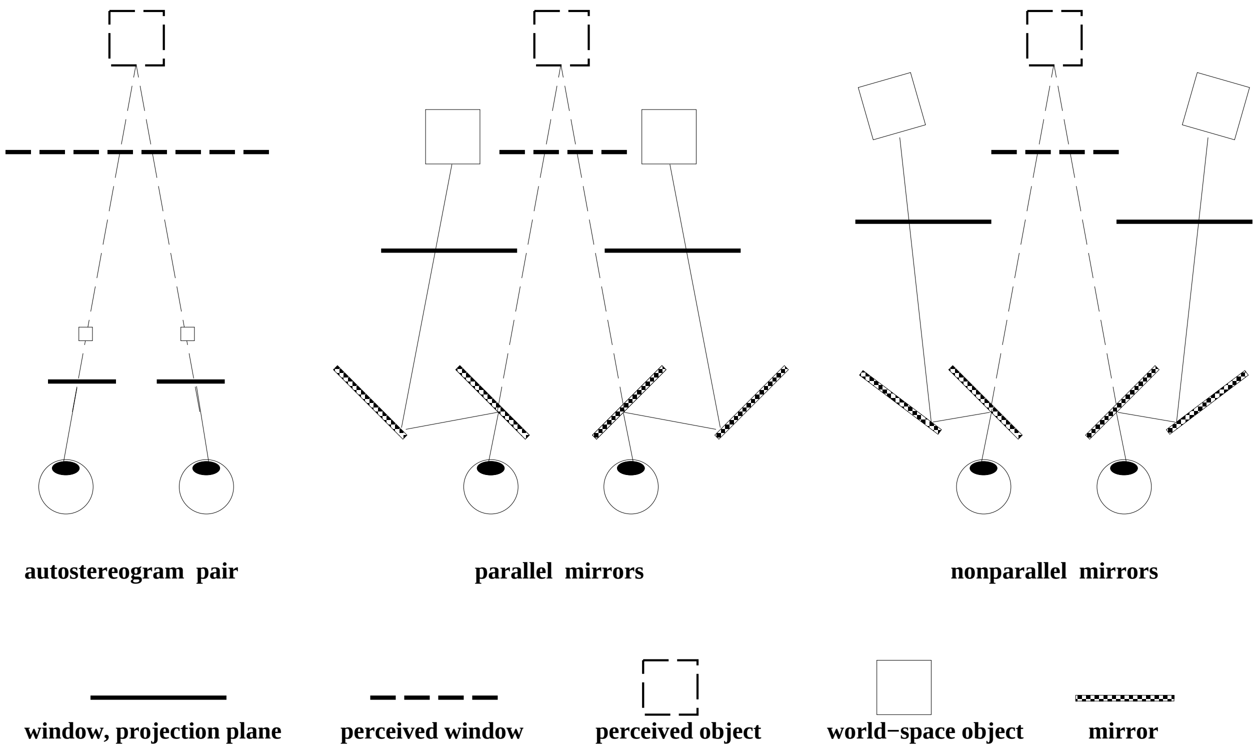

Figure 1: Three options for stereoscopic viewing of side-by-side images: autostereograms, parallel mirrors; and nonparallel mirrors.

1 Introduction

Visualization of three-dimensional scenes on graphics workstation monitors is greatly enhanced in many cases by the use of stereoscopic images. Accurate interpretation of such scenes is vital in visualization of scientific data, where abstract shapes lack many of the 3-dimensional clues present in more realistic scenes.

Two common methods of stereoscopy are based on superimposing the left and right images, but allowing each eye to see only one image by use of special equipment. Such images are often called anaglyphs. In one technique the images are orthogonally polarized, and in the other certain red and green colors are used [HM85, RA90, FDFH90, Tes90, Hod92]. A method based on side-by-side images is less commonly used, but has certain advantages. The latter method is studied in this paper.

Autostereograms are pairs of images that are intended for viewing side-by-side with little or no special equipment. The images are arranged so each eye sees only one, and they must be close enough together so the eyes can be directed straight ahead or at a positive convergent angle when looking at corresponding points in the two images. (See Figure 1.) This imposes a strict size limit on image size: the width cannot exceed the pupil-to-pupil distance of the eyes, about 6–7 cm. For example, 35 mm. stereo slides can be viewed very effectively in this manner. However, this size limit makes the method impractical for stereoscopic displays on graphics monitors. Indeed, even the images in Figure 3 are too far apart to be viewed without equipment.

Older stereo viewing instruments consisted of a system with two parallel mirrors for each eye (Figure 1).

These parallel mirrors increase the effective eye separation by translating the lines of sight away from each other without altering their directions. This allows larger images to be viewed without "looking wall-eyed". Rogers and Adams show a photograph of such a bench-mounted device that appears suitable for 15 cm. photographs or drawings [RA90]. The outer mirrors must be about 25 cm. apart to permit the full scenes of the two images to be viewed realistically. The size of this device pretty much requires it to be on a fixed mounting, but it could be adapted for graphics monitors.

A newer device uses pairs of adjustable nonparallel mirrors to reduce the separation of the outer mirrors without requiring the images to be closer to each other. As shown schematically in Figure 1, the nonparallel mirrors effectively bend as well as translate the lines of sight away from each other. Even though the outer mirrors are approximately 12 cm. apart (all distances are center-to-center unless specified otherwise), the images may be 20 cm. apart, for example. This device, which looks a little like binoculars, is conveniently hand held for viewing side-by-side images on a standard graphics workstation monitor. Side-by-side images may be up to 17 cm. each on such monitors.

Because a pair of nonparallel mirrors bends, or rotates, the sightlines, the usual off-axis stereographic projections [RA90, Tes90, Wat92, Hod92] are not accurate for this situation. This inaccuracy matters in practice. We experienced noticeable difficulties in fusing the two images of a polygonal scene using these known projections, whose mathematical justification is based on parallel sight lines (assuming the eyes are converging at infinity). The parallel-sightline assumption is appropriate for anaglyphic images, which will be superimposed and possibly multiplexed. It can also be correct for autostereograms and parallel-mirror devices.

Experience and earlier experiments have shown that incorrect projections make it difficult for the eyes (brain, actually) to fuse a pair of stereo images. Hodges and McAllister analytically studied artifacts due to the use of a simplified, but slightly incorrect, projection method, the so-called "rotation method" [HM90]. This study was based on the geometry for parallel sight lines (assuming the eyes are converging at infinity), which is applicable to anaglyphic images. Their findings confirmed earlier experimental studies. (See cited paper for further bibliography.)

This paper analyses the projection transformations needed for bent-sightline viewing with nonparallel mirrors. The principal difference from the parallel-sightline case is that a scene rotation *is* needed to compensate for the fact that the sightlines to infinity are no longer perpendicular to the plane of projection. Calculation of the correct rotation angle is somewhat involved. The geometry is studied in Section 2. The correct projection can be obtained with standard graphics library (GL) procedures, making the method simple to implement.

2 Viewing Geometry for Nonparallel Mirrors

This section describes the geometry and transformations required for stereoscopic viewing with nonparallel mirrors. We assume a right-handed coordinate system with both eyes looking in the negative Z direction. The key to finding the correct transformation is to assume the eyes are focused at infinity so their sightlines are parallel to the Z axis. If the projected scene for each eye is consistent with this assumption, then slightly crossing the eyes, as in normal vision, will permit objects at finite perceived distances to fuse properly.

Figure 2 shows the geometry and indicates the viewing parameters. The direction toward the screen is negative Z, to the right is positive X, and positive Y "comes out of the paper". Each eye, looking straight ahead, looks into an *inner mirror* set at 45°, which effects a 90° Y-rotation on the sightline. This sightline enters an *outer mirror*, which (in general) is somewhat off from 45°. After the Y-rotation effected by the outer mirror, the sightline proceeds to the screen at an (in general) oblique angle.

We denote by θ the angle between the twice-reflected sightline and the negative Z axis. As usual, a positive angle denotes a counterclockwise rotation from the axis to the sightline. The commonest case with this viewing method is that θ is positive for the left eye and negative for the right eye. We assume that θ has the same magnitude for both eyes, in other words, that the outer mirrors are adjusted simultaneously by a single control.

As the figure shows, we can back up the sightline to find a virtual eye location: the position where the eye would be if it were looking along the final sightline without mirrors. The goal of the projection transformation is to simulate the 3-dimensional scene on the 2-dimensional window in the plane of the monitor screen.

The scene should appear to be centered at a specified distance d directly in front of a point between the two eyes; the separation of the eyes is denoted as $2e$. The apparent orientation of the scene should be that its negative Z axis is parallel to the final sightline, its positive X axis is horizontally to the right, and the Y axis is vertically up. We further specify that the center of the scene should appear to be in the center of the window.

To be compatible with the foregoing paragraph, viewing transformations should place the scene in screen space as described. These transformations should be designed on the assumption that the scene will be viewed from eyes located at $(\pm e, 0, d)$. Further, the visible part of the scene will *approximately* be the part that projects into a window of width w and height h in the plane $Z = 0$.

To complete the geometric specification, we denote certain distances (see Figure 2), as follows:

1. S_W: center-to-center separation of left and right windows.

2. S_M: center-to-center separation of outer mirrors.

3. a: distance from outer mirror to the screen (actual projection plane).

Recall that we specified that the center of the scene (origin, in screen space) should appear in the perceived center of the window on the screen. Due to the oblique sightline, a perspective projection to the virtual eye position must map the origin into a position on the actual window that is offset in X by an amount δ from the actual window center. (See Figure 2, where δ is indicated for the right window, and is negative there; it is the same sign as θ, in general.) This value is analyzed in the appendix. A closed form solution may not exist, but for reasonable viewing parameters, the following approximation is acceptable:

$$\delta = \frac{w^2 \tan\theta}{4d}$$

Let us analyze the constraint on θ. The final sightline of the left eye therefore intersects the screen (Q in the figure) at a distance of $(e/\cos\theta) - \delta$ to the left of the window center. The total horizontal displacement of the point on the screen directly in front of the left outer mirror (P in the figure) from Q, the point where the sightline intersects, is therefore $\frac{1}{2}(S_W - S_M) + (e/\cos\theta) - \delta$. It follows immediately that

$$a \tan\theta = \tfrac{1}{2}(S_W - S_M) + \frac{e}{\cos\theta} - \frac{w^2 \tan\theta}{4d}$$

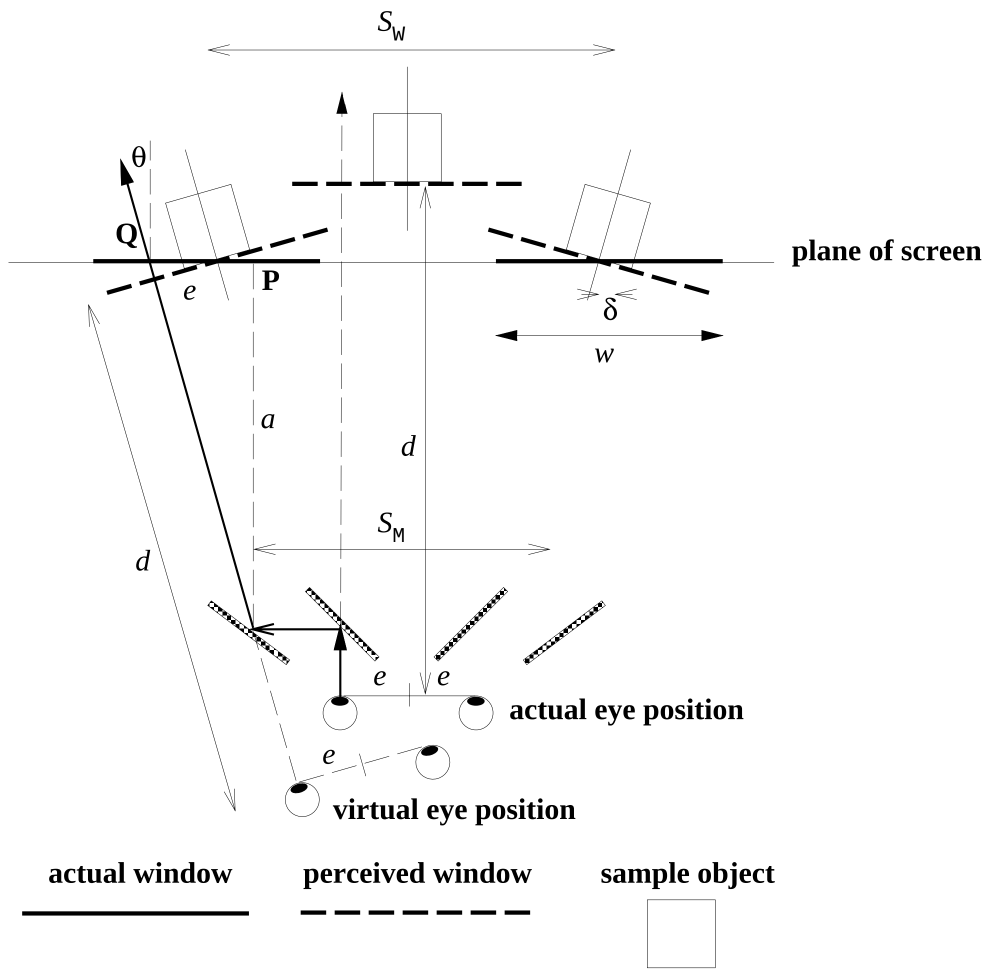

Figure 2: Geometry of nonparallel mirrors, discussed in Section 2. Sample object is a cube of side e with corners at $Z = 0$ and $Z = -e$, and $X = \pm e/2$ in screen space. Eyes focused on infinity should look past the object on either side.

To manipulate this formula more easily we introduce two nondimensional variables:

$$c = \frac{(S_W - S_M)}{2(a + w^2/4d)} \qquad \rho = \frac{e}{(a + w^2/4d)}$$

Observe that all the quantities on the right-hand sides can be expressed in "real world" units, such as centimeters. However, d, e and w will later be used in "screen space" units for transformations; the two representations must not be confused during implementation.

Multiplying through by $\cos\theta$ and substituting the new variables yields:

$$\sin\theta = c\cos\theta + \rho$$

A value for $\sin\theta$ may be found by solving the quadratic equation:

$$(\sin\theta - \rho)^2 = c^2(1 - \sin^2\theta)$$

The sign of the radical is chosen to give a positive root when c is positive:

$$\sin\theta = \frac{\rho + c\sqrt{1 + c^2 - \rho^2}}{1 + c^2}$$

Although it is possible to substitute values such that a solution does not exist, such values would represent an unrealistic viewing situation.

3 Viewing and Projection Transformations

Once the values of θ and δ are determined as described in the preceding section, the following sequence of transformations produces the left eye image. The modifications for the right eye are also mentioned.

1. Apply appropriate viewing transformations to bring the center of the scene to the origin, assuming it will be viewed from $(\pm e, 0, d)$. (These transformations are identical for left and right eyes.)
2. Translate by $(e, 0, -d)$. (Use $(-e, 0 - d)$ for the right eye.) The virtual eye is now at the origin.
3. Rotate by θ around the Y axis. (Rotate by $-\theta$ for the right eye.)
4. The key step is the following perspective transformation. Compute

   $$\begin{aligned} z' &= -d\cos\theta - e\sin\theta \\ x' &= -d\sin\theta + e\cos\theta \end{aligned}$$

   (For right eye, reverse signs of e and θ. Thus z' is the same and x' reverses sign.) Note that the origin of screen space as of the end of step 1 is now at $(x', 0, z')$.

   Project with perspective onto the plane $Z = z'$ (assuming the eye is at the origin, as usual), placing the left window in the plane $Z = z'$, with left edge at $(x' - \delta - \frac{1}{2}w)$ and right edge at $(x' - \delta + \frac{1}{2}w)$. (For the right window, reverse the sign of δ and use the x' computed for the right eye.) The top and botton edges are $\pm\frac{1}{2}h$.

The correctness of these transformations is evident from Figure 2.

The standard off-axis stereographic projection method (suitable for anaglyphic images) can be obtained as a special case of the above method: simply set $S_M = 2e$ and $S_W = 0$. (This fact also confirms the correctness of the calculations.)

Also, for parallel-mirror devices, set S_M to its natural value and set $S_W = S_M - 2e$. This causes θ to come out 0, and places the left and right window borders in consistent positions for both eyes. Notice that the constraint $S_W \geq w$ imposes a limit on window size: $w \leq S_M - 2e$.

Figure 3 shows the difference in images on a test scene suggested by Rogers and Adams [RA90]. The black image was computed as described here. The gray image was computed with the "standard" off-axis projection stereo transformation. The purpose of the picture is to illustrate the magnitude of the discrepancies, not to demonstrate quality. These discrepancies are of the same order of magnitude as those between the "standard" method and the "rotation" method.

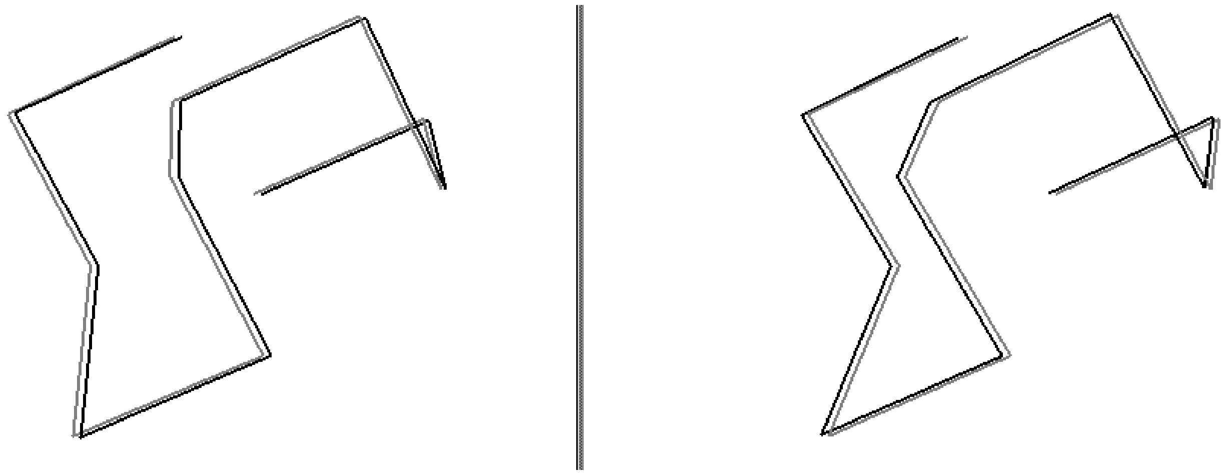

Figure 3: Comparison of transformations on a test image; *not* intended for autostereo viewing. Black is the transformation described here; gray is the "standard" stereo transformation.

3.1 Using the Graphics Library

The foregoing transformations can be accomplished with standard graphics library (GL) procedures, keeping in mind that the procedures are called in the reverse order of the sequence in which transformations are to be applied. These GL calls generalize the method of Tessman [Tes90] by providing for side-by-side images and by providing for the near clipping plane to be closer than the monitor screen.

The only transformation requiring special attention is step 4. The simplest way to accomplish this would be

$$\texttt{window}(x' - \delta - \tfrac{1}{2}w,\ x' - \delta + \tfrac{1}{2}w,\ -\tfrac{1}{2}h,\ \tfrac{1}{2}h,\ -z',\ -1000 * z')$$

The first point to notice is that the last two parameters must have their signs reversed, as they are interpreted as *distances* of near and far clipping planes, not as the Z coordinates of those planes.

A more subtle problem is that the scene will be clipped in the plane of the screen by the specification of $-z'$. Thus objects whose natural 3-dimensional position would be closer than the screen will be clipped, and actually, the clipping will be inconsistent between the two eyes when $\theta \neq 0$.

To avoid undesirable near clipping, we simply observe that the perspective transformation is invariant under scaling of all parameters of `window`. That is, unclipped objects' images are not altered by scaling all parameters. Therefore we simply compute a positive r, say 0.10, and actually call:

$$\texttt{window}(r * (x' - \delta - \tfrac{1}{2}w),\ r * (x' - \delta + \tfrac{1}{2}w),\ r * (-\tfrac{1}{2}h),\ r * (\tfrac{1}{2}h),\ r * (-z'),\ -1000 * z')$$

to move the near clipping plane 9/10 of the way from the screen to the eyes, while leaving the images unchanged.

To reiterate, the sequence of GL calls for each eye would be `loadmatrix(identity)`, `window(...)`, `rotate(`θ`, 'y')`, `translate(`e`, 0, `$-d$`)`. This matrix should be saved (by `pushmatrix()`) and restored (by `popmatrix()`) before and after scene-specific viewing transformations.

3.2 Projection Matrix

Using standard methods, the matrices representing the transformations of steps 2–3 (left) and of step 4 (right) are shown below.

$$\begin{bmatrix} \cos\theta & 0 & -\sin\theta & 0 \\ 0 & 1 & 0 & 0 \\ \sin\theta & 0 & \cos\theta & 0 \\ x' & 0 & z' & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ \dfrac{(\delta - x')}{z'} & 0 & 1 & \dfrac{1}{z'} \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{1}$$

where x' and z' are as given earlier in the description of step 4. (Recall that z' is negative.)

For ray-casting purposes, the above matrix can be used to determine the directions in which to cast rays. With the projection given, the window is bounded by $\pm\frac{1}{2}w$ in X and $\pm\frac{1}{2}h$ in Y, and in the plane $Z = z'$.

4 Advantages and Disadvantages of Nonparallel Mirrors

One advantage of methods based on nonparallel mirrors is that such devices are commercially available for under \$200, and can be used on a variety of media, including still photographs and VCRs, as well as graphics monitors.

Side-by-side stereo image pairs can take full advantage of the monitor's natural color quality and resolution. Also, the stereo image pair can share the screen with other windows in the normal manner, and one image can be viewed without equipment. The viewing environment is generally natural and comfortable. In contrast, the red-green filter method greatly limits the range of colors. The polarizing filter method delivers only half of the normal vertical resolution. (See Hodges and McAllister [HM85] or Hodges [Hod92] for discussion of the trade-off between vertical resolution and refresh rate for polarizing filters.) Currently, the entire monitor must be put in a special mode to display multiplexed images. Viewing with special glasses, whether color or polarized, is generally regarded as slightly stressful and uncomfortable.

A disadvantage of the side-by-side method is that greater distortions arise when the viewing position is off to one side. This is due to one window being farther from the eyes than the other. Variations of 20° off center become uncomfortable. Although distortions arise in anaglyphic methods due to incorrect viewpoints, they seem to be somewhat smaller and more tolerable.

Another source of confusion is tipping the head slightly to one side. With mirrors, the vertical alignment goes out and fusion breaks down quickly. With polarized glasses, slight amounts of the wrong image creep in. Color glasses are impervious to this error.

Also, when $\theta \neq 0$, clipping in the near projection plane is inconsistent between the two images. Scenes should be planned so that such clipping does not occur, or is not an important part of the effect. However, if one wanted to clip on a surface at an exact value to see the contour, it would have to be programmed into the rendering; use of the near clipping plane would not suffice.

Another disadvantage is that the field of view is limited, making it difficult to watch the picture while adjusting it via the mouse or keyboard.

5 Conclusion

Nonparallel mirror devices provide a low cost and flexible method to use stereoscopy in 3-dimensional visualization. However, to get the most accurate images requires acknowledgement of the fact that the effective sightline is normally oblique to the actual projection plane, which is the monitor screen. A good approximation to the exact transformation can be computed in closed form (provided square roots, trig functions, and inverse sin are available). Iteration could be used for greater accuracy, but appears unnecessary. The projection transformation can be implemented with standard graphics library procedures, including `window()`.

Further work is needed to study the sensitivity of the method to incorrect assumptions about viewing parameters, and errors in viewer position.

Also, experience is needed to determine whether perspective transformations are appropriate for visualization of scientific data. Do the distortions of scale give misleading readings? There is a case for this attitude in monocular scenes, but it needs to be re-evaluated in the context of stereo.

Acknowledgements

This research was supported in part by NSF PYI grant CCR-8958590, NSF New Technologies Program grant ASC-9102497, NSF Infrastructure grant CDA-9115268, and a NASA-Ames Research Center Cooperative Agreement Interchange No. NCC2-717.

Appendix A Calculation of Center Offset

The calculation of δ (see Figure 4), the X-offset from the center of the actual window to the projection of the center of the scene (i.e., screen space origin) is derived here. We find that the closed form exists, but is too complicated a function of θ to permit θ to be found in closed form. However, a suitable approximation for realistic viewing parameters can be found, and this is a simple enough function of θ.

The geometry is shown in more detail in Figure 4. Let v denote the width of the perceived window, which is defined to be in a plane perpendicular to the sightline. As shown, the sightline intersects this plane at a negative offset of e. The angles α and β are from the sightline to the left and right borders of the (perceived or actual) window, respectively. We have

$$\tan\alpha = \frac{\frac{1}{2}v - e}{d} \qquad \tan\beta = \frac{\frac{1}{2}v + e}{d}$$

Further, by the law of sines from plane geometry, together with the identity $\sin(\frac{1}{2}\pi \pm \psi) = \cos\psi$, we obtain

$$\frac{\sin\alpha}{\frac{1}{2}w + \delta - e/\cos\theta} = \frac{\cos(\alpha+\theta)}{d + e\tan\theta} \qquad \frac{\sin\beta}{\frac{1}{2}w - \delta + e/\cos\theta} = \frac{\cos(\beta-\theta)}{d + e\tan\theta}$$

Manipulating the latter two equations leads to

$$\frac{v(d + e\tan\theta)}{d\cos\theta(1 + \tan\beta\tan\theta)(1 - \tan\alpha\tan\theta)} = w$$

as well as

$$2\delta - \frac{2e}{\cos\theta} = \frac{(d + e\tan\theta)(2\tan\alpha\tan\beta\tan\theta - \tan\beta + \tan\alpha)}{\cos\theta(1 + \tan\beta\tan\theta)(1 - \tan\alpha\tan\theta)}$$

which eventually reduces to

$$\delta = \frac{vw\tan\theta}{4(d + e\tan\theta)}$$

Expressions for $\tan\alpha$ and $\tan\beta$ in terms of v, d and e were used to get this equation.

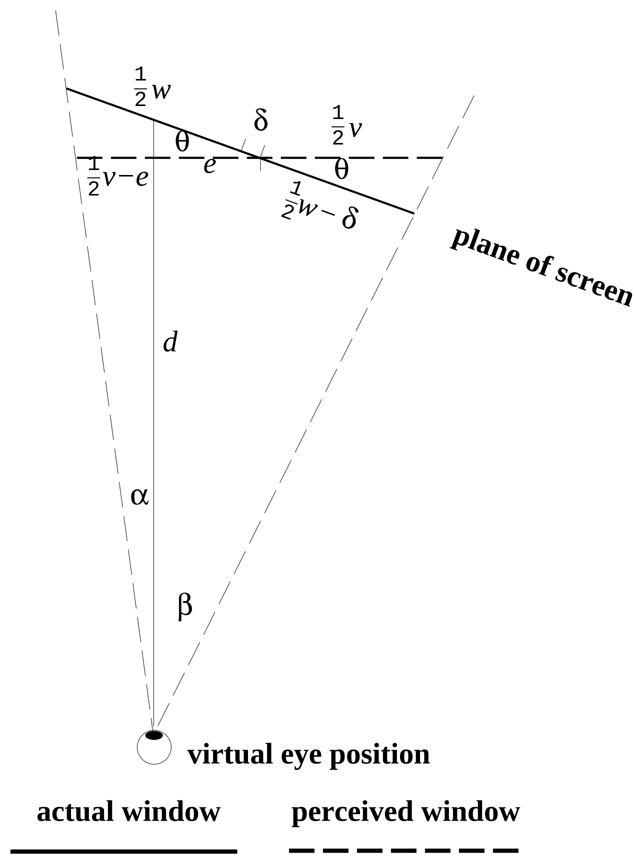

Figure 4: Geometry for deriving offset δ. The "perceived window" is considered to be orthogonal to the sightline, and has width v. Its center maps to a point that is offset by δ from the center of the actual window, which has width w.

To complete the solution for δ, v must be expressed in terms of the parameters θ, d, e, and w. An exact solution is given below.

However the exact solution for δ is intractable for the eventual goal of finding θ in closed form. Instead we use an approximation that ignores second order terms in $\tan\theta$, $\tan\alpha$, and $\tan\beta$.

$$v = w\left(1 + \left(\frac{e}{d}\right)\tan\theta\right)$$

This gives the tractable approximation used in Section 2:

$$\delta = \frac{w^2 \tan\theta}{4d}$$

For completeness, we give the exact solution for v. Noting that $\tan\alpha$ and $\tan\beta$ are linear functions of v, we obtain the quadratic equation

$$v^2\left[\frac{w\tan^2\theta\cos\theta}{4d^2}\right] + v\left[1 + \frac{e}{d}\tan\theta\right] - \left[w\cos\theta\left(1 + \frac{e}{d}\tan\theta\right)^2\right] = 0$$

The positive root is

$$v = w\left(1 + \frac{e}{d}\tan\theta\right)\left[\frac{\sqrt{1 + \left(\frac{w\tan\theta\cos\theta}{d}\right)^2} - 1}{\left(\frac{w^2\tan^2\theta\cos\theta}{2d^2}\right)}\right]$$

The expression in square brackets, for small $\tan\theta$, has the expansion

$$\left[1 - \left(\frac{1}{2} + \frac{w^2}{4d^2}\right)\tan^2\theta + \ldots\right]$$

which justifies approximating it by 1.

References

[FDFH90] James D. Foley, Andies Van Dam, Steven Feiner, and John Hughes. *Computer Graphics: Principles and Practice*. Addison-Wesley Publishing Company, Reading, Mass., 2 edition, 1990.

[HM85] Larry F. Hodges and David F. McAllister. Stereo and alternating-pair techniques for display of computer-generated images. *IEEE Computer Graphics and Applications*, 5(9):38–45, 1985.

[HM90] Larry F. Hodges and David F. McAllister. Rotation algorithm artifacts in stereoscopic images. *Optical Engineering*, 29(8):973–976, 1990.

[Hod92] Larry F. Hodges. Time-multiplexed stereoscopic computer graphics. *IEEE Computer Graphics and Applications*, 12(2):20–30, 1992.

[RA90] David F. Rogers and J. Alan Adams. *Mathematical Elements for Computer Graphics*. McGraw-Hill, New York, 2 edition, 1990.

[Tes90] T. Tessman. Perspectives on stereo. In J. O. Merritt and S. S. Fisher, editors, *Proceedings of SPIE 1256: Stereoscopic Displays and Applications*, pages 22–27, 1990.

[Wat92] Val Watson. Stereo simplified. Nasa-Ames Research Center, Moffett Field, CA, 1992.