



Parallel vs Converged *an industry op-ed by Steve Wright*

the debate over stereo camera setups

The recent explosion in popularity of stereo 3D films has triggered a raging debate over the best stereo camera arrangement - parallel or converged (toe-in). The discussion has taken on religious fervor with each stereographer trying to out-stereograph the other to establish himself as the high priest of stereography. What is obviously needed here is some good old-fashioned R & D to nail down standards and practices with facts, math and science so we can get on with the business of producing quality stereo movies.

In fact, some very serious science is being set up even now and will eventually result in established practices we can all agree on, because to disagree would be addlepatated folly. In the meantime, however, the debate rages on, shedding more heat than light. This article does not intend to settle the debate. Its about the yin and the yang, the pro and the con, and merely tries to equip the reader with some understanding of the issues in order to join the fray. Have fun.

Key Stereo Concepts

As we all know, to see a stereo image there needs to be two slightly different views of a scene, one seen by the left eye only and one by the right eye only. If the two views of an object are shifted left or right relative to each other, that shift is called parallax, or disparity. If there is no shift between the two views of an object then it has zero parallax. It is the amount and direction of parallax in a stereo image that places an object in front of, at, or behind the theatre screen.

Zero parallax places an object right on the movie screen. If the left eye's view of an object is shifted left at the screen relative to the right view, then the object is perceived as behind the screen in "screen space". If the left view is shifted to the right the object is perceived as in front of the screen in "theatre space". The movie screen itself is referred to as the "stereo window", referring to objects appearing to be either inside or outside of this "window".

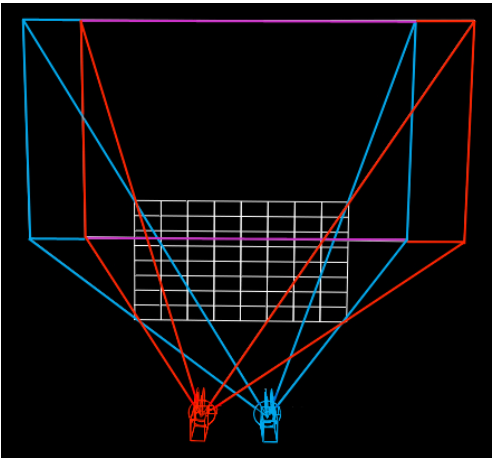
Another key concept in stereography is "fusion" - the brain's ability to fuse the left and right views into a single stereo view. The brain does some very impressive image processing and feature matching to synthesize a stereo image. But to work properly, the two views must have a matching set of features to correlate. No fusion, no stereo.

How much stereo depth a movie exhibits is affected by the distance between the two cameras, called the inter-axial distance. Human eyes are about 2.5 inches apart, so the inter-axial distance is usually between 2 and 3 inches, but can be varied depending on lenses and artistic intent.

Many of the images in this article are red and blue anaglyph images created with Nuke to help tell the story. You will need an attractive pair of anaglyph glasses to see them properly.




Converged Cameras



Converged camera setup

The converged camera setup is shown in the illustration at left. The two cameras are set apart by the inter-axial distance then converged, or angled towards each other (sometimes called a “toe-in”). The region where the two cameras’ fields of view converge, or intersect, defines the stereo window (the white grid) where the parallax is zero. This places objects that are at that distance right on the theatre screen. Everything in front juts out into theatre space, anything behind is pushed back into screen space.

The stereo render below (glasses on) was done using converged cameras with the front head at the point of convergence placing it on the screen plane and everything else behind it in screen space. With anaglyph images like this you can tell when an object has zero parallax by taking your glass off and noting that it has no red or blue outline. This means its left and right views are not shifted, the very definition of zero parallax. If you move your cursor over the front head’s nose (glasses on) it appears to sit right on it. Moving it over other parts of the picture it appears to float over it.




Stereo render with converged cameras

Note the two other heads on the left and right. They clip the edge of the stereo window so the left and right views have some missing regions. With the anaglyph glasses on, close one eye then the other and you can see how one view will be cut off but the other is complete. Based on our rules of fusion, the brain obviously cannot fuse the two views

in the regions of missing information and should create visual distress. However, because the missing information is *behind* the stereo window the brain does not complain.

Why is this so? Because the brain is well adapted to this situation. Just look out a window. Closing your eyes one at a time you will see the window frame jump left and right hiding and revealing information about the outside scene, but with both eyes open the scene looks fine. No complaints from Mr. Brain. The brain is designed to ignore foreground obstructions to the background view so we could spot a tiger in the weeds from behind a tree. However, if the missing information is from objects *in front* of the screen we get a *stereo window violation* and the brain objects. We will see this effect shortly.



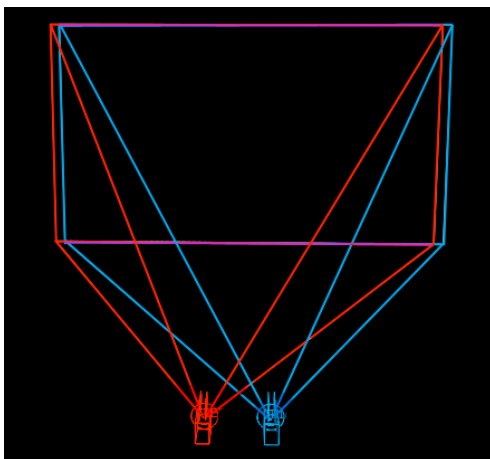
left view right view
Keystone effect from converged cameras

The downside of converged cameras is that they distort the two views with what is called the “keystone effect” illustrated here. If converged stereo cameras were used to photograph a brick wall “flat on” each camera

would be viewing the wall from a slight angle introducing different perspectives into the images which results in keystoning. Depending on the scene content, this may or may not be a problem. The good news is that all feature films are completely digitized nowadays so the keystoning can be corrected digitally.

The converged cameras also “bake in” the convergence point that defines the screen plane. However, this can be adjusted in post by shifting the two views horizontally, but this introduces a third problem. When the images are shifted like this they lose a bit of picture from their left and right edges and need to be scaled up to push the black edges out of frame. To avoid this stereo films are often photographed with an “overscan”, meaning the scene is shot wider than needed to leave room for a modest crop.

Parallel Cameras




Parallel camera setup

With all the problems of converged cameras then the solution must be parallel cameras. Sorry, this just introduces a different set of problems. Some stereographers prefer these problems, others do not. Of course, the real world is not made up of absolutes, so some brave souls actually use parallel cameras with a *slight* toe-in. Iconoclasts!

A parallel camera setup is illustrated at left. The two cameras are separated by the inter-axial distance but both are aimed straight ahead with no convergence or toe-in. The point of convergence becomes essentially infinity. Each camera sees a slightly different view of the

scene so we get our stereo views. There is no keystoneing, and a convergence point is not baked in. This must be the perfect solution, right?


Of course not. When the convergence point is infinity, then infinity is placed at the screen level and everything else thrusts into theatre space. You can see what that looks like in the rendered image below. If you place your cursor on the sky in the far background (glasses on) you will see that it is on the screen plane. Place your cursor over any of the heads or ground plane and they appear to float over the cursor because they are in theatre space.



Stereo render with parallel cameras


The problem with this “everything in theatre space” arrangement can be seen with the two heads at the left and right edges. They have unmatched information for the two views at the screen edge just like we saw for the converged cameras, but this time they are in *front* of the stereo window resulting in a severe stereo window violation (yikes!). Mr. Brain objects to this strenuously and punishes the viewer with eye strain and headaches. Since the foreground head itself does not touch the stereo window it pokes into theatre space without creating a window violation - unlike the two other heads and the ground plane. Note how difficult it is to focus on the bottom or side edges of this image. However, occasional stereo window violations are unavoidable (i.e. the talent exits frame left) and if brief and well-chosen, are considered acceptable.

The fix is to shift the two views horizontally to place the point of convergence wherever desired. The formal term is *Horizontal Image Translation*, or HIT. The HIT fix has been applied to the parallel image rendered below. Unfortunately, this has now introduced red and blue strips on the left and right edges of the picture with missing information from each view. However, (glasses on) the convergence point is now on the head in the foreground, the far horizon is pushed back into deep screen space, and there is no stereo window violation. The results are just as pleasing as the converged stereo image, except for the hideous red and blue strips on the edge.



Stereo parallel render with HIT fix (overscan guide)

To avoid having to scale the image up to push the problem edges out of frame scenes can be shot at high resolution with an overscan (dotted guide illustrated above) so the image can be cropped and resized down to the final show format with no loss of picture. The example below shows the finished shot cropped to the overscan guide.




Stereo render with cropped HIT fix

While this works fine for shows captured in stereo, when doing a stereo conversion on a finished "flat" movie there is no extra margin around the edge of the frame so picture edges may be lost on some scenes.

Parallel with Offset Image Planes

In the perfect world of CG there appeared an ideal solution. Using parallel cameras to avoid keystone, the camera's image planes are offset slightly in opposite directions to set the point of convergence at the item of interest - the foreground head in this case. As a result, in the stereo render below we have the best of both worlds. Observe (glasses on) the deep stereo scene with no stereo window violations since the two heads and the ground plane clip the stereo window in screen space. Rendered with parallel cameras, the camera image planes were offset to set the convergence point to the foreground head like all the previous examples with no need for overscan or cropping.



Parallel stereo renders with image plane offset

The good news is that modern cameras now have adjustable film backs (image planes) that can duplicate the adjustments possible with the CG cameras. The DP or on-set stereo supervisor can now factor in the lens choice, inter-axial distance, film back size, and the distance to the convergence target to calculate the proper overscan so there will be no HIT or cropping needed. Remember when the DP only needed a lens chart and a light meter?

Steve Wright