3D Production and Post

Barry Clark
barryclark@telenova.us
Rev. 04-15-11

3D Production and Post

Barry Clark Rev. 04-15-11

1. Real World 3D

When a viewer's eyes focus on a real object, they automatically converge on the object. From the separate perspectives seen by the two eyes, the viewer fuses a coherent 3D image of the object. All of the objects that the viewer sees in 3D occupy a cone that is bounded by the edges of the overlapping fields of focus and convergence of the viewer's eyes (Fig.1). Everything outside of this cone is seen by the viewer in 2D. As the viewer's eyes focus on progressively more distant objects, the zone of convergence shifts with the zone of focus and the cone shrinks in width until an outer limit of distance is reached—a distance of 100-200 yards in the average adult—beyond which the viewer can no longer distinguish the perspectives seen by the left and right eyes. Everything that is located further away from the viewer seems to lie on a flat, 2D plane. To judge the relative position in space of objects that lie beyond this *stereoscopic limit,* a viewer must rely on *monoscopic depth cues*, including *motion* cues (nearby objects seem to shift position more rapidly than distant objects), *atmospheric* cues (the hue of objects shifts toward blue as they move into the distance), and *occlusion* cues (near objects obscure the view of more distant objects).

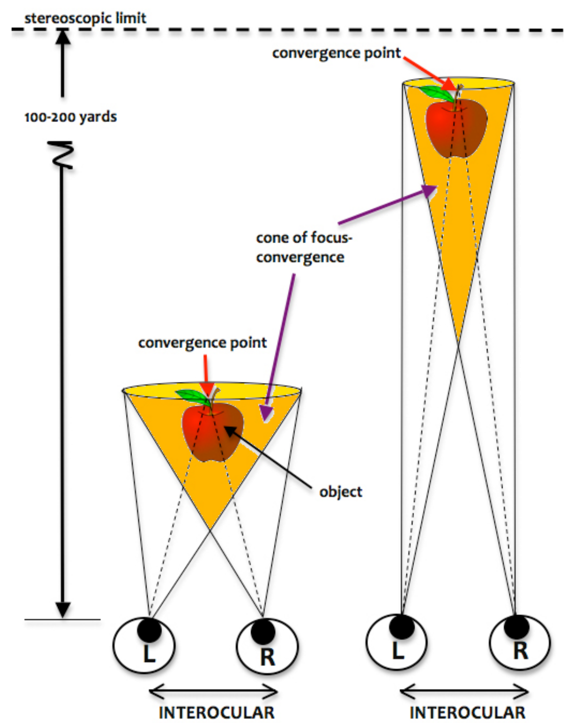

Fig.1 – Real World 3D

2. Simulated 3D

The experience of viewing a 3D film is significantly different from the way we see 3D in the real world. The most obvious differences between real world 3D and the simulated 3D that is viewed on a screen are consequences of the fixed depth-of-field and the fixed point-of-view of the lenses that capture the images. As a result of these constraints, viewers watching simulated 3D can no longer alter their point-of-view simply by shifting the position of their heads, as they can in the real world. And when turning their attention from one object of interest to another, they can no longer simply refocus their eyes, as they can in the real world. In a 3D film, the point-of-view and the focus are invariables established on the set. In addition, when looking at a 3D object displayed on a screen, a viewer's eyes must focus on the screen while, at the same time, they converge on a point in space that may be located *beyond* the screen, *on* the screen, or *in front of* the screen. As a result of this process—which differs from the way a viewer sees the world—the viewer has the sensation that the object is located either in the space beyond the screen, on the screen plane, or in front of the screen. Key objects that appear to be located on the screen plane are perceived to be in 2D, and are relatively easy to watch. But, over time, a viewer may experience eyestrain from the effort involved in fusing coherent 3D images of objects that appear to reside far beyond or far in front of the screen.

3. Depth and Scale Effects

Along with the limitations noted above, simulated 3D offers some valuable tradeoffs. In addition to the opportunity to use most of the creative effects that are familiar to 2D filmmakers (including color effects, lens distortions, and a wide depth-of-field), 3D filmmakers gain the opportunity to exploit **depth** and **scale** effects that are unique to the 3D medium. By choosing a wide I/O or interocular—the horizontal displacement between their lenses—3D filmmakers are able to lend depth to objects that lie beyond the stereoscopic limit—objects that otherwise would appear to be flat. And by significantly widening or narrowing the I/O, filmmakers can diminish or exaggerate the apparent size of objects—a consequence of the fact that the fixed interocular of the human eye—about 2.5” in the average adult—is the primary frame of reference that viewers use to evaluate the depth and size of objects they see in the world. Scenes shot with a narrow I/O display the effects of *hypostereo*; scenes shot with a wide I/O display the effects of *hyperstereo*; and scenes shot with an average I/O (scenes which most closely replicate the scale of objects in the real world), are said to be in *orthostereo.* Strong degrees of convergence can also introduce significant scale effects. Familiar objects that are placed by the filmmaker far beyond the screen plane may seem gigantic relative to the size of the screen while the same objects, when placed far in front of the screen, may appear miniaturized relative to the size of the screen. In addition, by strongly converging their lenses, 3D filmmakers can cause the spatial geometry of a scene to seem to warp, as if the space were bending toward or away from the viewer. Effects such as these, though unwanted in productions that aim to portray the world as it is, can be useful tools for a filmmaker producing a surrealistic or fantasy film. In addition, since digital 3D images are fused by the viewer from two separate 2D images--each with its own set of noise levels and video scanlines—these images are perceived to possess greater spatial resolution than comparable 2D images. The increment in perceived resolution relative to the same images in 2D is further enhanced by the psychophysical effect upon the viewer of experiencing images in 3D—images that seem to be endowed with a heightened sense of "presence" or "actuality." With its unique capabilities, 3D is not merely a tool for adding verisimilitude to a film—a way to persuade audiences that they are "there." It is, above all, a medium that offers filmmakers an opportunity to communicate with their audiences on a richer, more compelling, and more intimate level than was afforded to them in 2D. When skillfully employed, the artifacts that are peculiar to simulated 3D—instead of being merely clever gimmicks or awkward limitations—can be powerful tools that creative filmmakers can use to more effectively tell dramatic, emotionally-engaging stories.

4. Interocular and Convergence

In addition to the variables to consider when capturing images in 2D (among them, focal length, focus, aperture, shutter angle, imager size, dynamic range, codec, colorspace, and frame rate) two important new options are available to 3D filmmakers (Fig.2): 1) the ability to adjust the *interocular*, or displacement between their lenses, and 2) the ability to adjust the *convergence* of the lenses. The term *convergence*, as it applies in 3D imaging, refers to the rotation, or toeing-in, of the lenses of the two cameras in the 3D rig. The location of the point of convergence—the point where the optical axes of the lenses cross—establishes the position of the virtual object relative to the *screen plane* when the image is displayed

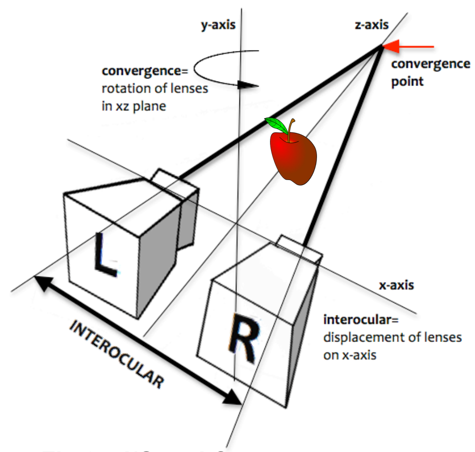

Fig.2 – I/O and Convergence

5. Convergence, Parallax, and Depth

Parallax is the difference (R-L) between the location on the screen plane of an object seen by the right eye (R) and the location on the screen plane of the same object seen by the left eye (L). When the left-eye image of the object is seen on the screen at a position that lies to the right of its right-eye image, the parallax is negative (Fig.3a). When a viewer fuses (converges) this image pair into a coherent 3D object, the object seems to be located *in front of* the screen plane. When the left-eye image of the object is seen on the screen at a position that lies to the left of its right-eye image, the parallax is positive (Fig.3b). When a viewer fuses this image pair into a coherent 3D object, the object seems to be located *beyond* the screen plane.

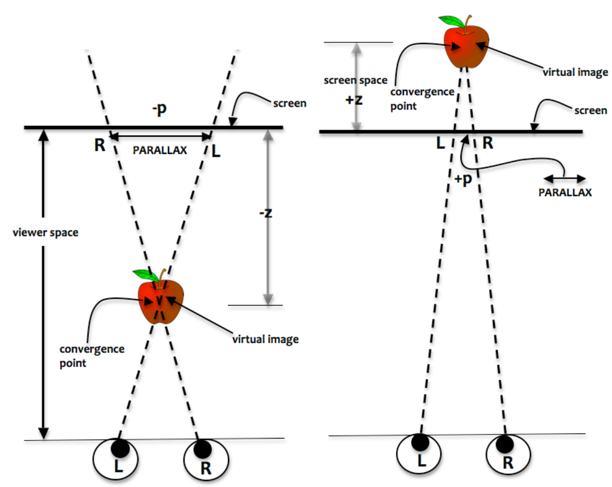

Fig.3a **Fig.3b**

Fig.3 – Convergence, Parallax, and Depth

Fig.3a: the eyes are focused on the screen and the left image is to the right of the right image (negative parallax).
The object seems to be located in the viewer's space (in front of the screen), or in "negative z-space."

Fig.3b: the eyes are focused on the screen and the left image is to the left of the right image (positive parallax).
The object seems to be located in the screen space (beyond the screen), or in "positive z-space."

When the parallax value is zero, the two images of the object overlap at the screen plane. In this circumstance the object is perceived to be in 2D and appears to be located on the screen plane. When a viewer's eyes are focused on the screen and converged on an object that appears to be *in front of* the screen (i.e. an object with negative parallax), the viewer's left and right eyes cross. And when a viewer's eyes are focused on the plane of the screen and converged on an object that appears to be *on* or *beyond* the screen (i.e. an object with zero or positive parallax), the viewer's eyes remain uncrossed. Because crossing the eyes for an extended period of time can cause physical discomfort, filmmakers are often wary about employing strong negative parallax, preferring instead to limit the frequency and degree to which objects are allowed to intrude into the viewer's space[1].

6. Preproduction Decisions

The essential decisions regarding I/O and convergence are often made in preproduction. While errors in convergence decisions can be relatively easy to correct in post, misjudgments in setting the interocular may be impossible to correct. The I/O and convergence decisions for a film may be incorporated into a *depth script*, a plot that tracks the filmmaker's intentions for the emotional impact of the scenes in the film (Fig.4). If a filmmaker wishes the audience to be drawn into the action on screen, the scene may be shot with strong positive parallax (staging the key action beyond the screen plane). But if a filmmaker wishes the audience to experience discomfort, the scene may be shot with strong negative parallax (staging the key action in front of the screen plane). In the same way, emotionally-neutral scenes may be shot with minimal parallax (with the key action staged on or near the screen plane). These *depth decisions* may be precisely defined in preproduction or broadly suggested in the shooting script with marginal notations such a C-5, C-1, C-0, C+1, C+5, etc.—where C-5 refers to a shot with the lenses converged far in front of the target object, thereby pushing the virtual object deep into the screen space; C-0 refers to a shot with the lenses converged on the target object, thereby placing the virtual object on the screen plane; and C+5 refers to a shot with the lenses converged far beyond the target object, thereby pulling the virtual object deep into the viewer's space.

To save time in production and to avoid the keystoning distortions that can result from acutely-converged lenses, 3D filmmakers often choose to shoot unconverged, with the lenses of the cameras parallel and all convergence decisions and adjustments postponed until post.

In terms of I/O, the use of wider-than-normal and narrower-than-normal camera configurations, as noted above, can introduce unwanted distortions of scale. However, the informed use of exaggerated interoculars can help filmmakers enhance the impact of their stories. A wider-than-normal interocular might be used to emphasize the relative insignificance of a character relative to the character's surroundings, while a narrower-than-normal interocular might be used to exaggerate the importance of a character relative to the character's surroundings. As with the depth decisions noted above, these *scale decisions* may be incorporated in the depth script via a plot that tracks the variation in the I/O settings over the running time of the film.

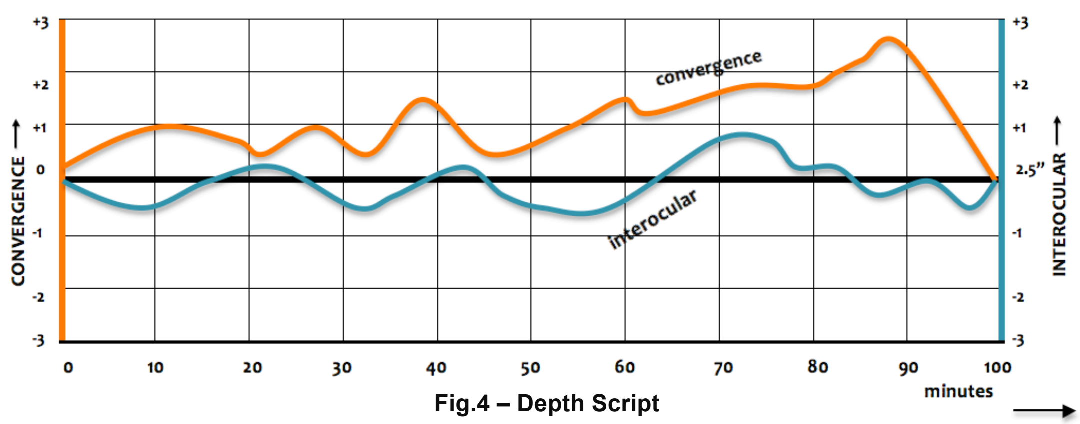

Fig.4 – Depth Script

[1] This caveat applies in particular to productions for TV and to feature-length films shown in typical cinema venues. In giant screen venues such as traditional (film-based) IMAX theaters—with films that are typically 40 minutes in length and screens up to five times the size of those in typical cinemas—negative parallax is considered to be easier to experience than it is in typical viewing spaces. In giant-screen IMAX venues, the exaggerated off-the-screen effects plus the relative loss of awareness of the edges of the screen affords viewers the sense that they are not simply *watching* a film but that they are *immersed* in it. It is generally considered that a feature-length film that relies excessively on negative parallax, whether it is displayed on a conventional or giant-sized screen, is likely to produce uncomfortable eyestrain.

The scale decisions referred to above may be precisely defined in preproduction or broadly suggested in the shooting script with marginal notations such as I-5, I-1, I-0, I+1, I+5, etc.—where I-5 refers to a shot in which the lenses are narrowly separated, making key subjects look larger than life; I-0 refers to a shot in which the lenses are separated by the average human interocular (2.5"), making key subjects look the same size as they do in life; and I+5 refers to a shot in which the lenses are widely separated, making key subjects look smaller than they do in life.

In order to ease the audience into the 3D experience at the beginning of a film and to ease the audience out of the 3D experience at the end of the film, 3D filmmakers may choose to shoot the first and last scenes of the film with convergence and I/O settings that are close to C-0 and I-0. Under these circumstances the key objects in the first and last scenes will be in "soft" 3D (or even 2D) and the objects will appear to have the same proportions relative to other objects in the scene that they have in real life.

7. Production Decisions

The cameras in the 3D rig may be configured to shoot (1) in parallel mode (unconverged); (2) converged beyond the key object of interest in the scene; (3) converged on the key object of interest in the scene; or (4) converged in front of the key object of interest in the scene (Fig.5). The process of setting convergence involves toeing-in the lenses of the cameras so that the right and left images of a target object overlap. The target for setting convergence may be an object in the scene or a slate held in front of, on the plane of, or beyond the plane of a subject.

The convergence point that is chosen by the filmmaker establishes the position of the zero parallax plane in the image. This plane coincides with the screen plane when the film is displayed (Fig.6).

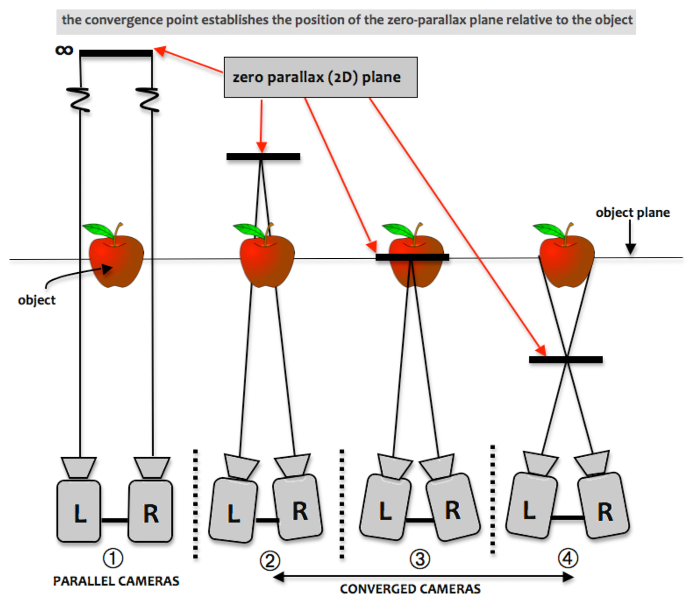

Fig.5 – Image Capture

8. Check As You Go

The avoidance of unwanted artifacts during production is best achieved on location by evaluating shots with a 3D monitor and by screening 3D dailies on a screen that matches the size of the display on which the film will be seen in its target market. As further insurance, simple computer programs permit filmmakers to determine acceptable interocular settings by entering values for parallax, focal length, imager size, subject distance, the average interocular of the target audience, and the width of the target screen.

9. Free Viewing

It is useful for 3D filmmakers—and especially stereographers--to develop the ability to *free view* the images displayed in the viewfinders of their cameras or on a 2D field monitor. The term *free viewing* refers to the ability to view 2D images in 3D without the use of glasses.

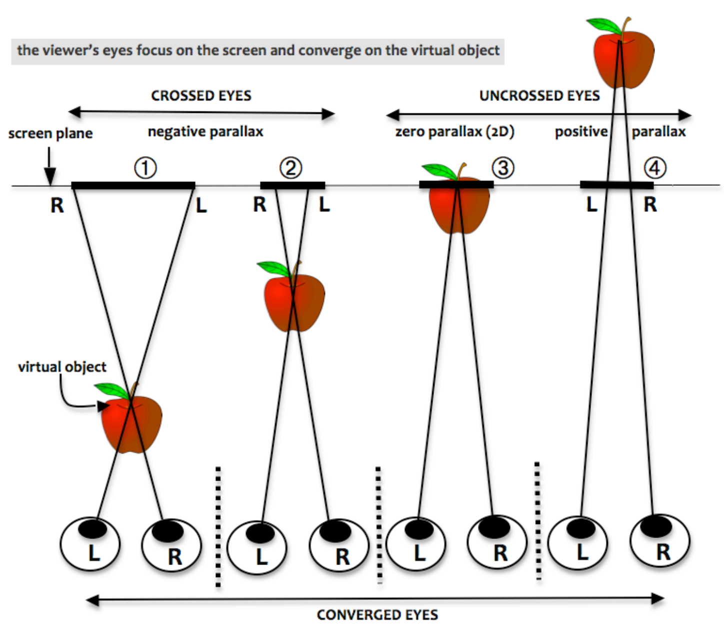

Fig.6 – Image Display

The effect can be achieved by crossing the eyes while viewing any 2D still or moving image—thereby converging the line-of-sight of the two eyes at a point that lies between the viewer and the image (Fig.7). For those who have difficulty crossing their eyes, the effect can be replicated, at the expense of image brightness, by viewing the 2D image while half-closing one eye, leaving the other eye open. In this case, the line-of-sight of the half-closed eye is deflected inward, to converge with the line-of-sight of the open eye at a point in front of the image. The result is the establishment of a "virtual screen plane" at the point of convergence of the eyes, with the subject seemingly located in positive parallax beyond this plane.

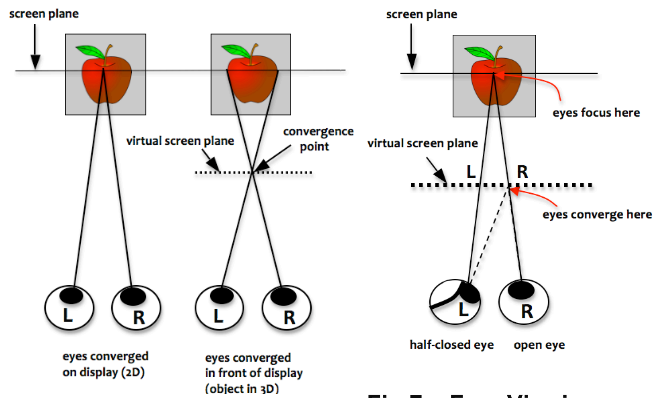

Fig.7 – Free Viewing

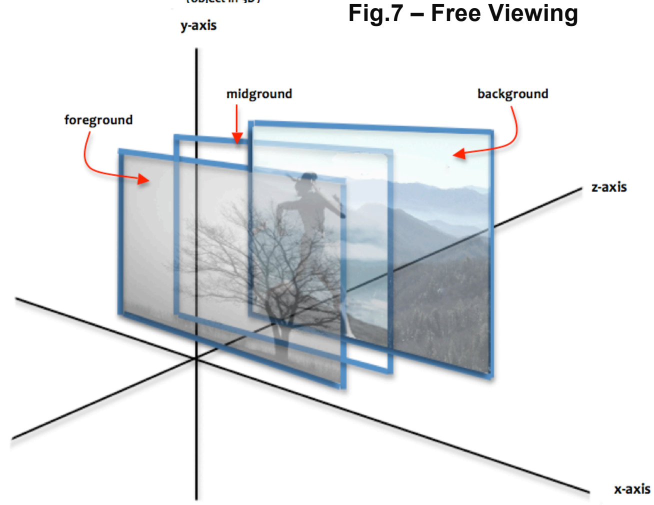

Fig.8 – Free Viewing Depth Planes

While cross-eyed free viewing may cause eyestrain if practiced for extended periods, it can be a useful technique for quickly evaluating the potential effects of the conversion of a 2D image to 3D in post. Among the potentially annoying artifacts that can be readily flagged are busy, soft-focus backgrounds that viewers—especially when watching a large 3D display—may struggle to try to resolve. In addition the technique can detect objects that will appear to be miniaturized in 3D—a consequence of the use of wide-angle lenses or the dwarfing of familiar objects in wide panoramic views. Fig.7 illustrates the convergence of the sight lines of the eyes during conventional viewing; during free viewing with both eyes crossed; and during free viewing with one eye half-open, causing its sight line to cross that of the open eye. The technique—unlike the Pulfrich Effect, which depends upon lateral motion on the screen--works equally well for still and moving images. The process depends on the decoupling of the act of converging and focusing the eyes, with the brain relying on monoscopic depth cues that are present in the image (among them focus cues, atmospheric cues, and occlusion clues) in order to stratify the 2D image into a series of planes that correspond to increasing degrees of depth (Fig.8).

10. Post Processing

If decisions regarding convergence are deferred until post, the images from the left and right eyes may be converged (shifted on the x-axis) by the process called *depth grading* or *horizontal image translation* (HIT). This technique involves horizontally displacing the right and left images to produce negative, zero, or positive parallax on the screen plane (see Fig.9 on the next page). Because of the potential need to enlarge and to crop the image during the depth grading process, filmmakers who are shooting with their lenses parallel (unconverged) generally choose to shoot in overscan mode, leaving a buffer around the sides of the frame.

11. Convergence Considerations

As noted above, strong convergence can produce unwanted warping effects (including keystoning effects, which result from vertical disparities at the edge of the right and left images), as well as scaling effects—a consequence of placing familiar objects far beyond or far in front of the screen plane. In addition, images that are strongly converged (images with high parallax values) may display more ghosting when displayed on a screen—a consequence of the unavoidable "leakage" of light between the two lenses of the 3D viewing glasses. However, converging the cameras has the mitigating effect of centering the screen image on the optical axes of the lenses, where any radial distortion in the lenses is likely to be at a minimum. In addition, by converging the cameras, a filmmaker may be able to include a larger area of the set within the scene, increasing the width of the 3D zone in the space closest to the camera.

While filmmakers are free to converge their lenses wherever they wish, audience tests suggest that a *comfort zone* exists within which parallax values are within an comfortable range. The depth of this zone is inversely proportional to the width of the viewing screen: a movie screen that is 30’ wide can afford a comfort zone that is only 10% of the depth of the comfort zone that is considered appropriate for a TV screen that is 3’ wide. Factoring in the average audience viewing distance from a 30’ movie screen, this comfort zone is considered to extend from about 30’ in front of the screen (where extreme effects may intrude into the viewer’s space) to about 30’ beyond the screen (where the furthest 3D objects reside in screen space). Stated another way, these tests have shown that negative parallax should extend into the viewer’s space a distance that is no more than 50% of the distance between the viewer and the screen. And the extreme limit of positive parallax (beyond the screen) should be a comparable distance. A larger negative parallax value may force viewers to cross their eyes to an uncomfortable degree, and a larger positive parallax value may force viewers to diverge their eyes to an uncomfortable degree—the so-called *wall-eye* effect. Because the depth of the comfort zone is dependent upon screen size, 3D films designed for distribution to multiple platforms need to be subjected to convergence adjustment (depth-grading) passes that tailor their parallax values to the size of the screens on which they will most likely be displayed (Fig.9).

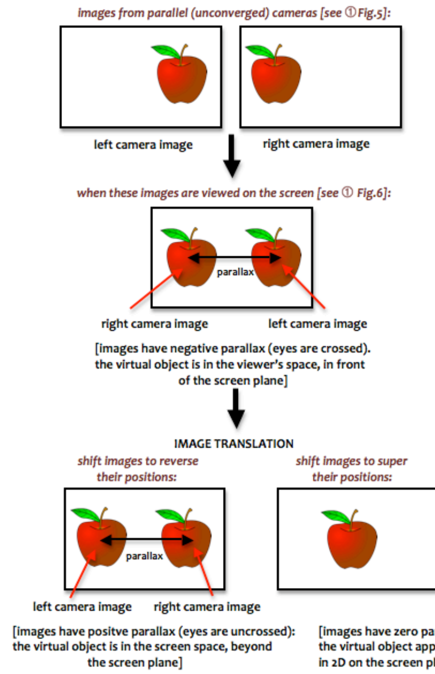

Fig.9 – Depth Grading

The range of parallax values employed in a film is termed the film’s *depth budget*. Typically, a filmmaker making an action or suspense film may choose to employ a larger depth budget than would be appropriate for a conventional drama, which may instead confine much of the on-screen action to depths that are close to the screen plane. In the case of a production designed for a 30’ screen, the separation of the right and left images (the parallax) may range from a negative extreme of -10” to a positive extreme of +2.5”, for an overall depth budget of 12.5” (3% of the width of the screen), while for a comparable production designed for display on a 3’ TV monitor, the appropriate depth budget might be 2” (6% of the width of the screen). Because of this scale-dependency, a production whose convergence is appropriate for a theater screen may lose much of its apparent depth when viewed on a small screen, while a production whose convergence is considered to be appropriate for a small screen may force the viewers to diverge their eyes to an uncomfortable degree when they watch the same production on a large screen. While efforts are underway to standardize 3D specifications to conform to the needs of the various projection and display systems in the marketplace (each of which uses light and the color spectrum in a distinctive way), the prospect of standardizing specifications to meet the needs of the vast range of possible *display geometries* – from IMAX screens to personal computers – remains a distant hope.

12. Maximizing the Illusion of Depth

When viewing 3D films, audiences accustomed to scanning the real 3D world tend to explore the entire contents of the frame, believing they are free to focus and converge their eyes wherever they choose. As a result, viewers may experience eyestrain when they try, without success, to resolve objects on a large screen that are out of focus—objects that are beyond the lenses’ depth-of-field. For this reason 3D filmmakers, when creating for bigger screens, generally opt to use lenses with a wider depth-of-field, lenses that, in the case of 1” camera imagers, have a focal length of around 50 mm[2]. Not only do longer lenses shrink the depth-of-field, but they foreshorten perspective. And

[2]extreme wide-angle lenses are generally avoided in 3D productions that seek to achieve naturalistic effects. Though such lenses offer the advantage of a wide depth-of–field, the miniaturization effect they introduce is dramatically exaggerated in 3D. And once this effect is baked-into the image, it is--like the effect of a wide interocular—virtually impossible to correct in post. While cinema audiences have grown accustomed to the effects of gigantism (such as those conveyed by an ECU of an actor), miniaturization effects (such as those produced by a wide I/O or by wide-angle lenses) are generally perceived to be unnatural. Because of this, a useful rule-of-thumb for filmmakers aiming for lifelike 3D is to avoid both wide interoculars and extreme wide-angle lenses.

because of this they can produce potentially disturbing *cardboarding* effects—the illusion that objects in the scene are merely cutouts, lacking modeling or depth. This artifact can be counteracted, to a degree, by choosing a wider interocular, thereby restoring some of the depth to the scene. But, as noted above, choosing this option risks introducing unwanted miniaturization effects. Unless, for some reason, a filmmaker specifically wishes to use artifacts such as soft focus, cardboarding, and miniaturization as storytelling devices, it is generally considered advisable to avoid long focal-length lenses in 3D productions. This restriction on the choice of lenses unfortunately deters most 3D filmmakers from using one of the most effective of 2D filmmaking tools—the application of selective focus as a way to direct the audience's attention to the subjects the filmmaker considers to be of primary importance on the screen. Because of this, 3D filmmakers must pay renewed attention to the power of light, color, and composition to draw attention to the subjects in the scene that they want the audience to see. For example, a street that winds away into the distance can help to set the actors apart from the background, or the use of smoke, dim lighting, or cooler hues can help to direct the audience's attention to foreground subjects that are more clearly defined, more brightly lighted, or shaded in warmer tones.

In addition, to take advantage of the unique capabilities of 3D, scenes may be composed in a way that emphasizes the depth of the set. In this regard, a set that includes multiple layers of depth may be preferable to one that includes only a few layers of depth. As an example, while a 2D cinematographer might stage a conversation between two actors as they stand against a wall, a 3D filmmaker, to take full advantage of the medium, might prefer to stage the same scene with the actors conversing as they stroll down a street, passing other pedestrians and moving through multiple layers of depth. As with any medium, however, the ability to add depth to a scene does not obligate a filmmaker to use the full complement of depth-enhancing tools. The 3D filmmaker is free to capture flattened images, or distorted images, if such images most effectively express the intent of a scene.

13. Managing Motion

Because 3D images require more time to scan and fuse (or *read*) than 2D images, viewers may be frustrated by overly-rapid camera and/or subject motion. To avoid this reading lag, 3D filmmakers often choose to slow the pace of the camera movement and to stage on-screen action either at a slower pace or at a diagonal relative to the camera—a strategy that not only allows more time for the viewer to fuse the pair of images but one that also minimizes the appearance of the dreaded 24 fps flicker, an artifact that is exaggerated in 3D. In the same way, because 3D is intrinsically more subjective, more interactive, and more *explorative* than 2D, filmmakers working in 3D generally opt to keep their cameras in motion, allowing the cameras to travel slowly and smoothly through successive layers of depth, relying less upon static shots and instead upon carefully-planned camera choreography to engage the audience in the story.

14. Watching the Edges

When composing films for smaller screens, filmmakers must take into account the prospect that objects with strong negative parallax (objects that intrude deeply into the viewer's space) may be truncated at the right or left sides of the screen, confusing the signals that viewers require in order to accurately place the objects in space. To correct this *edge violation* effect, movable mattes (or *floating windows*) may be applied in post, temporarily masking objects that reside for more than a second or two at the right and left edges of the frame. But, to avoid edge violations, 3D filmmakers often choose to compose their scenes so that objects with strong negative parallax are not allowed to linger for long at the edges of the frame.

15. Minimizing Disparities

To achieve effective 3D, every possible disparity between the images captured by the right and left lenses of the cameras must be minimized—with the obvious exception of the distinct point-of-view of the two lenses. Accomplishing this goal requires maintaining precise synchronization between the cameras, in addition to achieving a precise match between the optics of the lenses and a precise match between the imaging systems in the cameras. Mismatches in sync or in optics, as well as stray reflections, light flares, or stray objects in the frame—essentially any artifact introduced during production or post that affects one "eye" but not the other—can introduce *retinal rivalry*—a conflict between the components of the stereo pair that disrupts the web of illusions that produce effective 3D.

16. Filming Live Events

When shooting live events such as news stories, stage performances, natural history scenes, and spectator sports, 3D filmmakers confront several limitations that can compromise their ability to capture effective 3D. Foremost among these are the need to make all of the decisions regarding I/O and convergence in real time, plus the restriction upon the placement and motion of the cameras. The first limitation can be addressed by a stereographer who has access to 3D feeds from the cameras and can relay I/O and convergence decisions to camera assistants who can adjust the I/O and convergence of the cameras. on the fly. But since effective 3D places the viewer *inside* the action of the film--as an active participant and not merely as a passive spectator--the second limitation is more difficult to address. While live performances generally permit cameras to have a limited presence on the stage, news stories, spectator sports, and natural events typically place strict limits on the space in which cameras can operate. And unfortunately, this *camera space* is outside the *event space*, the locus of the key action that the filmmaker wishes to capture (Fig.10).

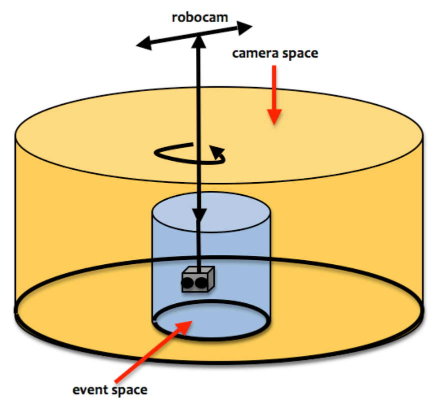

Fig.10 – Camera Space and Event Space

As a result of these limitations, 3D filmmakers who wish to effectively capture live events may need to depend upon the deployment of robocams—remote-controlled 3D cameras that can unobtrusively dip in and out of the subject space without interfering with the movement of the performers, subjects, or players while, at the same time, managing to avoid obstructing the views of any spectators. In addition to these considerations, 3D filmmakers who shoot live events are unlikely to want the distorted depth and scale artifacts that may find a useful place in the palette of the dramatic filmmaker. In particular, by shooting with a very wide I/O or with wide-angle lenses, live event filmmakers may inadvertently capture images that suffer from miniaturization. In this regard, it is often advisable for live event filmmakers to shoot most of the action in *orthostereographic* mode, with their cameras mounted in side-by-side configuration, with the I/O of their lenses set at 2.5," and their lenses converged a short distance in front of or behind the key subject of interest. An exception to this rule applies to establishing shots—wide vistas from a vantage point at the back of an arena, high in the stands, high on a hill, or high in the air. In this case the stereographer may opt for a slightly wider-than-normal I/O in order to lend some depth to the performers on the stage, the players on the field, or the wildlife in the scene—while taking care to avoid the miniaturization effect that results from a wide-angle lens or an overly-wide I/O. Along with the above, live event filmmakers, like all other 3D filmmakers, are generally advised to avoid long focal length lenses in favor of those with a nearly-normal depth-of-field.

17. Pacing the Cut

The caveats that apply to 3D production also apply to the art of composing 3D films in post. In addition to carefully matching convergence and interocular distances from shot to shot in order to maintain continuity of depth and scale over the running time of the film, experienced 3D picture editors tend to make allowances for an extended audience reading time—especially in the case of adults. Until their audiences become accustomed to rapidly reading 3D images, editors working on films that are intended for general audiences are advised to play scenes for durations that are up to two times longer than they normally would play in 2D.

18. Weighing the Cost

As more films are produced in 3D, filmmakers will gain a clearer perspective on the impact in time and dollars of shooting in 3D relative to 2D. In general, however, it seems safe to estimate that a 3D film will take 30% longer to shoot and 50% longer to post than a comparable 2D film and will cost 30% to 50% more overall. Ultimately, the decision to shoot in 3D instead of 2D, while subject to creative considerations, must depend upon the ability of the marketplace to return that incremental investment.

19. Terminology

parallax – the separation on the screen plane between the right and left images of an object. The degree of parallax determines the perceived depth of objects within the viewer's space or the screen space.

negative parallax – objects are perceived to be positioned within the viewer's space (or "personal space"), i.e. in front of the screen plane.

positive parallax – objects are perceived to be positioned within the screen space (or "world space"), i.e. beyond the screen plane.

zero parallax – objects are positioned at the screen plane and appear to be in two dimensions.

depth budget – the combined values of maximum positive and maximum negative parallax in a production.

screen plane – the plane of the display (the zero-parallax plane): the plane that corresponds to the surface of the movie screen, TV screen, or computer screen.

convergence – (or "toeing-in"): the inward rotation of the lenses to shift the parallax of the objects in the scene and the apparent proportions of objects relative to the screen.

point of convergence – the position where the optical axes of the lenses intersect, defining the position of the zero parallax plane or screen plane.

divergence – the unnatural outward rotation of the human eyes to view images with extreme positive parallax, resulting in an uncomfortable "wall-eye" effect.

interocular – (or "interaxial"): the horizontal displacement of the lenses of the cameras.

hyperstereo – the effect of an interocular that is larger than that of the average human eye (2.5").

miniaturization – an artifact that results from use of wide-angle lenses or a larger interocular than that of the average human eye (hyperstereo).

hypostereo – the effect of an interocular that is smaller than that of the average human eye (2.5").

gigantism – an artifact that results from use of a smaller interocular than that of the average human eye (hypostereo).

orthostereo – the effect of shooting with an interocular that approximates that of the average human eye.

cardboarding – an artifact that can result from the use of long focal length lenses.

keystoning – (or "trapezoidal distortion"): an artifact that results from excessive convergence of the lenses.

wall-eye – an uncomfortable condition that results from the attempt to fuse objects with strong positive parallax and a wider-than-normal interocular.

free-viewing – the technique of viewing a 2D image in 3D by crossing the eyes in order to cause the eyes to converge in front of the 2D image.