Stereographics

Mike Bailey

mjb@cs.oregonstate.edu

Oregon State University

Stereovision is not new –
It's been in common use since the 1950s,
and sporadically even before that

Life Magazine

Binocular Vision

In everyday living, part of our perception of depth comes from the slight difference in how our two eyes see the world around us. This is known as *binocular vision.* We care about this, and are discussing it, because stereo computer graphics can be a great help in de-cluttering a complex 3D visualization scene.


The Cyclops Model

In the world of computer graphics, the two eye views can be reconstructed using standard projection mathematics. The simplest approach is the *Cyclops Model.* In this model, the left and right eye views are obtained by rotating the scene plus and minus what a Cyclops at the origin would see. Looking at everything from the top:

The left eye view is obtained by rotating the scene an angle +Ø about the Y axis. The right eye view is obtained by rotating the scene an angle -Ø about the Y axis. In practice, a good value of Ø is 1-4˚.

The Vertical Parallax Problem

This seems too simple, and in fact, it is. This works OK if you are doing orthographic projections, but if you use perspective, you will achieve a nasty phenomenon called *vertical parallax*, as illustrated below:

Same face seen from different eye positions

The fact that the perspective shortening causes points A and B to have different vertical positions in the left and right eye views makes it very difficult for the eyes to converge the two images. For perspective projections, we need a better way.


The Vertical Parallax Problem

Why not just keep using orthographic projections? Mathematically this is fine, but in practice, the two depth cues, stereo and no-perspective, fight each other. This will bring on an optical illusion. A good example of this is a simple cube, drawn below using an orthographic projection:

Because of the use of stereographics, the binocular cues will say that the Near face is closer to the viewer than is the Far face. However, our visual experience says that the only way a far object can appear the same size as a near object is if it is, in fact, larger. Thus, your visual system will perceive the Far face as being larger than the Near face, when in fact they are the same size.

Diversion #1 – Specifying the Viewing Frustum

The OpenGL glFrustum call can be used in place of gluPerspective:

glFrustum(left, right, bottom, top, near, far);

Rather than having to specify the left, right, bottom, and top limits at the near clipping plane (which is what glFrustum expects), let's setup a way to specify those limits at a particular distance in front of us. (This is derived using similar triangles.)

```
void
FrustumZ( float left, float right, float bottom, float top, float znear, float zfar,   float zproj )
{
        if( zproj != 0.0 )
        {
                left     *= ( znear/zproj );
                right    *= ( znear/zproj );
                bottom *= ( znear/zproj );
                top      *= ( znear/zproj );
        }

        glFrustum( left, right, bottom, top, znear, zfar );
}
```

So, if you wanted to view a car from 30 feet away, you could say:

FrustumZ(-10., 10.,   -10., 10.,   .1, 100.,   30.);


Diversion #2 – Where does a 3D Point Map in a 2D Window?

Take an arbitrary 3D point in the viewing volume. Place a plane parallel to the near and far clipping planes at its Z value (i.e., depth in the frustum). The location of the point on that plane shows proportionally where the 3D point will be perspective-mapped from left to right in the 2D window.

Two Side-by-side Perspective Viewing Volumes

The best stereographics work is done with perspective projections. To avoid the vertical parallax problem, we keep both the left and right eyes looking straight ahead so that, in the vertical parallax example shown before, points A and B will project with exactly the same amount of shortening.

The left eye sees the box towards the far right side of its display

The right eye sees the box towards the far left side of its display

The left eye view is obtained by translating the eye by -E in the X direction, which is actually accomplished by translating the scene by +E instead. Similarly, the right eye view is obtained by translating the scene by -E in the X direction. We now have a *horizontal parallax* situation, where the same point projects to a different horizontal position in the left and right eye views.

Note that this is a *situation*, not a *problem*. The difference in the left and right eye views requires a at least some horizontal parallax to work. You can convince yourself of this by alternately opening and closing your left and right eyes. We just need a good way to *control* the horizontal parallax.


Two Side-by-side Perspective Viewing Volumes

We do this by defining a distance in front of the eye, z0p, to the *plane of zero parallax*, where a 3D point projects to the same window location for each eye. To the viewer, the plane of zero parallax will be the glass screen and objects in front of it will appear to live in the air in front of the glass screen and objects behind this plane will appear to live inside the monitor. The plane of zero parallax is handled by:

1. Set the distance from the eyes to the plane of zero parallax based on the location of the geometry and the look you are trying to achieve.

2. Looking from the Cyclops eye at the origin, determine the left, right, bottom, and top boundaries of the viewing window on the plane of zero parallax as would be used in a call to glFrustum( ). These can be determined by knowing Z0p and the field-of-view angle Φ:

Two Side-by-side Perspective Viewing Volumes

Cyclops eye:

$$L0p = -Z0p * tan(ø/2)$$
$$R0p = Z0p * tan(ø/2)$$
$$B0p = -Z0p * tan(ø/2)$$
$$T0p = Z0p * tan(ø/2)$$


Two Side-by-side *Non-symmetric* Perspective Viewing Volumes

Use the Cyclops's left and right boundaries as the left and right boundaries for each eye, even though the scene has been translated. In the left eye view, the boundaries must then be shifted by +E to match the +E shift in the scene. In the right eye view, the boundaries must be shifted by -E to match the -E shift in the scene.

Left eye:
$$R0p = Z0p * tan(ø/2) + E$$
$$L0p = -Z0p * tan(ø/2) + E$$

Right eye:
$$R0p = Z0p * tan(ø/2) - E$$
$$L0p = -Z0p * tan(ø/2) - E$$

```
void
Stereopersp( float fovy, float aspect, float znear, float zfar, float z0p, float eye )
{
    float left, right;              // x boundaries on z0p
    float bottom, top;             // y boundaries on z0p
    float tanfovy;                 // tangent of y fov angle

    // tangent of the y field-of-view angle:

    tanfovy = tan( fovy * (M_PI / 180.) / 2. );

    // top and bottom boundaries:

    top = z0p * tanfovy;
    bottom = -top;

    // left and right boundaries come from the aspect ratio:

    right = aspect * top;
    left  = aspect * bottom;

    // take eye translation into account:

    left  -= eye;
    right -= eye;

    // ask for a window in terms of the z0p plane:

    FrustumZ( left, right, bottom, top, znear, zfar, z0p );

    // translate the scene opposite the eye translation:

    glTranslatef( -eye, 0.0, 0.0 );

}
```

An Example

Parallel viewing stereo

Left   Right

Cross-eye viewing stereo

Right   Left

mjb – January 27, 2011

7

Acquiring Stereo Video

ESPN's 3D camera

ESPN's 3D camera

Panasonic's 3D Camcorder

Quad-Buffered OpenGL

Remember double buffering, where you draw into the back buffer and display from the front buffer? OpenGL actually has two back buffers and two front buffers, one for each eye. So, draw the left eye view into GL_BACK_LEFT and the right eye view into GL_BACK_RIGHT. First you need to tell GLUT that you are doing stereo graphics. In InitGraphics( ):

glutInitDisplayMode(GLUT_RGBA | GLUT_DOUBLE | GLUT_DEPTH | GLUT_STEREO);

Then go ahead and create the window as normal. After creating the window, you can also expand it to be the full screen with:

glutFullScreen( );

In Display( ), you need to clear both buffers:

glDrawBuffer(GL_BACK_LEFT);
glClear(GL_COLOR_BUFFER_BIT | GL_DEPTH_BUFFER_BIT);

glDrawBuffer(GL_BACK_RIGHT);
glClear(GL_COLOR_BUFFER_BIT | GL_DEPTH_BUFFER_BIT);

Oregon State University
Computer Graphics

Quad-Buffered OpenGL

In Display( ), you also need to draw into both back buffers:

```
for( int eye = 0; eye <= 1; eye++ )
{
        glMatrixMode( GL_PROJECTION );
        glLoadIdentity();

        if( eye == 0 )            // left eye view
        {
                glDrawBuffer( GL_BACK_LEFT );
                Stereopersp( fovy, 1.0,  znear, zfar,  z0p,  - eyesep );
        }
        else                     // right eye view
        {
                glDrawBuffer( GL_BACK_RIGHT );
                Stereopersp( fovy, 1.0,  znear, zfar,  z0p,   eyesep );
        }

        glMatrixMode( GL_MODELVIEW );
        glLoadIdentity( );

   << draw the 3D scene >>
}

glutSwapBuffers( );              // this goes outside the eye loop!
```

Separating the Left and Right-eye Views –
Shutterglasses

Infrared transmitter to synchronize the left-right of the glasses to the left-right of the screen refresh

9

Separating the Left and Right-eye Views –
## Head-mounted Goggles

Separating the Left and Right-eye Views –
the Stereo Mirror

Separating the Left and Right-eye Views –
Dual Projectors ("GeoWall")

Two filters statically provide the polarization

Separating the Left and Right-eye Views –
Stereo Movie Projectors

Carmike Theatres, Corvallis

One filter dynamically provides the polarization

Separating the Left and Right-eye Views – Stereo Movie Projectors

For movies and sporting events

Separating the Left and Right-eye Views

12

**Separating the Left and Right-eye Views –
Left-Right 3DTV**

Shutterglasses

mjb – January 27, 2011

**Separating the Left and Right-eye Views –
Top-Bottom 3DTV**

Shutterglasses

mjb – January 27, 2011

13

Separating the Left and Right-eye Views –
Interlaced 3DTV

Shutterglasses

Separating the Left and Right-eye Views –
Red-Cyan Anaglyphs

Separating the Left and Right-eye Views – ColorCode™

This is what was used at the 2009 Super Bowl halftime!

ColorCode 3·D

OSU **Oregon State University**
Computer Graphics

mjb – January 27, 2011

Encoding Stereo in a Single Image – ChromaDepth™

ChromaDepth picture

Blue virtual image

Red virtual image

OSU **Oregon State University**
Computer Graphics

mjb – January 27, 2011

Encoding Stereo in a Single Image –
ChromaDepth™

Separating the Left and Right-eye Views –
Lenticular

Mash up image from multiple eye position images

Separating the Left and Right-eye Views –
Other Ways

Stereographics Rules of Thumb

• Stereographics is especially good for de-cluttering wireframe displays.

• Use perspective, not orthographic, projections to avoid the optical illusion.

• Use an eye separation, E, of approximately: $E = Z0p * tan(1°-4°)$

• Use the far clipping plane well. The stereo effects are enhanced when the scene is not complicated by a lot of tiny detail that is far away. The interactive response is improved too.

• Because you are drawing the scene twice, using display lists is especially important.

• It is fun to set Z0p = Zfar so that the image appears to be hanging out in the air in front of the monitor. However, in real life we rarely see anything hanging out in the air that has its sides clipped for no apparent reason, as the scene is likely to have. Perceptually, it is often better to set Z0p =   Znear so that the entire scene looks like it is inside the monitor and that you are viewing it through a rectangular hole cut through the glass. This situation is common in everyday life, so we are used to seeing things that way.

• Intensity depth cueing (glFog) nicely enhances the stereo illusion.

• If you are using texture mapping, be sure to use GL_LINEAR, not GL_NEAREST, for the texture filtering.

From what you now know, real stereo images have to be generated from the original data – they cannot be as effectively retro-generated from a mono image, at least not without a lot of work.

Beware the multiple plane effect !

Watch out for a lot of monoscopic movies being "re-released" in stereo !